Notes on behavioural economics

Jason Collins

20 May 2025

Table of contents

Preface		2
Ι	Economic foundations	5
1	Preferences	7
2	Rationality	9
3	Utility	14
4	Economic foundation exercises	16
II	Decision making under risk	20
5	Notation and mathematical background	22
6	Expected value	29
7	Axioms for Expected Utility Theory	32
8	Expected Utility Theory	44
9	Expected utility examples	54
10	Anomalies in expected utility theory	65
11	Decision making under risk exercises	80
II	I Prospect theory	93
12	Reference dependence	95

13 Loss aversion	98
14 The reflection effect	103
15 The value function	109
16 Probability weighting	113
17 Prospect theory implementation	117
18 Prospect theory examples	120
19 Prospect theory applications	147
20 Prospect theory exercises	152
IV Intertemporal choice	167
21 Exponential discounting	170
22 Exponential discounting examples	174
23 Exponential discounting anomalies	185
24 Present bias	188
25 Present bias examples	193
26 Sophisticated present bias	202
27 Commitment	215
28 Delayed gratification, spread and variation	227
29 Intertemporal choice applications	229
30 Intertemporal choice exercises	237
V Probability foundations	256
31 Probability theory	258
32 Bayes' rule	265
33 Probability theory exercises	273
34 Subjective expected utility theory	284

VI Heuristics and biases	289
35 Heuristics	292
36 Biases in probability judgment	297
37 Heuristics and the bias-variance trade-off	315
38 Overconfidence	320
39 Heuristics and biases exercises	324
VII Game theory	334
40 Simultaneous-move one-shot games	337
41 Sequential games	347
42 Asymmetric information	356
43 Strategic moves and commitment	361
44 Game theory exercises	366
VIII Behavioural game theory	376
45 Level-k thinking	378
46 Asymmetric information and the curse of knowled	lge 385
47 Emotions	
48 Behavioural game theory exercises	393
IX Social preferences	404
49 People care about resource distribution	409
50 Reputation	425
51 Reciprocity	428
52 Social preferences exercises	433

X Practice problems	437
53 Expected utility theory problems	439
54 Prospect theory problems	451
55 Intertemporal choice problems	457
56 Beliefs problems	464
57 Behavioural game theory problems	469
References	480

Preface

Welcome to these notes on behavioural economics.

These are the notes for the UTS undergraduate subject 23005 Behavioural Economics.

I take a traditional approach. I start with the basic economic and game theory foundations. I then examine what happens when we introduce a richer view of human behaviour. I look at how we make decisions under risk and uncertainty, and over time, how we judge probability and how we interact with others.

The result is a new set of predictions as to how humans might behave and what economic phenomena might emerge.

There are no mathematical prerequisites for this subject. Hence, the mathematics is kept at a basic level.

I have created videos to accompany these notes. You can find links to them and other teaching materials on my website at www.jasoncollins.blog/teaching/.

These notes cover the following areas:

- **Economic foundations**: The core economic concepts that we will use and modify in exploring behavioural economics.
- Decision making under risk and uncertainty: The traditional economic approach to decision making under risk and uncertainty, and some empirical anomalies that arise when using this approach.
- **Prospect theory**: The pre-eminent alternative to expected utility theory, prospect theory, with examples and possible applications.
- Inter-temporal choice: Decision making involving costs and benefits occurring at different times. I look at two types of time preference: exponential discounting and quasi-hyperbolic discounting.
- **Beliefs**: Some core concepts on beliefs and probability judgment and the empirical evidence on how humans do not adhere to these foundational principles.
- **Game theory**: The game theoretic concepts that we use in our analysis of behavioural game theory and social preferences.

- Behavioural game theory: How heterogeneous agents and bounded rationality can change our analysis of economic games.
- **Social preferences**: How agents might incorporate the outcomes and beliefs of others into their decisions.

What is behavioural economics?

Behavioural economics is a discipline that seeks to increase the explanatory power of traditional economic approaches by incorporating more realistic psychological foundations.

Behavioural economics retains the general framework and tools used by economists, but deviates from some assumptions to generate new insights and better predictions. These deviations are generally grounded in observed behaviour rather than abstract principles.

What is not behavioural economics?

You have probably heard the term behavioural economics in the popular press and books. Many of those references are not what I would consider to be behavioural economics. As a result, it is worth identifying what is not behavioural economics.

The word economics is the key. Economics is the study of how economic agents make decisions under conditions of scarcity and the study of the interactions between those agents.

As I noted, behavioural economics involves the introduction of more realistic psychological foundations to that economic approach.

Behavioural economics is not the general study of human behaviour. Behavioural science is a better term for that general study.

Similarly, psychology is the study of the human mind, decisions and behaviour. Psychology is part of the behavioural sciences. Behavioural economics draws on a small subset of psychology to develop a better understanding of human behaviour.

Three thought experiments

Many early ideas in behavioural economics emerged from thought experiments concerning human behaviour that was hard to explain with traditional economic frameworks.

Here are three thought experiments to illustrate the types of problems that concern behavioural economists. In this course I will examine potential explanations for each of these behaviours.

Thought experiment 1

This first thought experiment comes from Kahneman and Tversky (1984).

A. Imagine that you have decided to see a play and paid the admission price of \$100 per ticket. As you enter the theatre, you discover that you have lost the ticket. The seat was not marked, and the ticket cannot be recovered.

Would you pay \$100 for another ticket?

B. Imagine that you have decided to see a play where admission is \$100 per ticket. As you enter the theatre, you discover that you have lost a \$100 bill.

Would you still pay \$100 for a ticket for the play?

The two situations appear equivalent, yet people are more likely to pay \$100 for a ticket in scenario B.

Thought experiment 2

This second thought experiment comes from R. Thaler (1980).

Mr. R bought a case of good wine ... for about \$5 a bottle. A few years later his wine merchant offered to buy the wine back for \$100 a bottle. He refused, although he has never paid more than \$35 for a bottle of wine.

Why is there such a large difference between the price at which he is willing to buy and the price at which he is willing to sell?

Thought experiment 3

This third thought experiment also comes from R. Thaler (1980).

A group of hungry economists is awaiting dinner when a large can of cashews is opened and placed on the coffee table. After half the can is devoured in three minutes, everyone agrees to put the rest of the cashews into the pantry.

Why did they agree to remove the cashews when they simply could have not eaten them? Why did they deliberately reduce their choice set?

Part I

Economic foundations

Behavioural economics builds on the frameworks and tools used by economists. In this part, I introduce some of the core economic concepts that we will build on or explicitly deviate from in our exploration of behavioural economics.

I will first examine the preference relation used by economists.

Then I will discuss how economists define rationality, and the foundational axioms of rationality, completeness and transitivity.

Finally, I will describe how the preference relation and axioms form the basis of utility functions.

Preferences

Summary

- Economists use preference relations to capture how agents order their choices.
- The strong preference relation (≻) indicates one option is "better than" another.
- The weak preference relation (≽) indicates one option is "at least as good as" another.
- The indifference relation (~) indicates two options are "as good as" each other. Indifference is linked to weak preference: x ~ y if and only if x ≽ y and y ≽ x.

Economists use the preference relation to capture the ordering that an agent gives to options between which they might choose. There are three forms of the preference relation.

The first is the strong (or strict) preference relation. We use the strong preference relation to indicate that an agent considers one option "better than" another. We represent strong preference with the symbol \succ .

The second is the weak preference relation. We use the weak preference relation to indicate that an agent considers one option "at least as good as" another. We represent weak preference with the symbol \succeq .

The third is indifference. Indifference occurs when an agent considers that each option is "as good as" the other or that the agent is "indifferent between" the options. We represent indifference with the symbol \sim .

Here are a couple of illustrations of these relations.

Suppose I strongly prefer bananas to apples. I would write:

• bananas \succ apples

Bananas are better than apples.

Similarly, if I am indifferent between bananas and oranges, I would write:

• bananas \sim oranges

Bananas are as good as oranges. Oranges are as good as bananas.

There is an important link between indifference and the weak preference relation. I am indifferent between x and y ($x \sim y$) if and only if I weakly prefer x to y ($x \succeq y$) and I weakly prefer y to x ($y \succeq x$). The weak preference relation includes the possibility of indifference.

In that case of the example of my indifference between oranges and bananas, I could also say that I weakly prefer bananas to oranges and that I weakly prefer oranges to bananas:

• bananas \succ oranges and oranges \succ bananas

Rationality

Summary

- In economics, rationality means that preferences respect certain principles, not that decisions are necessarily reasonable or self-interested.
- The completeness axiom states that an agent can always compare any two options: for all x and y, either $x \succeq y$ or $y \succeq x$ (or both).
- The transitivity axiom states that if an agent prefers A to B and B to C, they will prefer A to C: if $x \succcurlyeq y$ and $y \succcurlyeq z$, then $x \succcurlyeq z$.
- These axioms allow for the construction of a preference ordering, but they provide minimal constraints on the nature of preferences themselves.

2.1 Introduction

A standard economic assumption is that decision makers are rational.

However, rationality in economics has a different definition to the 'lay' definition of rationality.

Rationality simply means that preferences respect some desirable principles. These principles are assumptions, not rules of behaviour.

Economists tend to keep the number of assumptions as small as possible. They choose the assumptions that they need for the particular analytical problem they have at hand.

In its most stripped-back form, analysis of consumer choice rests on just two assumptions. These are completeness and transitivity.

2.2 Completeness

Completeness means that an agent can always compare any two options. The agent cannot fail to have a preference between two options (although that preference may be indifference).

For example, if an agent was presented with a choice between an apple and a banana, or a choice between a Mercedes and a BMW, they will always strictly prefer one of them or be indifferent between the two. They will never not know what to choose or be unable to make a choice. They cannot be indecisive.

Formally, we can state the completeness axiom as follows:

For all x and y, either $x \succeq y$ or $y \succeq x$ (or both).

Completeness means people always prefer x or y, or are indifferent between the two.

While the completeness axiom appears sensible, it is possible to develop examples where it may not hold. Consider a choice between two possible holiday destinations. Or two potential dates or spouses. If you are torn between the options and unable to make up your mind, this would represent a breach of the completeness axiom.

Incomplete preferences are different from indifference. If you are indifferent, you would be able to decide by, say, flipping a coin. You will be equally satisfied whatever the outcome. Incompleteness makes choice impossible.

2.3 Transitivity

Under transitivity, if a person prefers A to B and B to C, they will prefer A to C.

Formally, we can state the transitivity axiom as follows:

For all x, y and z, if $x \succcurlyeq y$ and $y \succcurlyeq z$, then $x \succcurlyeq z$.

One classic argument for transitivity is the concept of a "money pump" (Davidson et al. (1955)).

Suppose you have a person who prefers A to B, B to C and C to A. That is, they have intransitive preferences. They have \$20 and an endowment of C.



They are offered B in exchange for their endowment of C for some small nominal cost (say \$1). If they make the trade they now have B and \$19.



They are then offered A in exchange for their endowment of B, again for a nominal cost.



Finally, they are offered C in exchange for their endowment of A for a further nominal cost. They now have an endowment of C and \$17. This process can be repeated until the agent has no money.



2.4 Preference Orderings

When preferences are transitive and complete, we can create a ranking of options from most to least preferred. We can construct a preference ordering.

Completeness guarantees that there will be only one ordering. There are no gaps or ambiguous comparisons. If you have three options (A, B and C), you must be able to rank them fully, such as A is better than B is better than C, or I am indifferent between A and B which are both better than C.

Transitivity guarantees that there will be no cycles in strict preference. Weak preferences can cycle, so that one can prefer A to B and B to C and C to A, but this entails indifference.

That preference ordering, a simple rank of which options an agent prefers, is all that is required for some analysis of consumer choice.

2.5 Defining rationality

This definition of rationality, accordance to the axioms of completeness and transitivity, provides little constraint on preferences, while allowing for diverse preferences and behaviours.

You could prefer less money to more. You could prefer sums of money divisible by seven with no remainder. You could prefer more for yourself or more for someone else.

These axioms do not lead to an assumption that people are selfish, unless you define selfishness to be simply acting in accordance with their preferences.

These axioms place some constraints on behaviour, and empirical evidence suggests those constraints are sometimes breached by decision makers. But they are constraints that allow much behaviour to be described as rational.

2.6 Independence of irrelevant alternatives

While completeness and transitivity are the most basic axioms of rationality, there are others that can be added to provide more structure to preferences. One such axiom is the independence of irrelevant alternatives.

The axiom of independence of irrelevant alternatives states that if an agent prefers x to y, the introduction of a third option z should not change the preference order between x and y. A new option shouldn't flip how you rank two existing options. For example, if you select fish rather than chicken from a restaurant menu, being told by the waiter that there is a vegetarian option should not lead you to change your selection to chicken.

Mathematically, we can state that:

If $x \succ y$ in choice set $S = \{x, y\}$, then:

 $x \succ y$ in choice set $T = \{x, y, z\}$

The axiom has important implications:

- It allows preferences to be context-independent. An agent's choice between two options can be analysed without needing to know the full choice set.
- It simplifies economic models by making preference orderings more stable and predictable. It eliminates the need to consider complex interactions between irrelevant and relevant options.
- It prevents the outcomes of decision-making processes from being manipulated by the strategic inclusion or exclusion of alternatives.

The axiom of independence of irrelevant alternatives is used in utility theory, mechanism design, and social choice theory, where it helps establish important results about preference aggregation and voting systems. For example, in voting systems, the axiom prevents a candidate's elimination from changing the relative ranking of other candidates.

Utility

Summary

- Economists use utility functions to represent strength of preference, assigning numbers to different options.
- Formally, a utility function u(·) maps alternatives to real numbers, assigning larger numbers to preferred alternatives.
- The preference relation can be expressed using the utility function:

 $x \succcurlyeq y \Leftrightarrow u(x) \ge u(y) x \succ y \Leftrightarrow u(x) > u(y) x \sim y \Leftrightarrow u(x) = u(y)$

Economists often use numbers to represent strength of preference. This is done through utility functions.

A utility function associates a number with each member of the universe. For example:

- Banana: 3
- Orange: 2
- Apple: 1

This does not mean that I rate bananas three times higher than apples. It simply means that I prefer bananas to apples. This utility scale is ordinal, not cardinal. The following is equivalent:

- Banana: 300
- Orange: 2
- Apple: 1

Formally, the utility function $u(\cdot)$:

- maps the set of alternatives into the set of real numbers
- assigns larger numbers to preferred alternatives.

For example, we might write:

$$\begin{aligned} u(\text{banana}) &= 3\\ u(\text{orange}) &= 2\\ u(\text{apple}) &= 1 \end{aligned}$$

The rank of those numbers gives us the preference relation:

$$\begin{split} x \succcurlyeq y & \Longleftrightarrow u(x) \ge u(y) \\ x \succ y & \Longleftrightarrow u(x) > u(y) \\ x \sim y & \Longleftrightarrow u(x) = u(y) \end{split}$$

Again, following from the above:

$$u(\text{banana}) = 3 > 2 = u(\text{orange}) \Longleftrightarrow \text{banana} \succ \text{orange}$$

This calculation of utility is not how the mind actually works. But under the axioms of completeness and transitivity, the consumer behaves as if they have a utility function $u(x_i)$ over outcomes x_i .

Economic foundation exercises

4.1 Buridan's ass

The paradox of Buridan's ass runs as follows: An ass that is equally hungry and thirsty is placed halfway between a pile of hay and a bucket of water. The ass cannot decide between the hay and water, so dies of dehydration and starvation.

What axioms of choice are relevant to this fable? Are any axioms violated?

? Answer

The axiom of completeness is violated. Under this axiom an agent cannot fail to have a preference between two options (although that preference may be indifference).

Incompleteness is different from indifference. Is the ass was merely indifferent it would be happy taking either option and be equally satisfied. Indifference does not make a choice impossible.

4.2 Picking a mobile plan

You are considering two mobile phone plans. Each has different monthly fees, data caps, excess data charges, international inclusions and 5G coverage. You realise it will take all day to work through the fine print to understand the plans.

You decide that your options are:

- Pick a plan by flipping a coin.
- Spend the day working through the plans and choose one.

• Avoid the work by reading a book instead.

Does this decision accord with the axioms we have discussed?

? Answer

All three choices could be considered to accord with the axioms we have discussed. What we have effectively done in each case is created a richer choice set. In no case are their preferences incomplete. You could think of the choice set as {Plan A, Plan B, invest to understand plans, do something else}.

You would only state that the preferences are incomplete if the agent wasn't able to express a preference between Plan A and Plan B. However, if they were forced to choose and happy to flip a coin, that would suggest indifference.

4.3 Rationality

Consider the following statements about rationality in economics. How do these criticisms relate to the definition of rationality we have discussed?

a) From Robert Frank in the New York Times:

TRADITIONAL economic models assume that people are selfinterested in the narrow sense. If "homo economicus" - the stereotypical rational actor in these models - finds a wallet on the sidewalk, he keeps the cash inside. He doesn't leave tips after dining in restaurants that he will never visit again. And he would never vote in a presidential election, much less make an anonymous donation of money or time to a presidential campaign.

💡 Answer

The traditional economic axioms assume self-interest in a narrow sense, in that people make decisions in accordance with their preferences. However, completeness and transitivity say nothing about the content of those preferences. A person might prefer to return the wallet or leave a large tip for the good feeling they get. They might enjoy voting and care about outcomes for others.

Even auxiliary axioms such as monotonicity or non-satiation leave these possibilities open in that while the agent will always want more, they do not require that it is purely for their own benefit.

b) From Brian Easton in interest.co.nz:

For the last 150 years much economic analysis has been based on homo economicus, an 'economic' man who is rational and narrowly self-interested and who pursues his subjectively defined ends optimally.

? Answer

First, "man who is rational" accords with the technical definition in economics but differs from how used in common speech (and likely that of readers of the article).

If we read "narrowly self-interested" to mean that they make decisions in accordance with their preferences, then we might agree with that statement. However, we cannot place any further content into that idea of self-interest.

Similarly, the statement "pursues his subjectively defined ends optimally" accords with the idea that under the axioms of completeness and transitivity, the consumer behaves as if they have a utility function $U(x_i)$ over outcomes x_i . They are able to choose between any two options. They also do not make errors (although randomness can be built into utility functions). This definition of "optimally" is narrower than might be used in common speech.

4.4 Apples and oranges

Elmer claims he has no preference between apples and oranges and is genuinely indifferent between the two. Is this stance consistent with the completeness axiom? Explain why or why not. (2 marks) [200 word limit.]

? Answer

Elmer's stance of being indifferent between apples and oranges is consistent with the completeness axiom.

The completeness axiom states that for any two options A and B, an individual can always specify exactly one of these three possibilities: A is preferred to B, B is preferred to A, or the individual is indifferent between A and B.

Elmer's claim of indifference between apples and oranges satisfies this axiom. He has a clear preference relation between the two options (indifference), which is one of the three possibilities allowed by the completeness axiom.

4.5 Small or medium coffee?

Sam's local cafe offers two sizes of coffee: small and regular. Each day, Sam orders a small coffee.

The cafe then introduces a new large size. Sam starts ordering the regular.

What axiom is Sam violating? Explain. (2 marks) [200 word limit]

? Answer

Sam's action is inconsistent with the independence of irrelevant alternatives.

Under the independence of irrelevant alternatives, if A is preferred to B, then this preference should remain unchanged in the presence of a third option C.

Initially, Sam prefers small to regular. After the introduction of the large size, Sam switches to preferring regular to small. This change in preference between small and regular, triggered solely by the presence of the large option (which Sam doesn't even choose), contradicts the axiom.

This scenario does not involve probabilities, as the independence axiom is defined, so we can't evaluate it in terms of the independence axiom.

Part II

Decision making under risk

On leaving university, you receive two job offers.

The first is from an established company offering a guaranteed salary of \$150,000.

The second is from a promising startup offering a base salary of \$100,000, plus stock options that will be worth either \$300,000 or nothing, with equal probability.

Despite the potential for a large windfall, many people would choose the established company's lower but guaranteed amount. This common preference reveals something fundamental about how people make decisions under risk.

In this part, I examine the basic frameworks by which economists analyse such decisions under risk. I will explore questions such as:

- How do people evaluate risky choices?
- How can we mathematically model risk preferences?
- Where do standard economic models fail in predicting decisions under risk?

I begin by setting out some mathematical foundations, including the concept of expected value. I then lay out the axioms that underlie Expected Utility Theory, the classical economic framework for analysing decisions under risk. Expected utility theory suggests people maximize their expected utility rather than expected value, helping explain phenomena like risk aversion.

However, research has identified anomalies where people's choices systematically deviate from Expected Utility Theory's predictions. We'll examine several anomalies, including the Allais Paradox, where people make inconsistent choices between pairs of gambles. We will look at how framing and reference points affect risk preferences. We will also see cases where small-stakes risk aversion implies implausibly extreme large-stakes risk aversion

Understanding the insights and limitations of Expected Utility Theory provides a foundation for us to later examine alternative theories from behavioural economics to explain decision making under risk.

Notation and mathematical background

Summary

- Lotteries can be represented as $L=(p_1,x_1;p_2,x_2;...;p_n,x_n),$ where p_i is the probability of outcome $x_i.$
- Mathematical functions used in expected utility theory include exponentiation (x^a) , exponential (e^x) , and logarithmic $(\ln(x))$ functions.
- Differentiation finds the rate of change of a function, written as $\frac{d}{dx}f(x)$ or f'(x).
- The second derivative $(\frac{d^2}{dx^2}f(x) \text{ or } f''(x))$ measures a function's curvature, indicating whether it's convex (f''(x) > 0) or concave (f''(x) < 0).

5.1 Notation

Before analysing decision-making under risk and uncertainty, I will introduce some notation.

Suppose we have a lottery L with n possible outcomes $x_1, x_2, ..., x_n$ each with probabilities $p_1, p_2, ..., p_n$. A shorthand way to write this is:

$$L = (p_1, x_1; p_2, x_2; ...; p_n, x_n)$$

For example, suppose you are offered a gamble with a 50% probability of winning \$200 and a 50% probability of losing \$100. We can write this as:

$$L = (0.5, -100; 0.5, 200)$$

The order of each outcome-probability pair does not matter. I could also write:

$$L = \left(0.5, 200; 0.5, -100\right)$$

You may also see gambles represented with the outcome and probability in a different order, such as:

$$L=(x_1,p_1;x_2,p_2;\ldots;x_n,p_n)$$

Or:

$$L = (x_1, x_2, ..., x_n; p_1, p_2, ..., p_n)$$

It is typically not difficult to determine which is which.

5.2 Mathematical background

We will use some basic mathematical concepts to analyse expected utility. I will briefly review these concepts here.

5.2.1 Functions

5.2.1.1 Exponentiation

Exponentiation is a mathematical operation where a number is multiplied by itself a certain number of times.

Exponentiation is written as $f(x) = x^a$. That is, x is multiplied by itself a times. For example, $2^3 = 2 \times 2 \times 2 = 8$.

The exponent *a* can be any real number, including fractions and negative numbers. For example, a plot of the function $f(x) = x^{0.5}$ is shown in Figure 5.1.

5.2.1.2 The exponential function

The exponential function is written as $f(x) = e^x$. The letter e is a constant, approximately equal to 2.71828. It is a special case of exponentiation where the base is e, which is multiplied by itself x times.

A plot of the exponential function is shown in Figure 5.2.



Figure 5.2: The exponential function

5.2.1.3 The logarithmic function

The logarithmic function is written as $f(x) = \ln(x)$ or $\log_e(x)$.

The logarithmic function is the inverse of the exponential function. That is, if $f(x) = e^x$, then $x = \ln(f(x))$.

A plot of the logarithmic function is shown in Figure 5.3.



Figure 5.3: The logarithmic function

Note that the logarithmic function is only defined for positive values of x. The logarithm of zero is undefined.

5.2.2 Differentiation

Differentiation is a mathematical operation that finds the rate of change (or slope) of a function. It is written as $\frac{d}{dx}f(x)$ or $\frac{dy}{dx}$ or f'(x).

There are several simple rules to differentiate a function. The rules relevant to these notes are as follows.

The derivative of a constant is zero.

$$\frac{d}{dx}c = 0$$

The derivative of an exponentiation is:

$$\frac{d}{dx}x^a = ax^{a-1}$$

For example:

$$\frac{d}{dx}x^2 = 2x$$

You can see from this that for any value of x greater than zero, the derivative of x^2 is greater than zero, signifying that the function $f(x) = x^2$ is increasing and has positive slope. For any value of x less than zero, the derivative is less than zero, signifying that the function is decreasing and has negative slope.

As another example:

$$\frac{d}{dx}x^{0.5} = 0.5x^{-0.5}$$

You can see from this that for any value of x greater than zero, the derivative of $x^{0.5}$ is greater than zero, signifying that the function $f(x) = x^{0.5}$ is increasing and has positive slope. The function is not defined for x < 0. This is shown in Figure 5.1.

The derivative of the logarithmic function is:

$$\frac{d}{dx}\ln(x) = \frac{1}{x}$$

This derivative is positive for all values of x for which $\ln(x)$ is defined. Therefore $\ln(x)$ is increasing in x. You can see this in Figure 5.3.

The derivative of a fraction is:

$$\frac{d}{dx}\frac{1}{f(x)} = -\frac{f'(x)}{f(x)^2}$$

For example:

$$\frac{d}{dx}\frac{1}{x} = -\frac{1}{x^2}$$

Where you have a function $\frac{1}{x^a}$, it is often easier to write it as x^{-a} and use the rule for exponentiation. For example:

$$\frac{d}{dx}\frac{1}{x} = \frac{d}{dx}x^{-1} = -1x^{-2} = -\frac{1}{x^2}$$

5.2.2.1 The second derivative

The second derivative of the function is a measure of the curvature of the function or the rate of change of the slope. We can calculate the second derivative by taking the derivative of the first derivative.

$$\frac{d^2}{dx^2}f(x) = \frac{d}{dx}\left(\frac{d}{dx}f(x)\right)$$

We can use the second derivative to determine whether a function is concave or convex. A function is concave if the second derivative is negative and convex if the second derivative is positive.

$$\frac{d^2}{dx^2}f(x) > 0 \text{ for all } x \Rightarrow f(\mathbf{x}) \text{ is convex}$$
$$\frac{d^2}{dx^2}f(x) < 0 \text{ for all } x \Rightarrow f(\mathbf{x}) \text{ is concave}$$

The second derivative of a function is written as $\frac{d^2}{dx^2}f(x)$ or $\frac{d^2y}{dx^2}$ or f''(x). For example, if $f(x) = x^2$, then:

$$\frac{d^2}{dx^2}x^2 = \frac{d}{dx}2x = 2$$

The second derivative is positive (equal to 2) for all values of x. This implies that $f(x) = x^2$ is increasing at an increasing rate. The function is convex.

The second derivative of $x^{0.5}$ is:

$$\frac{d^2}{dx^2}x^{0.5} = \frac{d}{dx}0.5x^{-0.5} = -0.25x^{-1.5}$$

The second derivative is negative for all values of x for which $x^{0.5}$ is defined. This implies that $x^{0.5}$ is increasing at a decreasing rate. The function is concave. You can see this in Figure 5.1.

The second derivative of the logarithmic function is:

$$\frac{d^2}{dx^2}\ln(x)=\frac{d}{dx}\frac{1}{x}=-\frac{1}{x^2}$$

This second derivative is negative for all values of x for which $\ln(x)$ is defined. This implies that $\ln(x)$ is increasing at a decreasing rate. The function is concave. You can see this in Figure 5.3. When working through these notes, you will not be asked to differentiate any functions. However, understanding what differentiation is and what it shows will help you understand the intuition behind the concepts I discuss. I will use the functions $f(x) = \ln(x)$ and $f(x) = x^{0.5}$ in future sections.

Expected value

Summary

- The expected value of a gamble is the average outcome in the long run, calculated as the probability-weighted sum of potential outcomes.
- For a gamble with n possible outcomes x_i , each occuring the probability p_i , the expected value is expressed as:

$$\begin{split} \mathbb{E}[X] &= p_1 x_1 + p_2 x_2 + \ldots + p_n x_n \\ &= \sum_{i=1}^n p_i x_i \end{split}$$

6.1 Introduction

The expected value of a gamble is the amount you can expect to win on average, in the long run, when you play a gamble.

Suppose I offer to flip a coin. I give you \$1 if it is heads and you will give me \$1 if it is tails. What is the expected value of this gamble?

The expected value is \$0. You lose \$1 half the time and gain \$1 half the time.

Formally, given a gamble, the expected value E[X] of the random variable X is the probability-weighted sum of the potential outcomes. That is, we calculate the expected value by multiplying each possible outcome by the probability with which it occurs.

For the coin flip example, you multiply the 50% probability of heads by the \$1 outcome and the 50% probability of tails by the -\$1 outcome.

Probability	Outcome
50%	+\$1
50%	-\$1

$$E[X] = 0.5 \times 1 + 0.5 \times (-1) = 0$$

We calculate the expected value of a gamble with n possible outcomes using the following equation:

$$\begin{split} E[X] &= p_1 x_1 + p_2 x_2 + \ldots + p_n x_n \\ &= \sum_{i=1}^n p_i x_i \end{split}$$

In this equation, p_i is the probability of outcome x_i .

For those unfamiliar with the mathematical notation in the second line, the large symbol sigma allows us to write what could be a long expression much more succinctly. It means that we sum the term $p_i x_i$ for each value of *i* for 1 through to *n*. We sum $p_1 x_1$ with $p_2 x_2$ and so on until we reach $p_n x_n$. Breaking it down this way shows that the second line is equivalent to what I wrote in the first line.

6.2 Expected value examples

I will now illustrate the concept of expected value with some simple examples.

6.2.1 Example 1

You are offered a bet with a 50% chance of winning \$10 and a 50% chance of losing \$8.

The expected value of the gamble X is:

$$E[X] = \sum_{i=1}^{n} p_i x_i$$
$$= 0.5 \times 10 + 0.5 \times (-8)$$
$$= \$1$$

Relating back to my earlier explanation of the summation symbol, here we have n = 2 outcomes. We sum $p_1 = 0.5$ multiplied by $x_1 = 10$ with $p_2 = 0.5$ multiplied by $x_2 = -8$.

Suppose your chance of winning increases to 60%. The expected value of the gamble is:

$$\begin{split} E[X] &= \sum_{i=1}^{n} p_i x_i \\ &= 0.6 \times 10 + 0.4 \times (-8) \\ &= \$2.80 \end{split}$$

6.2.2 Example 2

You are offered a bet with a 50% chance of winning 50% of your wealth and a 50% chance of losing 40% of your wealth.

The expected value of the bet is:

$$\begin{split} E[X] &= \sum_{i=1}^n p_i x_i \\ &= 0.5 \times 0.5W + 0.5 \times (-0.4W) \\ &= 0.05W \end{split}$$

The expected value is 5% of your wealth.

Axioms for Expected Utility Theory

Summary

- Expected Utility Theory is based on four main axioms: completeness, transitivity, continuity, and independence.
- The continuity axiom states that if an agent prefers option A to B, and B to C, there exists some probability mixture of A and C that the agent finds equally desirable to B.
- The independence axiom asserts that a preference between two lotteries should not be affected by mixing both with an identical third lottery.
- These axioms allow for the construction of a cardinal utility function that represents an agent's preferences over lotteries.
- Additional auxiliary axioms often adopted include a reference point of zero wealth, non-satiation, monotonicity, convexity, and diminishing marginal utility.

7.1 Introduction

In my discussion of rationality, I noted that in its most basic form analysis of consumer choice rests on just two assumptions: completeness and transitivity.

For decision-making under risk, we require two additional axioms of desirable behaviour to develop a predictive or descriptive model. These additional axioms are continuity and the independence.

That gives us four axioms:
- Completeness
- Transitivity
- Continuity
- Independence

Under these axioms, a decision maker behaves as if choosing between risky prospects by selecting the one with the highest expected utility.

The four axioms are often called the von Neumann–Morgenstern axioms for a rational agent. This gives us another benchmark of "rationality". An agent is rational if they conform with these four axioms.

One important feature of preferences under these assumptions is that utility is cardinal. The magnitude, not just rank, of the numbers matters.

If you look at other resources on the axioms underlying expected utility theory, you may come across an axiom called the Archimedean property. The Archimedean property is an alternative assumption to continuity. Only one of continuity or the Archimedean property need be assumed. I will not cover the Archimedean property in these notes.

Beyond the axioms of completeness, transitivity, continuity and independence, some additional axioms are often adopted for practical purposes. These include using a reference point of zero wealth, non-satiation, monotonicity, convexity and diminishing marginal utility.

In the following sections, I discuss each of the von Neumann-Morgenstern axioms and the auxiliary axioms we use in examining decision making under risk.

7.2 Continuity

The idea behind continuity is that people have similar preferences for similar bundles. If x is preferred to y, bundles close to x are preferred to bundles close to y. There are no "jumps" in utility.

Continuity guarantees that every preference relation can be represented by a continuous utility function (and vice versa).

Here are two formal definitions.

7.2.1 Definition 1

A preference relation is continuous if for any $x \succ y$ there exists a number $\epsilon > 0$ such that every bundle *a* that is less distant from *x* than ϵ and every bundle *b* that is less distant from *y* than ϵ results in $a \succ b$.

To put this another way, a preference relation is continuous if for any $x \succ y$ there are *some* neighbourhoods $N_{\epsilon}x$ and $N_{\epsilon}y$ around x and y such that for every $a \in N_{\epsilon}x$ and $b \in N_{\epsilon}y$ we have $a \succ b$. One way to picture this is to imagine a circle around bundles x and y of radius ϵ . These circles represent the neighbourhood. There will always exist some circle - even if very small - within which every bundle a within the neighbourhood of x is preferred to bundle b within the neighbourhood of y.



The intuition behind this definition is that a very small change in your bundle should not result in a sudden switch of your preferences. If you prefer 5 bananas to 2 oranges, you will likely prefer 4.9 bananas to 2 oranges. (And if not, there will be some amount of bananas between 4.9 and 5 that you prefer over 2 oranges.)

Here's another intuitive example: if you prefer a Mercedes to a Toyota, there will be *some* level of defect in the Mercedes that you would be willing to accept while still preferring the Mercedes to the Toyota.

7.2.2 Definition 2

If x, y and z are lotteries with $x \succeq y \succeq z$, the continuity axiom requires that there exists a probability p such that y is equally as good as a mix of x and z. That is, there exists p such that:

$$px + (1-p)z \sim y$$

The below diagram illustrates continuity under this definition.

On the diagram are three bundles: x, y and z, and each sits on a different indifference curve. The indifference curve that x is on is higher than that of y which is higher than that of z. That is, $x \succeq y \succeq z$.

Now consider a gamble that pays x with probability p and z with probability 1-p. Each value of p would result in a gamble with utility falling between that of x and z. If we were to draw a line between x and z, you could think of the utility of the gamble for each value of p as having the same utility as a bundle on that line. Under the continuity axiom, there would be no holes in that line. For some value of p, that gamble will be on the same indifference curve for y. At that point, $px + (1-p)z \sim y$.



7.2.3 Example of discontinuous preferences

Lexicographic preferences occur where an agent prefers any amount of a good x to any amount of another good y. If choosing between bundles of goods, the agent will choose the bundle with the most x, regardless of the amount of y. They will only consider the amount of y if the amount of x in two bundles is identical.

Consider an agent with lexicographic preferences who is offered the following combinations of x and y.

A. (1, 1)

B. (1, 2)

C. (1.1, 1)

Their preference ranking will be $C \succ B \succ A$. They prefer C as it has more x than the other two options. As A and B have the same amount of x, the agent distinguishes them based on the quantity of y, preferring B.

This function is not continuous as there is a "jump" whenever there is an increase in x, even if y is large. Add an infinitesimal amount ϵ of x to bundle A and the preference relation between bundle A and bundle B flips.

These three bundles A, B and C are represented graphically below.

First, let's consider these preferences in terms of the first definition, being that there are *some* neighbourhoods around A and B such that we will always prefer another bundle of goods within the neighbourhood of B to any bundles within the neighbourhood of A. Around A I have drawn a circle of radius ϵ , which we can consider to be the neighbourhood. No matter how small I draw this circle - that is, no matter how small ϵ - any bundle within the circle that lies to the right of A (that is, contains x > 1) is preferred to bundle B. There is a jump in preferences to the right of A.



One interesting feature of lexicographic preferences is that you cannot draw indifference curves on this figure. If a bundle differs from another, it must be strictly preferred to the other as no amount of y can make up for any amount of x.

We can also consider lexicographic preferences in terms of the second definition

of continuity. There is no p for which:

$$pA + (1-p)C \sim B$$

When p = 1, $B \succ A$. For any p < 1, $pA + (1 - p)C \succ B$ as any non-zero share of C makes the combination of A and C preferred.

7.3 Independence

Consider the following scenarios.

In the first, a person has a choice between an orange and an apple. They state that they strictly prefer the orange.

In the second, they are offered a choice between two gambles. The first gamble is a 50% chance of an orange and a 50% chance of a pear. The second is a 50% chance of an apple and a 50% chance of a pear. They state that they strictly prefer the gamble with a 50% chance of an orange.

Compare the two scenarios. The choice between an apple or an orange from the first scenario is mixed with a 50% probability of a pear in the second.

Under the axiom of independence, a person who prefers the orange in the first will prefer the gamble with the orange in the second. Mixing those two lotteries (a 100% chance of an orange or a 100% chance of an apple) with a third lottery - in this case, a pear - will not change their order of preference.

More generally, under the axiom of independence, a person who mixes two lotteries with a third lottery will maintain the same order of preference when the lotteries are mixed as they had for the two original lotteries when presented independently of the third.

A formal definition states that if:

- x and y are lotteries with $x \succcurlyeq y$ and
- p is the probability that a third option z is present, then:

$$pz + (1-p)x \succcurlyeq pz + (1-p)y$$

The third choice, z does not change the preference ordering. The order of preference for x over y holds. It is independent of the presence of z.

7.3.1 Example of the axiom of independence

Let us put our earlier example into this formal definition.

Suppose x is a 100% probability of an orange and y is a 100% probability of an apple. I strictly prefer an orange to an apple.

Suppose there is now a third possibility z of receiving a pear, which will be present with p = 50% probability.

Under the axiom of independence, if I prefer oranges to apples, I will prefer a gamble with a 1 - p = 50% chance of getting an orange and p = 50% chance of receiving a pear to a gamble with a 1 - p = 50% chance of getting an apple and a p = 50% chance of receiving a pear.

That is:

orange \succ apple \Longrightarrow 50% chance of orange + 50% chance of pear \succ 50% chance of apple + 50% chance of pear

7.3.2 Distinguishing the independence of irrelevant alternatives from the independence axiom

The independence axiom is distinct from the principle of the independence of irrelevant alternatives.

The independence of irrelevant alternatives states that if an agent prefers x to y, the introduction of a third option z should not change the preference order between x and y. For example, if you select fish rather than chicken from a restaurant menu, being told by the waiter that there is a vegetarian option should not lead you to change your selection to chicken.

The independence axiom is specific to lotteries. The logic behind this specificity is that the outcomes of a lottery are never realised together. They can be treated as independent. In my illustration involving apples, oranges and pears, there is no outcome where the agent receives more than a single piece of fruit. They will receive an apple, an orange or a pear. They will not receive a mix of fruit.

This is not the case for goods. Consider the following example with goods drawn from Page (2022). You are again in a restaurant and have a choice between chicken with mashed potato and beef with mashed potato. You choose the chicken. You are then told that the restaurant has run out of mashed potato, and the options are now chicken or beef with peas. Under the axiom of independence, you would not change your choice to beef. However, beef may go better with peas than chicken. There is an interaction between the two, with the options realised together. Due to this interaction, the axiom of independence is less compelling for the case of goods than it is for lotteries.

7.4 Auxiliary axioms for expected utility theory

Beyond completeness, transitivity, continuity and independence, economists often adopt other axioms. These are not required for expected utility theory, but make analysis more practicable. These include:

- Reference point of zero wealth
- Non-satiation
- Monotonicity
- Convexity
- Diminishing marginal utility

I provide further detail on these.

7.4.1 Reference point of zero wealth

When people are considering whether to accept or gamble or compare options, they do not decide from a blank slate. They come with an existing set of resources (wealth), and that wealth may affect their decision. A gamble may be more attractive if someone has more or less wealth.

This necessitates the setting of a "reference point" from which utility is calculated. In Expected Utility Theory, that reference point is typically considered to be zero wealth.

The way this is implemented is we typically calculate utility over total wealth. For example, if offered a gamble where they could win or lose \$10, we do not calculate the utility of each option as U(\$10) and U(-\$10). Rather, the utility of each option is calculated as U(W + \$10) and U(W - \$10).



The practical impact of this implementation is that people's choices may differ

depending on their wealth. The same gamble may be accepted or rejected at different levels of wealth.

7.4.2 Non-satiation

The idea behind non-satiation is that no matter what you have, there is always another (nearby) bundle that you would rather have. There is no "maximum" level of utility that you can achieve. Whatever your utility now, you can always increase it.

In this diagram I have plotted an indifference curve. Point x is on the curve. For non-satiation, there will always be a point, such as y, that is strictly preferred to x.



7.4.3 Monotonicity

Preferences are monotone if more of any good in the bundle makes the agent strictly better off. Non-satiation is implied by monotonicity, but not the other way around.

Monotonicity implies downward-sloping indifference curves. This is because any increase of a good in your bundle would take you to a higher indifference curve. A horizontal indifference curve is not feasible as moving along that indifference curve implies more of the good, but that is not possible as monotonicity implies you are better off and hence on a higher indifference curve. This can be seen in the following diagram. Point x lies on the indifference curve. Increasing the amount of either good will take you to a higher indifference curve. That is true for all points on that indifference curve.



7.4.4 Convexity

Convexity means that people have a preference for variety or combination (indifference curves bulge toward the origin). Averages are better than extremes.

In many contexts this makes sense. For example, suppose you are indifferent between two beers and two meat pies. Under this assumption, any mix of the two, such as a beer and a pie will be at least as preferred as the two beers or two pies.

Formally, a preference relation is convex if, for any $x \succeq y$ and for every $\theta \in [0, 1]$:

$$\theta x + (1 - \theta)y \succcurlyeq y$$

This definition is illustrated in the following diagram. The curve represents an indifference curve for different combinations of two goods. There are two bundles, x and y. In this case, $x \succeq y$ (as $x \sim y$). Any weighted combination of x and y, which would be on the line between the two, can be seen to be strictly preferred to either x or y as it would be on a higher indifference curve (a curve further from the origin).



One point to note from this diagram is that if the indifference curve were a straight line, any point on a line between x and y would also be on that line, and weakly preferred to x and y. This would still be a convex curve.

This contrasts with strict convexity. Strict convexity is where, for any $x \sim y$, $x \neq y$ and for every $\theta \in [0, 1]$:

$$\begin{aligned} \theta x + (1-\theta)y \succcurlyeq y \\ \theta x + (1-\theta)y \succcurlyeq x \end{aligned}$$

For two equivalent goods or bundles, a weighted average of the two bundles is better than each of those bundles.

7.4.5 Diminishing marginal utility

Suppose you want some chocolate. You eat a piece. You then eat another. And another. How much utility do you imagine getting from the first piece compared to the 100th piece? The first piece will likely be much more satisfying than the 100th. This is the idea of diminishing marginal utility.

Marginal utility is how much utility you get or lose from an incremental decrease or increase in consumption. Under diminishing marginal utility, each successive additional unit of consumption delivers a smaller (diminishing) amount of utility than the last.

This concept is illustrated in the following diagram. The curve represents an indifference curve for the good x. The curve is concave, which means that the slope of the curve decreases in x and the marginal utility of each additional unit of good x decreases as you consume more of it. One additional unit of good x when the agent has a units of the good leads to a much larger increase in utility than one additional unit when the agent has b units of the good.



Diminishing marginal utility leads to risk-averse preferences. Someone is risk averse if they strictly prefer the expected value of a gamble to the gamble itself.

Diminishing marginal utility is related to the axiom of convexity. Diminishing marginal utility will lead to convex indifference curves. However, the reverse relationship does not always hold.

Chapter 8

Expected Utility Theory

Summary

- Expected utility theory states people maximise expected utility, not expected value, in choosing between risky prospects.
- Expected utility is calculated by weighting the utility of each outcome by its probability and summing.

$$\mathbb{E}[U(X)] = \sum_{i=1}^n p_i U(x_i)$$

• Outcomes typically represent final wealth positions, not just payoffs, so decisions depend on current wealth. The formula incorporating initial wealth used for calculating expected utility is:

$$\mathbb{E}[U(W+X)] = \sum_{i=1}^n p_i U(W+x_i)$$

- Expected utility theory allows the examination of an agent's attitude toward risk, which can be risk aversion, risk neutrality, or risk seeking.
- A risk-averse person prefers a sure amount to a gamble with the same expected value.

$$U(\mathbb{E}[X]) > \mathbb{E}[U(X)]$$

• A risk-seeking person prefers a gamble to a sure amount of the same expected value.

$$U(\mathbb{E}[X]) < \mathbb{E}[U(X)]$$

• A risk-neutral person is indifferent between a gamble and receiving its expected value with certainty.

$$U(\mathbb{E}[X]) = \mathbb{E}[U(X)]$$

• Certainty equivalent (CE) is the amount of money that makes a person indifferent between taking a gamble and taking the money. For risk-averse individuals, $CE < \mathbb{E}[X]$; for risk-neutral, $CE = \mathbb{E}[X]$; for risk-seeking, $CE > \mathbb{E}[X]$.

8.1 Introduction

Under expected utility theory, people do not seek to maximize expected value but instead, maximize expected utility.

Under expected utility theory, people choose between risky prospects (prospect being another word for lottery or gamble common in the literature) by comparing expected utility values. An agent would pick the option that maximises their expected utility.

8.2 Calculating expected utility

Consider a prospect with final outcomes $x_1, ..., x_n$, with each outcome occurring with probability $p_1, ..., p_n$. Each outcome would deliver utility $U(x_i)$.

$$X = (p_1, x_1; \ p_2, x_2; \dots \ ; \ p_n, x_n)$$

For example, the prospect might be a coin flip that delivers a win of \$10 for heads and a loss of \$10 for tails. We would represent this prospect as:

$$X = (0.5, \$10; 0.5, -\$10)$$

Expected utility, E[U(X)], is calculated using the following formula:

$$\begin{split} E[U(X)] &= p_1 U(x_1) + p_2 U(x_2) + \ldots + p_n U(x_n) \\ &= \sum_{i=1}^n p_i U(x_i) \end{split}$$

You can think of this formula as comprising the following steps:

- 1. Define utility $U(x_i)$ over final outcomes $x_1, ..., x_n$
- 2. Weight the utility of each outcome $U(\boldsymbol{x}_i)$ by the probability p_i of outcome \boldsymbol{x}_i
- 3. Add the weighted utilities.

For the coin toss that delivers a win or loss of \$10, we would write:

$$E[U(X)] = 0.5 \times U(\$10) + 0.5 \times U(-\$10)$$

There is an important note regarding outcomes $x_1, ..., x_n$. Typically, these outcomes are not just the payoffs from the gamble, but rather the agent's final position. If the agent has wealth of \$100 and is offered a coin flip to win or lose \$10, the outcomes are typically taken to be \$90 and \$110. Their decision depends on their current wealth. As a result, expected utility is often represented as in this equation:

$$\begin{split} E[U(W+X)] &= p_1 U(W+x_1) + p_2 U(W+x_2) \\ &+ \ldots + p_n U(W+x_n) \\ &= \sum_{i=1}^n p_i U(W+x_i) \end{split}$$

Returning to our coin toss example, if our starting wealth was \$100, the expected utility of the coin toss is:

$$E[U(X)] = 0.5 \times U(\$110) + 0.5 \times U(\$90)$$

8.3 Attitudes toward risk

Expected utility theory allows us to examine an agent's attitude toward risk.

There are three possible attitudes toward risk: risk aversion, risk neutrality, and risk seeking.

If a person prefers a sure amount to a gamble with the same expected value, they are risk averse. That is, the utility of the expected value of X is greater than the expected utility of X.

$$U(E[X]) > E[U(X)]$$

If a person prefers a gamble to a sure amount of the same expected value, they are risk-seeking. That is, the expected utility of X is greater than the utility of the expected value of X.

$$U(E[X]) < E[U(X)]$$

If a person is indifferent between a gamble and receiving the expected value of the gamble with certainty, they are risk neutral. That is. the utility of the expected value of X is equal to the expected utility of X.

$$U(E[X]) = E[U(X)]$$

8.3.1 Certainty equivalent

One useful concept in the analysis of attitudes to risk is the certainty equivalent. The certainty equivalent (CE) of a gamble X is the amount of money such that you are indifferent between taking the gamble and taking the money. That is, the utility of the certainty equivalent is equal to the expected utility of the gamble.

$$u(CE) = E[U(X)]$$

For a risk-averse person, the certainty equivalent of the bet is less than the expected value of the gamble. The certainty equivalent is equal to the expected value in the case of risk neutrality. A risk-seeking person would have a certainty equivalent higher than the expected value of the gamble.

I will now look at each of the attitudes in turn.

8.3.2 Risk aversion

A risk-averse person prefers a sure amount to a gamble with the same expected value. If I strongly prefer \$10 for certain to a gamble with an expected value of \$10, I am risk averse. The certainty equivalent of the prospect for this person would be less than \$10.

The following chart illustrates. The x-axis is the amount of the good over which the person receives utility. In the case of monetary gambles of the type we are talking about, the x-axis is the amount of money. The y-axis is the utility the person receives from that money.

The utility curve is a plot of the utility function for each amount of money.

The gamble shown in this chart has two possible outcomes, x_1 and x_2 . An outcome of x_1 would result in utility $U(x_1)$. An outcome of x_2 would result in utility $U(x_2)$.



Figure 8.2: Risk aversion 2

I have drawn a straight dash-dot line between the points on the utility curve for each of those outcomes. The expected utility from any gamble involving those two outcomes would fall on that line. Where on that line would depend on the probability of each outcome, and it would occur at a point in line with the expected value of the gamble. You can see the expected value of the gamble marked on the x-axis. If we extend up from that point, we can read the expected utility of the gamble, E[U(X)], from the y-axis.



Figure 8.3: Risk aversion 3

One way to think about why projecting the expected value onto the expected utility line identifies the expected utility of the bet is that each of the expected value and expected utility are weighted by the same probabilities. They both lie the same horizontal distance between the two potential outcomes.

In this example, the line between the utility of the two outcomes is always below the utility curve. That is, for any probabilities involving those two outcomes (except one of those outcomes with certainty), the expected utility of the prospect is less than the utility of the expected value of the prospect. In mathematical terms:

$$U(E[x]) > E[U(x)]$$

The expected utility line between the two outcomes is always below the utility

curve as the curve is concave. A concave curve leads to risk aversion. The greater the curvature, the more risk averse the agent is.

Finally, the horizontal dashed line identifying E[U(X)] allows us to identify the certainty equivalent of the gamble. At the point on the x-axis marked CE, the utility from the certainty equivalent is equal to the expected utility of the prospect with expected value E[X]. As the certainty equivalent is less than the expected value, this provides another way of saying that the person is risk averse.



Figure 8.4: Risk aversion 4

8.3.2.1 Absolute and relative risk aversion

In studying economics, you may encounter two measures of risk aversion: absolute risk aversion and relative risk aversion.

Absolute risk aversion is a measure of how risk averse a person is at any particular level of wealth. A person with constant absolute risk aversion (CARA) responds to a fixed sum bet in the same way, whatever their wealth.

For example, when faced with a 50-50 bet to win \$20 or lose \$10, they will make the same decision regardless of whether they have \$100 or \$1,000,000 in their bank account.

Relative risk aversion is a measure of how an individual's risk aversion scales with their wealth. A person with constant relative risk aversion (CRRA) would respond to a bet risking a certain proportion of their wealth in the same way, whatever their wealth. This means that as their wealth increases, the absolute amount they're willing to risk increases, but it remains a constant proportion of their total wealth.

For example, someone with constant relative risk aversion would always respond in the same way to a 50-50 bet to win 50% of their wealth, lose 40% of their wealth, regardless of whether that loss involves a possible \$400 out of \$1,000 or \$400,000 out of \$1,000,000.

Constant absolute risk aversion utility functions are often used in theoretical models due to their mathematical simplicity. However, constant relative risk aversion utility functions are generally considered more realistic for modelling human behaviour, as people tend to adjust their risk-taking based on their wealth.

One feature of constant relative risk aversion utility functions, such as log utility, is that people become less risk averse (as measured by absolute risk aversion) as their wealth increases. For a bet of a certain sum, they will be more likely to accept that bet at higher wealth.

This above plot is of a utility function with constant relative risk aversion. You can see that as the level of wealth increases, the utility function becomes increasingly linear. This reduction in curvature reflects the declining absolute risk aversion and results in bets of fixed value being more likely to be accepted at higher wealth.

8.3.3 Risk neutrality

A risk-neutral person is an expected value maximiser. They are indifferent between \$10 for certain and a gamble with an expected value of \$10. The certainty equivalent of the prospect for a person considering this gamble would also be \$10.

The following chart illustrates. Again, we have two possible outcomes, x_1 and x_2 , with resulting utility $U(x_1)$ and $U(x_2)$.

A line between the points on the utility curve for each of those outcomes lies on the utility curve itself. For any probability, the utility of the expected value and the expected utility are the same.

8.3.4 Risk seeking

A risk-seeking person prefers a gamble to a sure sum equal to the expected value of that gamble. The certainty equivalent is more than the expected value of the gamble. The gamble has value in and of itself.

The following chart illustrates. Again I have drawn a dash-dot-dot line between the points on the utility curve for each of those outcomes. That line is always



Figure 8.5: Risk neutral

above the utility curve. That is, for any probabilities involving those two outcomes (except one of those outcomes with certainty), the expected utility of the prospect is more than the utility of the expected value of the prospect.

The expected utility line between the two outcomes is always above the utility curve as the curve is convex. A convex curve leads to risk-seeking behaviour.



Figure 8.6: Risk seeking

Chapter 9

Expected utility examples

Summary

- In the following examples I demonstrate the operation of expected utility theory and concepts such as expected value, expected utility and certainty equivalent through various betting scenarios, illustrating how individuals with risk-averse utility functions make decisions under uncertainty.
- In a 50:50 bet example with a logarithmic utility function, a gamble with zero expected value reduces utility, demonstrating risk aversion.
- An 80:20 bet example shows that even risk-averse individuals may accept favourable bets, but the certainty equivalent is still less than the expected value.
- A bet involving proportions of wealth demonstrates that positive expected value does not always lead to acceptance of a gamble for risk-averse individuals.
- The St. Petersburg game demonstrates a paradox where the expected value is infinite, yet most people would not pay an infinite amount to play. This paradox can be resolved by introducing expected utility theory, which shows that risk-averse individuals with diminishing marginal utility would only pay a finite amount to play the game.

9.1 Introduction

In this section, I present a series of mathematical examples of expected utility theory.

9.2 A 50:50 bet

Suppose your utility function is $U(x) = \ln(x)$.

You have a 50% chance of winning \$10 and a 50% chance of losing \$10. Assume your starting wealth is \$20.

What is the expected value of this game?

$$\begin{split} E[X] &= \sum_{i=1}^{n} p_i x_i \\ &= 0.5 \times 10 + 0.5 \times (-10) \\ &= 0 \end{split}$$

The expected value of the game is \$0.

What is the expected utility of this game?

$$\begin{split} E[U(W+X)] &= \sum_{i=1}^n p_i U(x_i+W) \\ &= 0.5U(20-10) + 0.5U(20+10) \\ &= 0.5 \mathrm{ln}(10) + 0.5 \mathrm{ln}(30) \\ &= 2.85 \end{split}$$

What does an expected utility of 2.85 mean? To make it tangible, we can ask what wealth would give that utility.

$$U(W) = \ln(W) = 2.85$$
$$W = e^{2.85} = \$17.30$$

This gamble with an expected value of zero reduces utility by an amount equivalent to \$2.70.

We could also say that the certainty equivalent of this gamble is the final wealth of \$17.30, or a loss of \$2.70.

Figure 9.1 illustrates the example.

On the x-axis, we have the outcomes and on the y-axis, we have the utility.

I have added points on the x-axis for the outcomes of the two gambles, being W-10 and W+10. They deliver utility U(W+10) and U(W-10) respectively.



Figure 9.1: A 50:50 bet

The expected utility of the gamble is the probability-weighted average of these two points. It sits on the straight dash-dot-dot line between those two outcomes.

You can see that the expected utility of the gamble is lower than the utility of the expected value (being current wealth).

Also plotted is the certainty equivalent. We can identify it as the point on the utility curve where the utility of that certainty equivalent is equal to the expected utility.

9.3 An 80:20 bet

Suppose your utility function is $U(x) = \ln(x)$.

You have an 80% chance of winning \$10 and a 20% chance of losing \$10. Assume your starting wealth is \$20.

What are the expected value and the expected utility of this game?

$$\begin{split} E[X] &= \sum_{i=1}^n p_i x_i \\ &= 0.8 \times 10 + 0.2 \times (-10) \\ &= \$6 \end{split}$$

The expected value of the game is \$6. What is the expected utility of this game?

$$\begin{split} E[U(W+x)] &= \sum_{i=1}^n p_i U(x_i+W) \\ &= 0.8 U(20+10) + 0.2 U(20-10) \\ &= 0.8 \mathrm{ln}(30) + 0.2 \mathrm{ln}(10) \\ &= 3.18 \end{split}$$

What does an expected utility of 3.18 mean? To make it tangible, we can ask what wealth would give that utility.

$$U(W) = \ln(W) = 3.18$$

 $W = e^{3.18} = 24.08

This gamble with an expected value of \$6 increases utility by an amount equivalent to \$4.08.

We could also say that the certainty equivalent of this gamble is the final wealth of \$24.08.

Figure 9.2 illustrates the example.

The expected utility of the gamble E[U(X)] is higher than the utility from current wealth but lower than the utility of the expected value. That is, they are risk averse but would still accept this highly favourable bet.

Also plotted is the certainty equivalent. We can identify it as the point on the utility curve where the utility of that certainty equivalent is equal to the expected utility. In this case, it is at \$4.08 above current wealth.



Figure 9.2: An 80:20 bet

9.4 Betting a proportion of wealth

Suppose your utility function is $U(x) = \ln(x)$.

You have a 50% chance of increasing your wealth by 50% and a 50% chance of decreasing your wealth by 40%.

What are the expected value and the expected utility of this game?

$$\begin{split} E[X] &= \sum_{i=1}^{n} p_i x_i \\ &= 0.5 \times 0.6W + 0.5 \times 1.5W \\ &= 0.3W + 0.75W \\ &= 1.05W \end{split}$$

The expected value of the gamble is 5% of your wealth. The gamble has a positive expected value.

$$\begin{split} E[U(X)] &= \sum_{i=1}^n p_i U(X_i) \\ &= 0.5U(0.6W) + 0.5U(1.5W) \\ &= 0.5 \mathrm{ln}(0.6) + 0.5 \times \mathrm{ln}(W) + 0.5 \mathrm{ln}(1.5) + 0.5 \times \mathrm{ln}(W) \\ &= -0.255 + 0.203 + \mathrm{ln}(W) \\ &= -0.053 + \mathrm{ln}(W) \end{split}$$

Here we have a gamble with a positive expected value, 5% of your wealth, but lower expected utility. Someone with log utility would reject this bet.

Figure 9.3 illustrates the example.



Figure 9.3: Betting a proportion of wealth

I have added points on the x-axis for the outcomes of the two gambles, a 40% reduction in wealth and a 50% gain in wealth. The expected utility of the gamble is the probability-weighted average of these two points. It sits on the straight dash-dot-dot line between those two outcomes.

You can see that the expected utility of the gamble is lower than the utility of current wealth. They would reject an offer of this bet.

9.5 The St. Petersburg game

The St. Petersburg game was invented by the Swiss mathematician Nicolas Bernoulli.

The game starts with a pot containing \$2. A dealer then flips a coin. The pot doubles every time a head appears. The game ends, and the player wins the pot when a tail appears.

- A tail on the first flip leads to a payment of \$2.
- A tail on the second flip leads to a payment of \$4
- A tail on the third flip leads to a payment of \$8

And so on.

Consider what you would be willing to pay to play this game. Would you pay \$5? \$10? \$25? \$50? More?

The expected value of this game is equal to the sum of the following series.

$$E[X] = \underbrace{\frac{1}{2} \times 2}_{\text{Tail first}} + \underbrace{\left(\frac{1}{2} \times \frac{1}{2}\right) \times 4}_{\text{Tail second}} + \underbrace{\left(\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2}\right) \times 8}_{\text{Tail third}}$$
$$+ \underbrace{\left(\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2}\right) \times 16}_{\text{Tail fourth}} + \dots$$
$$= 1 + 1 + 1 + 1 + \dots$$
$$= \sum_{k=1}^{\infty} 1$$
$$= \infty$$

The first term in the series captures the 50% chance of a tail on the first flip, paying \$2. The second term represents the 50% chance of a head on the first flip, followed by the 50% chance of the tail second flip, paying \$4. The third term represents the 50% chance of a head on the first flip, followed by the 50% chance of a head on the second flip, followed by the 50% chance of a tail on the third flip, paying \$8. And so on.

Multiplying out each of those terms results in a series of 1s.

The \sum operator means "sum for k = 1 to $k = \infty$ ".

Contrast this expected value of ∞ with the sum you would pay to play the game. You were likely not willing to pay an infinite amount.

This "paradox" is often resolved by introducing an expected utility function. The expected utility of this game is equal to:

$$E[U(X)] = \underbrace{\frac{1}{2} \times U(W+2)}_{\text{Tail first}} + \underbrace{\left(\frac{1}{2} \times \frac{1}{2}\right) \times U(W+4)}_{\text{Tail second}} + \underbrace{\left(\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2}\right) \times U(W+8)}_{\text{Tail third}} + \underbrace{\left(\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2}\right) \times U(W+8)}_{\text{Tail fourth}} + \dots = \frac{1}{2}U(W+2) + \frac{1}{4}U(W+4) + \frac{1}{2}U(W+8) + \frac{1}{12}U(W+8)$$

$$\begin{split} &= \frac{1}{2}U(W+2) + \frac{1}{4}U(W+4) + \frac{1}{8}U(W+8) + \frac{1}{16}U(W+16) + \dots \\ &= \sum_{k=1}^{k=\infty} \frac{1}{2^k}U(W+2^k) \end{split}$$

Similar to the calculation of the expected value, the first term in the series captures the 50% chance of a tail on the first flip, paying \$2. The second term represents the 50% chance of a head on the first flip, followed by the 50% chance of the tail on the second flip, paying \$4. And so on. But here, we are using the utility function U(x).

In the second line, I multiplied the probabilities of each coin flip together.

In the third line, I expressed this infinite sum more compactly.

To take this equation further, we need to consider the particular utility function of the decision maker.

What maximum sum would a risk-neutral player with U(x) = x be willing to pay to play the game?

One strategy to determine this sum is to ask what sum would result in the player being indifferent between paying and rejecting a chance to play. That is the maximum sum c that they would be willing to pay. They will be indifferent when U(W)=E[U(X-c)].

We can solve this equation as follows.

$$\begin{split} U(W) &= E[U(X-c)] \\ &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} U(W + \$ 2^k - c) \\ W &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} (W + 2^k - c) \qquad (\text{substituting in the utility function}) \\ &= W - c + \sum_{k=1}^{k=\infty} 1 \qquad \left(\text{as } \sum_{k=1}^{k=\infty} \frac{1}{2^k} = 1 \right) \\ c &= \sum_{k=1}^{k=\infty} 1 \\ &= \infty \end{split}$$

In the second line, we use the sum we created earlier. In the third line, I substitute the utility function U(x) = x. We can then simplify as in the fourth line, which allows us to see that, given the infinite expected value of the game, the player would be willing to pay an infinite amount to play.

That is, a risk-neutral player would pay any amount c to play.

We could also have inferred this from the game's expected value being infinite.

What is the maximum sum a risk-averse player with $U(x) = \ln(x)$ would be willing to pay to play the game? How does their wealth affect their willingness to pay?

Again we will determine at what c the player is indifferent between accepting and rejecting a chance to play, which occurs when U(W) = E[U(X - c)].

$$\begin{split} U(W) &= E[U(X-c)] \\ U(W) &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} U(W + \$ 2^k - c) \\ \ln(W) &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} \ln(W + \$ 2^k - c) \end{split}$$

There is no closed-form solution to this equation to enable us to determine c. We need to solve via numerical methods (such as testing and iterating to a solution).

If we did solve this, we would find that someone who has wealth of \$0.01 would be willing to pay up to \$2.01. They would need to borrow. Someone with wealth \$1000 would be willing to pay \$10.95. A person with a wealth of \$1 million would be willing to pay \$20.87.

We cannot solve for a person with no wealth as $\ln(0)$ is undefined.

Why does willingness to pay increase with wealth?

With log utility, as wealth increases, the slope of the log function increasingly approximates a linear function (the second derivative approaches zero). Hence, the gambler displays less risk-averse (closer to risk-neutral) behaviour.

One way to gain an intuition for why this gamble now has a finite value is to calculate the utility of a risk-averse player whose only asset is the opportunity to play this game.

$$\begin{split} E[U(X)] &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} U(\$2^k) \\ &= \sum_{k=1}^{k=\infty} \frac{1}{2^k} \ln(2^k) \qquad \text{(substituting in the utility function)} \\ &= \sum_{k=1}^{k=\infty} \frac{k}{2^k} \ln(2) \qquad \text{(using the rule } \ln(x^a) = a \ln(x)) \\ &= \left(\frac{1}{2} + \frac{2}{4} + \frac{3}{8} + \frac{4}{16} + \frac{5}{32} + \dots\right) \ln(2) \\ &= 2\ln(2) \end{split}$$

The change in the utility from each flip rapidly declines. Ultimately the series of fractions sum to two.

We can then calculate what wealth is equivalent to this expected utility.

$$U(W) = \ln(W) = 2\ln(2)$$
$$W = e^{2\ln 2} = 4$$

The expected utility from the game is equal to the utility of \$4.

Chapter 10

Anomalies in expected utility theory

Summary

- The Allais Paradox demonstrates that people's choices between pairs of bets often violate expected utility theory, as their preferences can be inconsistent when a common consequence (a shared outcome with equal probability in both options) is changed.
- Rejection of small-scale bets implies unreasonably high risk aversion for larger bets under expected utility theory. For example, rejecting a 50:50 bet to win \$110 or lose \$100 at all wealth levels implies rejecting a 50:50 bet to win \$1 billion or lose \$1,000.
- People's choices are influenced by how options are framed, contradicting expected utility theory's assumption of description invariance.
- People evaluate outcomes relative to reference points rather than absolute wealth.

In this part, I show several anomalies in expected utility theory.

10.1 The Allais Paradox

The Allais paradox is one of the most famous anomalies in expected utility theory.

The paradox was first identified by Maurice Allais (1953). It emerges from the pattern of response to two pairs of bets. The following example comes from Kahneman and Tversky (1979).

For choice 1, the player is asked to choose one of the following bets:

Under Bet A, the player wins:

- \$2500 with probability 33%
- \$2400 with probability 66%
- \$0 with probability 1%

Under Bet B, the player wins:

• \$2400 with probability 100%

Which do you prefer?

When Kahneman and Tversky (1979) ran this experiment, 82% of participants chose option B.

For choice 2, the player is again asked to choose one of two bets:

Under Bet C, the player wins:

- \$2500 with probability 33%
- \$0 with probability 67%

Under Bet D, the player wins:

- \$2400 with probability 34%
- 0 with probability 66%

Which do you prefer?

When Kahneman and Tversky (1979) ran this experiment, 83% of participants chose option C.

Let's examine this pair of preferences, with over 80% of experimental participants selecting B in Choice 1 and C in Choice 2.

According to expected utility theory, if an agent selects B, the expected utility of B must be greater than the expected utility of A. That is:

U(2400) > 0.33U(2500) + 0.66U(2400) + 0.01U(0)

We can simplify that to:

$$0.34U(2400) > 0.33U(2500) + 0.01U(0)$$

We can do the same analysis with the second choice. According to expected utility theory, if an agent selects C, the expected utility of C must be greater than the expected utility of D. That is:

$$0.33U(2500) + 0.67U(0) > 0.34U(2400) + 0.66U(0)$$

We can simplify that to:

$$0.33U(2500) + 0.01U(0) > 0.34U(2400)$$

This is a contradiction. The two inequalities point in opposite directions. Under expected utility theory, if an agent chooses A it should choose C. And if the agent chooses B, it should choose D.

Why does this occur? What axiom is being breached?

To understand this, I will show you another representation of the choices in this table. The left half of the table shows the bets for choice 1, and the right half for choice 2. Within each choice, the bets are represented as a payoff-chance pair. For example, I can read from the table that bet A involves a 66% chance of \$2400, a 1% chance of \$0, and a 33% chance of \$2500. Bet B involves a 100% chance of \$2400.

Choice 1					Choice 2			
	Α		в		С	D		
Payoff	Chance	Payoff	Chance	Payoff	Chance	Payoff	Chance	
\$2400	66%	\$2400	100%	\$0	67%	\$0	66%	
\$0	1%					\$2400	34%	
\$2500	33%			\$2500	33%			

I can then break up these payoff-chance pairs to create an equivalent representation as in this second table. I have split the outcomes in bets B and C. For example, I have written the 100% chance of \$2400 in option B as a 66% chance of \$2400 and a 34% chance of \$2400. I have written the 67% chance of \$0 in bet C as a 66% chance of \$0 and a 1% chance of \$0.

Choice 1					Choice 2			
A		В		C		D		
Payoff	Chance	Payoff	Chance	Payoff	Chance	Payoff	Chance	
\$2400	66%	\$2400	66%	\$0	66%	\$0	66%	
\$0	1%	\$2400	34%	\$0	1%	\$2400	34%	
\$2500	33%			\$2500	33%			

With this split, you can see that the bets in the bottom two rows of choice 1 and choice 2 are the same. Both choice 1 and choice 2 involve a choice between, in one bet, a 1% chance of nothing and a 33% chance of \$2,500 and in the other bet, a 34% chance of \$2,400.

That shared bet in choice 1 and choice 2 is paired with a 66% chance of the same payoff regardless of the preferred bet. For choice 1 that "common consequence" across bet A and bet B is \$2400. For choice 2, that common consequence across bet C and bet D is \$0.

Choice 1					Choice 2			
A		В		С		D		
Payoff	Chance	Payoff	Chance	Payoff	Chance	Payoff	Chance	
\$2400	66%	\$2400	66%	\$0	66%	\$0	66%	
\$0	1%	\$2400	34%	\$0	1%	\$2400	34%	
\$2500	33%			\$2500	33%			

This representation allows us to see that preferring bet B to bet A and bet C to bet D violates the axiom of independence. Under that axiom, two gambles mixed with a third gamble will maintain the same order of preference as when the two are presented independently of the third gamble. In this case, the two gambles are contained in the last two rows. The third gamble is the 66% chance of \$2400 or \$0. The third gamble is called a "common consequence" as the payoff is the same regardless of whether you choose A or B, or C and D.

I can express this in terms of the formal definition of the independence axiom. The formal definition states that if:

- x and y are lotteries with $x \succeq y$ and
- p is the probability that a third option z is present, then:

$$pz + (1-p)x \succeq pz + (1-p)y$$

For each of the choices in our lottery:

- x is a 1 in 34 chance of \$0 and a 33 in 34 chance of \$2500
- y is a 100% chance of \$2400
- z is \$2400 in choice 1 and \$0 in choice 2.

If p = 0, we simply have $x \succeq y$. For any non-zero value of p, such as the 66% in both choices, the preference between x and y should not change.

Here's another intuitive way to think about this bet.

Suppose I am going to generate one number between 1 and 100 randomly.

If a number between 1 and 66 is generated, you win the prize in the first row. If number 67 is generated, you win the amount in the second. If a number from 68 to 100 is generated, you win the sum in the third.

Suppose that you know that the number generated is between 1 and 66. Would you prefer bet A or B in choice 1? As you would win \$2400 with either choice, you will be indifferent. You will similarly be indifferent between bet C and D in choice 2, winning \$0 no matter what.

Suppose instead that a number between 67 and 100 is generated, but you don't know which. If you prefer A to B, you should also prefer C to D. In each choice, you effectively face the same bet. Let's assume for the moment that you prefer A and C.
Finally, suppose you don't know what number will be generated. We have just determined that if you know the ticket is between 1 and 66, you are indifferent between the options, but if between 67 and 100 is drawn, you prefer A and C. You do not prefer B or D when the ticket range is 1 to 66 or 67 to 100, so you should not prefer B or D when the ticket number is unknown.

However, the responses to the bets generated by Kahneman and Tversky (1979) and many other experimentalists suggest that when the number is unknown, the size of the common consequence for numbers 1 through 66 does matter. This common consequence is changing the preferences of the experimental participants.

10.2 Absurd rates of risk aversion

An important anomaly in expected utility theory concerns the level of risk aversion required to explain observed behaviour.

People reject many small-scale bets. The crux of the anomaly is that if we are expected utility maximisers, rejection implies that we would reject some highly favourable larger-scale bets so favourable almost no one would reject them.

Consider the following one-off bet involving the flip of a coin:

Head: You win \$110

Tail: You lose \$100

Suppose you reject. You would not be alone in doing this. There is ample evidence that people reject favourable low-stakes bets, even when they have material wealth. Barberis et al. (2006) described an experiment where they offered a 50:50 bet to win \$550 or lose \$500 to a group of wealthy experimental participants. These participants included clients of a bank's private wealth management division, with a median wealth above \$10 million. Seventy-one percent of the private wealth clients turned down the bet.

Under the axiom of diminishing marginal utility, we could conclude that you rejected it as you are risk averse. However, the minimum utility function curvature required to reconcile an expected utility maximiser declining bets of this nature when you hold any material level of wealth implies that you would reject immensely favourable bets.

Examples in Rabin (2000) and Rabin and Thaler (2001) illustrate this. Suppose a person who acts consistent with expected utility theory always turns down a 50:50 bet to win \$110 or lose \$100, whatever their wealth. That person will also turn down a 50:50 bet to win \$1 billion, lose \$1,000. Another expected utility maximiser turns down a 50:50 bet to win \$11, lose \$10 at all levels of wealth. That person will turn down any 50:50 bet where they could lose \$100, no matter the upside.

At face value, that is ridiculous. But that is the crux of the argument. Turning down a low-value bet with a positive expected value implies that the marginal utility of money must decline quickly for small changes in wealth. Rejection of a low-value bet to win \$110 or lose \$100 would lead to absurd responses to higher-value bets. This leads Rabin (2000) to argue that risk aversion or the diminishing value of money has nothing to do with the rejection of low-value bets.

The intuition behind the rejection of small-scale bets leading to the rejection of immensely favourable bets is as follows.

Suppose we have an expected utility maximiser with a weakly concave utility curve. That is, they are risk neutral or risk averse at all levels of wealth. If we drew their utility curve, it would always be increasing at a constant or decreasing rate, depending on where you were on the utility curve.



Figure 10.1: Weakly concave utility curve

This person rejects a 50:50 bet to gain \$150, lose \$100. (I have chosen this size bet so I can draw this visually to scale and make the point clear. The argument follows the same logic for a win \$110, lose \$100 bet.) We also assume this person would reject the bet regardless of their level of wealth at the time the bet was offered to them.

I can plot this on a chart. I will develop this chart step by step so that you can understand each element.

In Figure 10.2, the horizontal axis is wealth and the vertical axis is utility. I have marked the current wealth, W, and utility of that person's wealth, U(W).



Figure 10.2: Utility of current wealth

I can then mark in Figure 10.3 the two possible outcomes of the bet, the gain of \$150 (W + 150) and the loss of \$100 (W - 100). The utility of each outcome will be a point on these vertical lines.

The expected value of the bet is a gain of \$25. That is also marked.

As the person rejected the bet, the expected utility of the bet must be less than or equal to the utility of current wealth. The point on the vertical line at W + 25 where I mark expected utility must be level with or below the point on the vertical line at W where we mark current utility (U(W)). I have placed a dot on the W + 25 line in Figure 10.4 indicating the highest that the expected utility can be.

The expected utility of the bet is the probability-weighted utility of each of the two possible outcomes. Therefore, If I draw a line between the utility of each outcome, that line will pass through the expected utility of the bet.

I have added a line in Figure 10.5 between the utility for the two possible outcomes, with the expected utility of the bet on that line (or more specifically, in the middle of the line as each outcome has probability 50%).

We now have three points on the utility curve: current wealth and their utility



Figure 10.3: Outcomes of the bet



Figure 10.4: Expected utility of the bet



Figure 10.5: Utility of the outcomes of the bet

if the person lost or won the bet.

As the person is risk averse at all levels of wealth, we know that all parts of the utility curve are at a minimum weakly concave. We can therefore draw in Figure 10.6 part of the utility curve as the least risk averse they could be while still rejecting the bet. This is drawn in the solid black line.

From the rejection of the bet, we know that:

$$U(W+150) - U(W) \le U(W) - U(W-100)$$

This means that you value each dollar between W and W + 150, on average, by at most 100/150ths (or 2/3rds) as much as you value each dollar, on average, between W and W - 100. We can see this in the slope of the two black lines forming the part of the utility curve we have drawn. The second part has two thirds the slope of the first.

Further, by weak concavity, we can say that you value your W + 150th dollar at most two thirds as much as you value your W - 100th dollar. The relative value of each dollar has declined by at least one third over the span of \$250.

Let us now consider what would happen if this person had \$250 more wealth: that is, they have W + 250. They are then offered the same bet.



Figure 10.7: Utility curve

We have assumed that they will reject the bet at all levels of wealth, so they will also reject at this wealth. We can, therefore, infer another piece of the utility curve (or, more specifically, a curve for the least risk averse they could be). I have marked in Figure 10.8 the utility of their wealth, the highest the expected utility of the bet can be if they reject the bet, and the utility of the two possible outcomes of the bet. The next section of solid black line represents another part of the utility curve that is at the minimum risk aversion they could be while still rejecting the bet.



Figure 10.8: Utility curve extended

Iterating the previous calculations, I can say that they will weight their W + 400th dollar only two thirds as much as their W + 150th dollar. This means they value their W + 400th dollar only (2/3)2 or 4/9 as much as their W - 100th dollar. Or put another way, over the span of \$500, the relative value of each dollar has declined by five ninths.

As we infer additional pieces, we can see in Figure 10.9 that this person rapidly declines in the rate at which they place utility on further wealth. The increase in slope for each additional sum becomes less and less.

Over time, the slope approaches a horizontal asymptote: effectively, a cap on the level of utility they can achieve, however wealthy they may become. This is shown in Figure 10.10.

Keep iterating in this way and you end up with some ridiculous results. You







Figure 10.10: Utility curve asymptote

value the 2500th dollar above your current wealth only 2% as much as your last current dollar of your wealth - (2/3)10 - reducing by a constant factor of 2/3 every \$250. This is an absurd rate of discounting.

Taking this iteration to the extreme, it doesn't take long for additional money to have effectively zero value. In Figure 10.11, we see the utility curve approaching a limit. Any gain of money from W + 3000 upwards is valued at almost nothing, meaning a gain beyond that level, no matter how large, could compensate for the possibility of losing \$1,000.



Figure 10.11: Utility curve limit

10.3 Framing

Under expected utility theory, a person's choices should not be affected by how the options are described or by how their preferences are elicited.

Kahneman and Tversky (1984) reported the following experiment.

A group of experimental participants were shown the following:

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved.

If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved.

Which of the two programs would you favour?

72% of participants chose option A.

Another group of experimental participants were shown the following:

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program C is adopted, 400 people will die.

If Program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

Which of the two programs would you favour?

22% of participants chose option C.

72% of participants chose A and 22% of participants chose option C. Yet these two options are equivalent. The only difference is the framing of the options, which under expected utility theory should not matter.

10.4 Reference points

An auxiliary axiom of expected utility theory is that people use a reference point of zero wealth. They consider the utility of the absolute outcomes.

However, consider the following two scenarios:

- You have not checked your share portfolio in a while. You expect that the portfolio is worth around \$40,000. Today when you check, it is worth \$30,000. Do you feel rich or poor?
- You have not checked your share portfolio in a while. You expect that the portfolio is worth around \$20,000. Today when you check, it is worth \$30,000. Do you feel rich or poor?

Under expected utility theory, those two scenarios should feel the same as you have U(\$30,000) in both cases.

However, in the first case, you feel poor and in the second case you feel rich. This is because you are comparing the outcome to your reference point of \$40,000 in

the first case and 20,000 in the second case. You are not assessing the absolute outcome but appear to be using a reference point.

That is, $U(W) \neq U(W')$ when W = W' but reference points differ.

Chapter 11

Decision making under risk exercises

11.1 Expected value of roulette

You are playing roulette at the casino. There are 37 numbered pockets around the edge of the wheel (0 through 36). If you make a straight up bet on one of the 37 single numbers, you are paid \$35 for every dollar you bet (in addition to receiving back your bet). What is the expected value of a \$20 bet.

? Answer

The expected value of the Roulette bet is:

$$\begin{split} E[X] &= \sum_{i=1}^n p_i x_i \\ &= \frac{36}{37} \times (-\$20) + \frac{1}{37} \times (35 \times \$20) \\ &= \$ - 0.54 \end{split}$$

11.2 A bet or a certain payment?

Anika is an expected utility maximiser with the following utility function:

$$U(x) = \sqrt{x}$$

Anika is offered the following choice:

- A) A 50% chance of winning 10 and a 50% chance of winning nothing
- B) \$4 for certain

? Answer

Anika has zero wealth besides this offer.

a) What is the expected value of option A?

The expected value of option A) is:

$$E[A] = \sum_{i=1}^{n} p_i x_i$$
$$= 0.5 \times \$10 + 0.5 \times 0$$

= \$5

b) Will Anika choose A or B? Why?

? Answer

We need to determine the expected utility of each option. Anika will selection the option with the highest expected utility. The expected utility of option A) is:

$$EU(A) = p_1 U(x_1) + p_2 U(x_2)$$

= 0.5 × $\sqrt{10}$ + 0.5 × $\sqrt{0}$
= 1.58

The expected utility of option B) is:

$$EU(B) = U(4)$$
$$= \sqrt{4}$$
$$= 2$$

Anika will choose option B) as it gives her higher expected utility. Anika is risk averse.

c) What is the certainty equivalent of option A?

? Answer

To calculate the certainty equivalent of option A, we calculate what payment with certainty would deliver equivalent expected utility. That is:

EU(CE) = 1.58 $\sqrt{CE} = 1.58$ $CE = 1.58^{2}$ = 2.5

The certainty equivalent of option A is \$2.50. That is, Anika would be indifferent between option A and a payment of \$2.50 for certain.

d) Draw a graph showing Anika's utility curve, the expected value of option A, the expected utility of options A) and B) and the certainty equivalent of option A).



11.3 A 50:50 gamble

Consider the following gamble:

(0.5; \$550; 0.5, -\$500)

This gamble provides a 50% chance of winning \$550 and a 50% chance of losing \$500.

a) Would a risk neutral agent (who maximises expected value) be willing to pay \$20 to play this gamble? What is the most they would be willing to pay to play?

? Answer

The expected value of the gamble is:

$$\begin{split} E[X] &= \sum_{i=1}^n p_i x_i \\ &= 0.5(550) + 0.5(-500) \\ &= 25 \end{split}$$

This is greater than \$20, so a risk neutral agent will be willing to pay \$20 to participate in the gamble. The most they would be willing to pay is \$25.

We could also have solved this by determining the expected value if they had paid \$20:

$$\begin{split} E[X] - c &= \sum_{i=1}^n p_i x_i - c \\ &= 0.5(550) + 0.5(-500) - 20 \\ &= 5 \end{split}$$

As the expected value is positive, the agent would be willing to pay \$20.

b) Would a risk averse expected utility maximiser with wealth \$1000 and utility function $U(x) = x^{1/2}$ be willing to pay \$20 to play this gamble? What is the most they would be willing to pay to play?

? Answer

The expected utility of the gamble for the risk averse agent if they paid \$20 to play is:

$$\begin{split} EU(x) &= p_1(W+x_1-c) + p_2(W+x_2-c) \\ &= 0.5(1000+550-20)^{1/2} + 0.5(1000-500-20)^{1/2} \\ &= 30.51 \end{split}$$

The expected utility of not playing the gamble is:

$$EU(x) = (1000)^{1/2}$$

= 31.62

They would not pay 20 as they would have higher utility if they turned down the gamble.

In fact, they would not pay any positive sum to participate in the gamble. If they were offered the gamble for free, their expected utility would be:

$$EU(x) = 0.5(1000 + 550)^{1/2} + 0.5(1000 - 500)^{1/2}$$

= 30.87

This is less than if they simply turned down the gamble. They would be willing to pay to avoid the gamble. How much?

We can determine this by asking what wealth a utility of 30.87 is:

$$W^{1/2} = 30.87$$

 $W = 30.87^2$

= \$952.67

The certainty equivalent of the gamble is \$952.67. The agent would be willing to pay up to \$47.33 to avoid the gamble.

c) Would the expected utility maximiser with utility function $U(x) = x^{1/2}$ change their decision if they had \$1 million in wealth? Explain.

? Answer

If they now have 1 million in wealth, we simply repeat the calculations above with the new wealth.

$$EU(x) = 0.5(1000000 + 550 - 20)^{1/2} + 0.5(1000000 - 500 - 20)^{1/2}$$
$$= 1000.00247$$

The expected utility of not playing the gamble is:

$$EU(x) = (1000000)^{1/2}$$
$$= 1000$$

They would be willing to pay \$20 as they would have higher utility if they accepted the gamble.

What is the most they would be willing to pay? If they were offered the gamble for free, their expected utility would be:

$$EU(x) = 0.5(1000000 + 550)^{1/2} + 0.5(1000000 - 500)^{1/2}$$
$$= 1000.0125$$

How much would they be willing to pay for this opportunity? We can determine this by asking what wealth a utility of 1000.0125 is:

$$W = (1000.0124655)^2$$
$$= \$1000024.93$$

The agent would be willing to pay up to \$24.93 for the gamble. This is close to the expected value of \$25.

Intuitively, as the agent's wealth increases their utility function becomes increasingly linear (second derivative approaches zero) and they become closer to risk neutral.

11.4 Purchasing insurance

An agent is considering insurance against bushfire for their house valued at H = \$1,000,000. The house has a 1 in 1000 (p = 0.001) chance of burning down. An insurer is willing to offer full coverage for a premium (R) of \$1100.

a) What is the expected value of purchasing insurance?

? Answer

If the agent purchases insurance, they pay the premium and do not suffer any loss regardless of whether there is a bushfire or not.

$$E[I] = -R = -\$1,100$$

The expected value of purchasing insurance is the guaranteed loss of the premium, \$1100.

You could also think of the expected value of purchasing insurance as involving:

• in 1 in 1000 instances, the loss of the house and the insurance payout in case of fire, minus the cost of the premium

• in the other 999 in 1000 instances, payment of the premium only. In that case, you would write:

$$\begin{split} E[\mathbf{I}] &= p \times (-H + H - R) + (1 - p) \times (-R) \\ &= 0.001 \times (-1000000 + 1000000 - 1100) + 0.999 \times (-1100) \\ &= -\$1, 100 \end{split}$$

This gives the same answer as the first method.

b) What is the expected value of not purchasing insurance?

? Answer

j

If the agent does not purchase insurance, they face the 1 in 1000 possibility of an uninsured loss.

$$E[\neg I] = p \times (-H)$$

= -0.001 × 1000000
= -\$1000

c) Would a risk-neutral agent purchase the insurance?

? Answer

Above, we calculated that purchasing insurance in this case has a lower expected value than not purchasing the insurance. A risk-neutral agent would not purchase the insurance.

d) Suppose an agent has a logarithmic utility function $(U(x) = \ln(x))$ and they have \$10,000 in cash in addition to their house, giving them wealth (W) of \$1,010,000.

Are they risk seeking, risk neutral or risk averse?

? Answer

The logarithmic utility function has diminishing marginal utility. Diminishing marginal utility is the principle the marginal utility from each additional unit decreases. In the context of wealth, this means that each additional dollar provides less satisfaction than the previous one. Diminishing marginal utility means that the agent will be risk averse. They will prefer a certain outcome to a gamble with the same expected

value.

e) Would this agent with logarithmic utility purchase the insurance?

? Answer

We need to compare the expected utility of purchasing insurance with the expected utility of not purchasing insurance.

The expected utility of purchasing insurance is:

 $E[U(\mathbf{I})] = \ln(W - R)$ = $\ln(1,008,900)$ = 13.8244

The expected utility of not purchasing insurance is:

$$\begin{split} E[U(\neg \mathbf{I})] &= 0.999 \times \ln(W) + 0.001 \times \ln(W - H)) \\ &= 0.999 \times \ln(1,010,000) + 0.001 \times \ln(10,000) \\ &= 13.8208 \end{split}$$

The expected utility of purchasing insurance is greater than the expected utility from not purchasing insurance. This agent will purchase insurance. The following diagram illustrates. The agent's utility function is plotted, with the outcome on the horizontal axis and the utility of each outcome on the vertical axis. Each outcome and the utility of that outcome is marked: wealth after losing the house when uninsured (W - H), wealth after paying the insurance premium (W - R), and wealth if uninsured but the house does not burn down (W).

The expected utility of not purchasing insurance is on the dash-dot line between U(W - H) and U(W). The location of this point is determined by the probability p of incurring a loss. This point lies at a distance of p from U(W) along the line (or equivalently, at a distance of 1 - p from U(W - H)). This point aligns with the expected value of leaving the house uninsured $E[\neg I]$.

The utility of purchasing insurance (U(W - R)) is greater than the expected utility of not purchasing insurance $(E[U(\neg I)])$. The agent will purchase insurance.



Figure 11.2: Insurance choice by a risk averse expected utility maximiser

What is the intuition for this agent's purchase of insurance? Consider the following scenario. You will experience one of two outcomes with equal probability: \$0 in one case and \$200 in the other. Your average wealth in this scenario is:

$$\frac{\$0 + \$200}{2} = \$100$$

Your average utility (which is expected utility) is:

$$\frac{U(\$0) + U(\$200)}{2}$$

What if you can insure (costlessly) to give you wealth of \$100 with certainty? In that case, your average wealth is \$100, as per the first scenario, with utility of U(\$100).

Does insurance result in higher utility?

For a risk-averse person with diminishing marginal utility, the jump from \$0 to \$100 provides a larger increase in utility than the jump from \$100 to \$200. Therefore:

$$U(\$100) > \frac{U(\$0) + U(\$200)}{2}$$

The utility of average wealth is greater than the average utility of wealth. This inequality demonstrates why people with diminishing marginal utility purchase insurance. When wealth is unevenly distributed across potential outcomes, the utility gained from the best outcome is relatively less than the utility lost from the worst outcome. Insurance helps distribute wealth more evenly across all possible outcomes. In good times, the person has slightly less wealth (due to the premium). In bad times, the person avoids a catastrophic loss. This even distribution of wealth across possible outcomes results in higher expected utility for the consumer.

11.5 Insurance but not a lottery ticket

In your own words but using concepts from this subject explain why a risk averse agent who makes decisions according to expected utility theory might purchase insurance but not a lottery ticket.

? Answer

Both lotteries and insurance have a negative expected value.

The risk averse agent will typically reject a lottery as it has a small probability of a large win for the price of a small loss. Due to diminishing marginal returns, the average weight given to each dollar in the large gain is weighted much less than the average weight given to each dollar in the small price. This makes the lottery unattractive.

In contrast, a risk averse agent may purchase insurance as for a small price they can avoid the possibility of a large loss. Due to diminishing marginal returns, the large loss can have much higher average weight given to each dollar than to the weight given to each dollar for the small premium.

11.6 An anomaly in expected utility

Consider the following two choices:

Choice 1: Choose one of the following bets:

Bet A:

- \$10,000 with probability: 11%
- 0 with probability: 89%

Bet B:

- \$50,000 with probability: 10%
- \$0 with probability: 90%

Choice 2: Choose one of the following bets:

Bet A':

• \$10,000 with probability: 100%

Bet B':

- \$50,000 with probability: 10%
- \$10,000 with probability: 89%
- \$0 with probability: 1%

Many people pick B for Choice 1 and A' for Choice 2.

Does this pair of choices conform with Expected Utility Theory? Why?

? Answer

According to Expected Utility Theory, if an agent selects B:

0.10U(50,000) + 0.90U(0) > 0.11U(10,000) + 0.89U(0)

0.10U(50,000) + 0.01U(0) > 0.11U(10,000)

According to Expected Utility Theory, if an agent selects A':

U(10,000) > 0.10U(50,000) + 0.89U(10,000) + 0.01U(0)

0.11U(10,000) > 0.89U(50,000) + 0.01U(0)

This is a contradiction. Under expected utility theory, if an agent chooses B it should choose B'. And if the agent chooses A it should choose A'.

	89%	10%	1%
Α	\$0	\$10,000	\$10,000
В	\$0	\$50,000	\$0
A'	\$10,000	\$10,000	\$10,000
B'	\$10,000	\$50,000	\$0

This occurs due to a breach in the principle of independence. Here is a representation of the choices.

The bets in the two shaded areas are the same. They are paired with an outcomes of either \$10,000 or \$0. Preferring B to A and A' to B' is a violation of the axiom of the independence of irrelevant alternatives. Under that axiom, two gambles mixed with an irrelevant third gamble will maintain the same order of preference as when the two are presented independently of the third gamble.

Using this representation in the table, here is another way of understanding why this combination of choices is an anomaly. Imagine there are 100 tickets numbered 1 to 100. One ticket will be drawn. If a ticket between 1 and 89 is drawn, you win the prize in the first column. If a ticket between 90 and 99 is drawn, you win the amount in the second. If a 100 is drawn, you win the sum in the third.

Suppose that you know the ticket that is drawn is between 1 and 89. Would you prefer A or B? As you would win \$0 with either choice, you will be indifferent. You will similarly be indifferent between A' and B', winning \$10,00 no matter what.

Suppose instead that a ticket between 90 and 100 is drawn, but you don't know which. You can see that if you prefer A to B, you should also prefer A' to B'. In each choice you are effectively facing the same bet. Let's assume for the moment that you prefer B and B'.

Finally, suppose you don't know what ticket will be drawn. We have just determined that if you know the ticket is between 1 and 89 you are indifferent between the options, but if between 90 and 100 is drawn you prefer B and B'. You do not prefer A or A' when the ticket range is 1 to 89 or 90 to 100, so you should not prefer A or A' when the ticket number is unknown.

Finally, using the formal definition for the independence of irrelevant alternatives axiom:

- if x and y are lotteries with $x \succcurlyeq y$ and
- p is the probability that a third option z is present, then:

$$pz + (1-p)x \succcurlyeq pz + (1-p)y$$

For each of the choices in our lottery:

• p = 89%

- x is a 100% chance of \$10,000
- y is a 0.01/(1-0.89) chance of \$0 and 0.10/(1-0.89) chance of \$50,000
 z is \$10,000 in choice 1 and \$0 in choice 2, although z's value does not matter due to its assumed irrelevance.

Part III

Prospect theory

Prospect theory is a descriptive theory of decision-making under risk. It was developed by Daniel Kahneman and Amos Tversky (1979) as a challenge to expected utility theory.

Prospect theory has the following features:

First, it has a value function - that is a function that ascribes a value to each possible outcome. The value function incorporates reference dependence, loss aversion and the reflection effect:

- Reference dependence means that the value of an outcome is judged relative to a reference point.
- Loss aversion means the value of a loss is greater than the value of an equivalent gain.
- The reflection effect means that agents are risk-averse in the gain domain and risk-seeking in the loss domain

Second, prospect theory has a probability weighting function, whereby the agent subjectively weights objective probabilities. Small probabilities are weighted relatively more heavily than moderate probabilities. Those weights are then applied to the value of each outcome.

Prospect theory is an "as if" model of decision-making. People do not perform the calculations implicit in the application of prospect theory. Rather, they act "as if" they are performing those calculations. That makes prospect theory a descriptive theory, not a theory of how people actually make decisions.

The following sections break down each of these elements of prospect theory and explain how they are incorporated in its implementation.

Chapter 12

Reference dependence

Summary

- Prospect theory posits that people evaluate choices based on a reference point, coding outcomes as losses or gains relative to this point.
- Several theories exist for reference point formation, including status quo, lagged consumption, goals, and recent expectations.
- The status quo as a reference point works well in lab experiments but may not fit many economic interactions or markets with intangible goods.
- Lagged consumption as a reference point introduces the concept of adaptation, where people react to shocks initially, but the effect fades over time.
- Goals as reference points explain how people can react differently to the same outcome based on their set objectives.
- Recent expectations suggests that beliefs about future outcomes form the reference point, potentially aligning with other theories depending on the stability of expectations.

12.1 Introduction

Under prospect theory, people assess choices based on a reference point instead of an absolute assessment of their position. Outcomes are coded as losses and gains relative to that reference point.

Consider the following two scenarios:

• Scenario 1: You have not checked your share portfolio in a while. You expect it to be worth around \$40,000. Today when you check, it is worth \$30,000. Do you feel rich or poor?

• Scenario 2: You have not checked your share portfolio in a while. You expect it to be worth around \$20,000. Today when you check, it is worth \$30,000. Do you feel rich or poor?

Under expected utility theory, those two scenarios should feel the same as you have U(\$30,000) in both cases.

In contrast, under prospect theory, the value function - value function being what the utility function is typically called in prospect theory - applies to changes relative to the reference point.

If their initial reference point is their expectation, the value of that change is v(\$10,000) or v(-10,000), depending on whether their expectation is exceeded or not. The importance of that distinction becomes apparent when we examine how people consider choices involving either losses or gains.

12.2 Theories of reference dependence

There are several theories on reference point formation. These include:

- The status quo
- Lagged consumption
- Goals
- Recent expectations

12.2.1 The status quo

A common assumption in prospect theory is that the reference point is the status quo, as it was for many examples in the original prospect theory paper by Kahneman and Tversky (1979). The status quo implies a preference for the current state. Any negative change is perceived as a loss.

The status quo appears straightforward and is a reasonable description in many contexts, such as lab experiments.

However, the status quo as a reference point does not appear to be a useful assumption for describing many economic interactions. Suppose you decide to sell your bike. Do you see the foregoing of the bike as a loss?

What if you run a bike shop? Does every sale involve a feeling of loss? For markets where intangible and fungible goods are exchanged (for example, the stock market) the status quo assumption appears a poor fit.

12.2.2 Lagged consumption

A second theory is that the reference point is lagged consumption.

Imagine you win the lottery.

How do you feel one week after the draw?

How do you feel one year after the draw?

Your reference point likely reflects more recent consumption.

Lagged consumption introduces adaptation into reference point determination:

- 1. First, we react to shocks
- 2. Then the effect of the shock fades in time

12.2.3 Goals

Another theory is that our goals are our reference points.

Consider the following problem from Heath et al. (1999):

Sally and Trish both follow workout plans that usually involve doing 25 sit-ups.

One day, Sally sets a goal of performing 31 sit-ups. She finds herself very tired after performing 35 sit-ups and stops.

Trish sets a goal of performing 39 sit-ups. She finds herself very tired after performing 35 sit-ups and stops.

What emotion is each person experiencing?

With goals as reference points, people see success or failure to achieve a goal as a loss or gain. Although both Sally and Trish have the same performance, Sally will have a positive emotional reaction and Trish a negative reaction.

12.2.4 Recent expectations

A fourth theory of reference point determination is recent expectations.

In this approach, the reference point is beliefs about future outcomes (Kőszegi and Rabin (2006)). For example, a 5% pay rise when expecting 10% may be perceived as a loss.

The expectations-based theory can produce the same predictions as alternative theories:

- 1. If expectations are stable, recent expectations will reflect the status quo.
- 2. Recent consumption will shape expectations, making lagged consumption a reasonable reference point.
- 3. Goals can also shape (or be shaped by) expectations

Chapter 13

Loss aversion

Summary

- Loss aversion is the concept that losses have a stronger psychological impact than equivalent gains, often causing people to reject gambles with positive expected value.
- A value function with loss aversion can be represented mathematically, where losses are experienced with greater intensity than gains (e.g., twice the force).
- The endowment effect, where people value items they own more highly, is considered an empirical expression of loss aversion.
- Kahneman et al.'s (1991) experiment demonstrated the endowment effect with mugs, showing a significant difference between willingness to accept and willingness to pay for the same item.

13.1 Introduction

Loss aversion is the concept that losses loom larger than gains. People feel more strongly about a loss than an equivalent gain, so they are often willing to reject gambles with a materially positive expected value.

For example, if someone feels losses with twice the feeling of gains, a 50:50 bet to win \$550, lose \$500 will be unattractive. This provides an alternative explanation to the absurd levels of risk aversion required to reject this bet (as discussed in Section 10.2).

13.2 A value function with loss aversion

This equation is an example of a value function with loss aversion:

$$v(x) = \begin{cases} x & \text{where} x \ge 0\\ 2x & \text{where} x < 0 \end{cases}$$

x is the outcome relative to the reference point.

In this value function, losses are experienced with twice the force of gains, with each loss multiplied by a factor of two.

As an example, suppose someone is given \$100. If their initial reference point is their wealth before receiving the \$100, x will be \$100. Therefore their change in value is +100. If the same person instead loses \$100, their change in value would be -200.

This plot shows the increased effect of the loss under this value function:



The greater slope of the curve in the loss domain, leading to a kink where the axes intercept, is indicative of the greater effect of losses.

13.3 The endowment effect

The endowment effect is often used to illustrate loss aversion.

Kahneman et al. (1991) ran one of the most famous and replicated experiments in economics.

They randomly assigned a free mug to members of a group and asked how much money they would accept for returning the mug (i.e. willingness to accept). The remaining participants were only shown a mug and asked about their willingness to pay for the mug.

Kahneman et al. (1991) found that the willingness to accept (\$5.75) was substantially higher than the willingness to pay (\$2.25).

The endowment effect is this phenomenon where people impute additional value to the items they own. The endowment effect is argued to be an empirical expression of loss aversion. Willingness to accept is higher as it is payment to incur the loss of the mug.

The endowment effect has been found in real estate markets, the stock market, with basketball tickets and in other domains.

13.3.1 Endowment effect example

Bruce has the following reference-dependent value function:

$$v(x) = \begin{cases} x & \text{where} x \ge 0\\ 2x & \text{where} x < 0 \end{cases}$$

x is the outcome relative to the reference point.

Assume Bruce has preferences over money (m) and mugs (c) as in this value function:

$$V(x) = v(m - r_m) + v(5c - 5r_c)$$

 r_m is Bruce's reference point as it relates to money and r_c is his reference point as it relates to mugs.

To illustrate how this value function works, imagine Bruce has two mugs, and he drops one. It breaks. His change in the value function is:

$$\begin{split} V(x) &= v(5c-5r_c) \\ &= v(5\times 1-5\times 2) \\ &= v(-5) \\ &= -10 \qquad (\text{as } v(x) = 2x \text{ when } x < 0) \end{split}$$

The loss of a mug results in value of -10.

At the beginning of the experiment, Bruce is given a mug. Assuming Bruce's reference point adapts such that he considers the mug his, how much would Bruce need to be paid to give up the mug?

We can calculate this by calculating what payment p would make Bruce indifferent between losing the mug and gaining p. That is the point where, after losing the mug and receiving payment, Bruce's change in value is equal to zero.

We assume a reference point for money of his current wealth and are instructed that his reference point for mugs is ownership of the mug.

$$\begin{split} V(x) &= 0 \\ V(x) &= v(m-r_m) + v(5c-5r_c) \\ &= v(W+p-W) + v(5\times 0 - 5\times 1) \\ &= v(p) + v(-5) \\ &= p-2\times 5 \\ &= p-10 \\ p &= 10 \end{split}$$

Bruce would need to be paid at least \$10 to give up the mug. Giving up the mug is seen as a loss and given greater weight than the money gained.

Now assume Bruce was not given a mug, but rather an opportunity to purchase a mug. How much would Bruce be willing to pay for the mug?

We can calculate this by calculating what payment p would make Bruce indifferent between gaining the mug and losing p. That is the point where, after receiving the mug and making payment, Bruce's change in value is equal to zero.

We assume a reference point for money of his current wealth and are instructed that his reference point for mugs is no mug.

$$\begin{split} V(x) &= 0 \\ V(x) &= v(m-r_m) + v(5c-5r_c) \\ &= v(W-p-W) + v(5\times 1-5\times 0) \\ &= v(-p) + v(5) \\ &= -2p+5 \\ p &= 2.5 \end{split}$$

The most Bruce would be willing to pay for the mug is \$2.50. He sees the foregone money as a loss, giving it greater weight than the mug gained.

Chapter 14

The reflection effect

Summary

- The reflection effect is an asymmetry in risk preferences: people are typically risk-averse for gains but risk-seeking for losses.
- This effect can be illustrated through framing experiments, such as the Asian Disease problem, where equivalent outcomes are perceived differently when framed as gains versus losses.
- The reflection effect can be represented in a value function with diminishing sensitivity to both gains and losses, resulting in a concave function in the gain domain (indicating risk aversion) and a convex function in the loss domain (indicating risk-seeking behaviour).

14.1 Introduction

The reflection effect involves an asymmetry in risk preferences in the gain and loss domains.

When people make a risky choice related to gains, they are risk averse. They prefer a certain option with value lower than the expected value of the risky choice. When choosing an option in the loss domain, they become risk-seeking. This phenomenon is called the reflection effect.

Put simply, people typically prefer to lock in a sure gain rather than risk it for more, but will take risks to avoid a sure loss.

The reflection effect might also be thought of as diminishing sensitivity to gains or losses in either direction. This contrasts with expected utility theory, where the pain of losses increases as they grow in size.

14.2 The Asian Disease problem

The reflection effect explains the framing effects in the following experiment by Kahneman and Tversky (1984).

One group of experimental subjects were asked the following hypothetical question that would be unlikely to be asked post-Covid-19.

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved.

If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved.

Which of the two programs would you favour?

72% of participants chose option A.

Another group of experimental participants were shown the following:

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program C is adopted, 400 people will die.

If Program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

Which of the two programs would you favour?

22% of participants chose option C.

Despite their differing presentation, these two scenarios are mathematically identical. Programs A and C both result in 200 people living and 400 dying. Similarly, Programs B and D offer a 1 in 3 chance of saving everyone and a 2 in 3 chance of everyone dying.

The difference is that the first scenario presents a 'gain frame', focusing on lives saved. The second scenario presents a 'loss frame', focusing on lives lost." In the gain frame, 72% chose the certain option A over the risky option B. In the loss frame, only 22% chose the certain option C, meaning 78% preferred the risky option D.
This shift in preferences demonstrates the reflection effect. When the choice is framed in terms of gains - lives saved - people prefer the certain option. When framed in terms of losses - lives lost - they prefer the risky option.

14.3 The reflection effect in the value function

The following value function is an example of a function with diminishing sensitivity to both gains and losses. This function generates the reflection effect.

$$v(x) = \begin{cases} x^{\frac{1}{2}} \text{where} x \ge 0\\ -(-x)^{\frac{1}{2}} \text{where} x < 0 \end{cases}$$

As x increases in magnitude in either direction, the marginal change in value from each incremental unit of x decreases. This decreasing marginal change is implemented via the power of one half. For the value of the gain or loss to double, the size of the gain or loss must quadruple.

When x is less than zero, we apply a negative sign to x as the power of one half is only defined for positive numbers. We then apply another negative sign after applying the power to indicate the negative value of losses.

This value function results in risk-averse behaviour in the gain domain and riskseeking behaviour in the loss domain. The following plot shows the diminishing effect in each direction.

In the gain domain, the function is concave, indicating risk aversion. In the loss domain, the convex function indicates risk-seeking behaviour. The result is an S-shaped curve.

14.4 An example

The following numerical example illustrates further.

Suppose an agent has a value function that expresses the reflection effect, with the outcomes in either direction subject to a power of one half.

$$v(x) = \begin{cases} x^{\frac{1}{2}} \text{where} x \ge 0\\ -(-x)^{\frac{1}{2}} \text{where} x < 0 \end{cases}$$

This agent is offered a choice between \$10 for certain and a 50:50 bet to win \$20 or end up with nothing. We calculate the value of each option by substituting the outcomes into the value function.

For the certain gain, we substitute \$10 into the value function.



Figure 14.1: The reflection effect

$$v(\text{certainty}) = v(10)$$
$$= 10^{\frac{1}{2}}$$
$$= 3.16$$

For the bet, we calculate the probability weighted value of the potential outcomes.

$$v(bet) = 0.5 \times v(20) + 0.5 \times v(0)$$
$$= 0.5 \times 20^{\frac{1}{2}} + 0.5 \times 0$$
$$= 2.24$$

The \$10 for certain has a higher value for the agent. This agent is risk averse in the gain domain and therefore prefers an amount for certain over a bet with the same expected value. The following chart illustrates. Plotted are the value of the certain payment of \$10 and the outcome from winning the gamble, \$20. A loss results in value of 0.

The value of the gamble itself is the probability-weighted average of the two gamble outcomes. Due to diminishing marginal utility in the gain domain, the value of the gamble is less than the value of \$10 with certainty. The extra \$10 over the certain outcome from winning the bet is less than the value of the first \$10. As a result, the agent does not want to risk the bet.



Figure 14.2: The reflection effect in the gain domain

Suppose the agent is now offered another choice. They can now have a certain loss of 10 or a 50:50 bet to lose 20 or to lose nothing.

Again, we calculate the value of each option by substituting the outcomes into the value function.

For the certain loss, we substitute -\$10 into the value function.

$$v(\text{certainty}) = v(-10)$$
$$= -10^{\frac{1}{2}}$$
$$= -3.16$$

For the bet, we calculate the probability weighted value of the potential outcomes.

$$\begin{split} v(\text{bet}) &= 0.5 \times v(-20) + 0.5 \times v(0) \\ &= -0.5 \times 20^{\frac{1}{2}} + 0.5 \times 0 \\ &= -2.24 \end{split}$$

This bet delivers higher value than the certain loss, despite the bet and the certain loss having the same expected value. The agent is willing to take a risk to avoid a loss. They are risk seeking in the loss domain.

The following chart illustrates. Plotted are the value of the certain payment of -\$10 and the outcome from losing the gamble, -\$20. A win results in value of 0.

The value of the gamble itself is the probability-weighted average of the two gamble outcomes. Due to diminishing marginal utility in the loss domain, the value of the gamble is greater than the value of -\$10 with certainty. The potential loss of another \$10 over and above the certain loss is given less weight than the first \$10. As a result, the agent wants to take the risk.



Figure 14.3: The reflection effect in the loss domain

Chapter 15

The value function

Summary

- Prospect theory's value function incorporates three phenomena: reference dependence, loss aversion, and the reflection effect.
- The function is defined on changes relative to a reference point, rather than final wealth levels, with diminishing marginal sensitivity to both gains and losses.
- The value function is characterised by a kink at the reference point, with losses weighted more heavily than gains.

15.1 Introduction

The three phenomena, reference dependence, loss aversion and the reflection effect, are each incorporated into the prospect theory value function.

First, the value function of prospect theory is defined on changes in wealth or welfare rather than on final wealth levels. In other words, gains and losses are defined relative to a reference point. The reference point might be the status quo, lagged consumption, a goal, or an expectation about the outcome.

Utility from an outcome depends on the distance to the relevant reference level. A multi-millionaire may reject a 50:50 bet to win \$550, lose \$500 because they are not comparing the outcomes to their total wealth but rather are making a judgment relative to the reference point of the status quo.

Second, perceptions of both gains and losses are characterised by diminishing marginal sensitivity in either direction. Successive incremental changes have a smaller marginal impact.

This is similar to decreasing marginal utility of wealth in expected utility theory. Prospect theory and expected utility theory differ in the baseline. In expected utility theory, the starting value is typically zero wealth, with increases from there decreasing in marginal utility. In prospect theory, the starting value is the reference point, with both increases and decreases having smaller marginal effects as they increase in magnitude.

Third, losses loom larger than gains. People feel more strongly about a loss than they do an equivalent gain. They are often willing to reject gambles with a materially positive expected value.

These phenomena result in the following famous figure from Kahneman and Tversky (1979).

- 1. The value function is centred on the reference point at the origin
- 2. There is a kink at the origin, with losses counting more than gains.

3. There is diminishing sensitivity to further changes from the reference point in both directions.



FIGURE 3.—A hypothetical value function.

The following diagrams illustrate the different shapes of the expected utility function and the Prospect Theory value function.

The expected utility function, in this case U(x) = ln(x), has diminishing marginal utility as utility increases. Utility is measured from a reference point of zero.



Figure 15.1: Expected utility function

For the prospect theory function, you can see the kink at zero, with losses weighted more heavily than gains, with gains and losses determined relative to a reference point. There is diminishing sensitivity to further changes in both directions

The following equation is an example of a value function incorporating both loss aversion and the reflection effect.

$$v(x) = \begin{cases} x^{\frac{1}{2}} \mathrm{where} x \ge 0\\ -2(-x)^{\frac{1}{2}} \mathrm{where} x < 0 \end{cases}$$

Losses are multiplied by 2, giving losses twice the weight of an equivalent gain in the value function. This is loss aversion.

The reflection effect is implemented through losses and gains being to the power of one-half. This leads to diminishing sensitivity to changes in both directions.



Figure 15.2: Prospect theory value function

Chapter 16

Probability weighting

Summary

- Prospect theory incorporates non-linear weighting of probabilities, contrasting with expected utility theory's linear approach.
- Empirical evidence shows that people tend to overweight certain events and very low-probability events (the certainty effect and possibility effect) while approximating linear weights for intermediate probabilities.
- This probability weighting helps explain phenomena like the Allais Paradox, where people's choices in different scenarios appear inconsistent with expected utility theory.
- The decision weight function, such as Prelec's (1998) formula, mathematically represents this non-linear probability weighting in prospect theory, capturing the overweighting of small probabilities.

16.1 Introduction

In expected utility theory, probabilities enter the expected utility function linearly. For example, if an event is twice as likely as another outcome, it has double the weight.

In contrast, prospect theory incorporates non-linear weighting of probabilities by applying decision weights to each potential outcome.

16.2 Empirical evidence

Experimental observation indicates that we approximate linear weights for intermediate probabilities when making decisions under risk. But, there is strong evidence that we overweight certain events, when the probability of the event is one, relative to near certain events, such as when the probability is, say, 99%. This overweighting of certainty is effectively the same as overweighting very low-probability events.

The following diagram from Tversky and Kahneman (1992) illustrates the relationship between objective probability and the decision weight applied to each outcome. On the x-axis is the probability of the outcome. On the y-axis is the weight applied to the value function for that probability. The straight line at 45 degrees represents linear weighting of probabilities. The curve represents the weighting function.



For this particular curve, where the probability is very low, such as around p = 0.05, the weight is around 0.15. Similarly, at high probability, such as p = 0.95, the weight is around 0.8. For intermediate probabilities, the observed weight is relatively closer to the objective probability.

Kahneman (2011) calls the large psychological value of the change from 0 to 5% (or some other small probability) the possibility effect. He calls the large psychological value of the change to 100% the certainty effect. We will pay a lot more for certainty than near certainty.

16.3 The Allais Paradox

Probability weighting is often offered as an explanation for the Allais Paradox, which I discuss in Section 10.1.

The Allais paradox can be illustrated as follows.

You are given the following pair of choices.

Choice 1: Choose one of the following bets:

Bet A:

- \$2500 with probability: 33%
- \$2400 with probability: 66%
- \$0 with probability: 1%

Bet B:

• \$2400 with probability: 100%

People tend to prefer Bet B.

Choice 2: Choose one of the following bets:

Bet C:

- \$2500 with probability: 33%
- \$0 with probability: 67%

Bet D:

- \$2400 with probability: 34%
- \$0 with probability: 66%

People tend to prefer Bet C.

It can be shown that this pair of preferences, Bet B and Bet C, does not conform with expected utility theory.

One explanation for this pair of decisions comes from probability weighting. If you look at Bet B, the outcome is certain. Certain events tend to be overweighted relative to near-certain events, such as the 99% chance of \$2400 or \$2500 in Bet A. An alternative way of thinking about this is that the 1% probability of nothing in Bet A is overweighted.

Conversely, the intermediate probabilities in Bet C and Bet D are weighted closer to linearly, which can result in the slightly higher expected value Bet C being preferred.

16.4 The decision weight

The weighting of probabilities is applied in prospect theory through a decision weight $\pi(p_i)$. The decision weight is a function of the probability of the outcome.

This decision-weighting function reflects the empirical regularity that people overweight certain events relative to near-certain events and overweight lowprobability events. An example probability weighting function of a type proposed by Prelec (1998) is as follows:

$$\pi(p) = e^{-(-ln(p))^{\alpha}} o < \alpha < 1$$

This function, with $\alpha = 0.6$, is plotted below.



Figure 16.1: Probability weighting function

Chapter 17

Prospect theory implementation

17.1 Summary

- Prospect theory involves two phases: editing and evaluation.
- The editing phase simplifies prospects through four main operations: coding (as gains or losses), combination (of identical outcomes), segregation (of riskless components), and cancellation (of shared components).
- In the evaluation phase, prospects are assessed using a formula that combines decision weights applied to probabilities and subjective values of outcomes.
- Prospect theory results in a fourfold pattern of risk attitudes: risk aversion for high-probability gains and low-probability losses, and risk seeking for low-probability gains and high-probability losses.
- This pattern explains real-world behaviours such as settlement negotiations in court cases, the attractiveness of lotteries despite typically low expected values, and the purchase of insurance against unlikely events.

17.2 Introduction

Under prospect theory, people assess the weighted value of a prospect in two phases: editing and evaluation.

17.3 Editing

Editing involves simplification of prospects for subsequent evaluation.

Kahneman and Tversky (1979) describe the editing phase as having four main operations: coding, combination, segregation and cancellation.

- In the coding operation, prospects are coded as gains or losses relative to a reference point.
- In the combination operation, prospects are simplified by combining probabilities for identical outcomes. For example, (0.25, 200; 0.25, 200; 0.50, 0) will become (0.50, 200; 0.50, 0).
- During segregation, riskless components are segregated from risky components. For example (0.80, 300; 0.20, 200) corresponds to a sure gain of 200 and the risky gamble (0.80, 100; 0.20, 0).
- Finally, in cancellation, components that are shared by two prospects are ignored.

17.4 Evaluation

In the evaluation phase, the prospects are evaluated, and the option with the highest weighted value is chosen.

The weighted value of a prospect is made up of:

- 1. a decision weight applied to each probability $\pi(p_i)$
- 2. the subjective value of each outcome $v(x_i)$

These are applied through the following formula to calculate V(X), the weighted value of the outcomes from gamble X.

$$\begin{split} V(X) &= \sum_{i=1}^n \pi(p_i) v(x_i) \\ &= \pi(p_1) v(x_1) + \pi(p_2) v(x_2) + \ldots + \pi(p_n) v(x_n) \end{split}$$

17.5 Fourfold pattern of risk attitudes

Prospect theory results in a four-fold pattern of risk attitudes, as shown in this table.

	Gains	Losses
High probability	Risk aversion	Risk seeking
Low probability	Risk seeking	Risk aversion

For high-probability gambles, the reflection effect leads to people being people risk averse in the domain of gains and risk seeking in the domain of losses. For the top left quadrant, the low possibility of missing out on the gain is overweighted, making the gamble less attractive and amplifying the risk-averse behaviour. For the top right quadrant, the low probability of avoiding the loss is also overweighted, amplifying the risk-seeking behaviour we see in the domain of losses.

These top two quadrants can be illustrated by considering an offer to settle a court case.

Imagine one party has a 95% chance of winning a large settlement. The shape of the value function in the gain domain and the certainty effect make a settlement offer attractive. Conversely, the other party overweights their 5% chance of victory and is risk-seeking in the loss domain, making them unlikely to seek settlement unless it is very favourable.

But for low-probability gambles, as in the bottom two quadrants, the probability weighting pushes decisions in the opposite direction to the value function. The possibility of a gain is overweighted, making gambles attractive and inducing risk-seeking behaviour. A similar effect occurs for a low probability of loss, with the overweighted probability making the gamble less attractive, inducing risk-averse behaviour.

The bottom two quadrants can be related to the purchase of insurance and lotteries.

Lotteries involve a small probability of gains. As people overweight that small probability of a win, people will be risk-seeking when considering whether to purchase a lottery despite the gamble being in the gain domain where they are typically risk averse.

Insurance involves a small probability of loss. As people overweight the small probability of a loss, people will be risk-averse when considering whether to purchase insurance despite normally being risk seeking in the loss domain.

Chapter 18

Prospect theory examples

18.1 Mathematical examples

Summary

- In the following mathematical examples, I demonstrate the operation of prospect theory through various scenarios, illustrating how individuals make decisions under uncertainty based on reference points, loss aversion, and diminishing sensitivity to gains and losses.
- The 50:50 gamble demonstrates loss aversion, with losses weighted more heavily than equivalent gains.
- The 60:40 gamble shows how diminishing sensitivity and loss aversion interact to influence decision-making.
- The gamble in the gain domain illustrates the reflection effect, where individuals are risk-averse for gains but become risk-seeking for losses.
- Changes in wealth or reference points can significantly alter risk preferences, as shown in the examples of accepting or rejecting gambles after losses or wins.

In this section, I present a series of mathematical examples of prospect theory.

18.1.1 A 50:50 gamble

Suppose Alby has the following reference-dependent value function:

$$v(x) = \begin{cases} x^{\frac{1}{2}} & \text{where} x \ge 0\\ -2(-x)^{\frac{1}{2}} & \text{where} x < 0 \end{cases}$$

x is the realised outcome relative to the reference point.

Assume that Alby's reference point is the status quo and that he weights outcomes linearly.

Alby is offered the gamble A:

$$(0.5,\$110;0.5,-\$100)$$

Accept or reject?

Will Alby want to play this gamble?

To determine this, we compare the weighted value of the gamble with the weighted value of rejecting the offer.

The weighted value of the gamble is:

$$\begin{split} V(A) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.5 \times v(110) + 0.5 \times v(-100) \\ &= 0.5 \times (110)^{\frac{1}{2}} - 0.5 \times 2 \times (100)^{\frac{1}{2}} \\ &= -4.76 \end{split}$$

Alby will not want to play this gamble as it has a negative value. He could receive a weighted value of 0 by simply not playing.

The reason for this negative value is that Alby is loss averse. The loss of \$100 is given twice the weight of an equivalent gain.

The following chart illustrates. The horizontal axis is the outcome x relative to the reference point. The vertical axis is the value of each outcome. The S-shaped curve represents Alby's value function, with a concave curve in the game domain and a convex curve in the loss domain. The curve is steeper in the loss domain, due to loss aversion.

The gain of \$110 and the loss of \$100 and their respective values are marked.

I have drawn a straight line between the two outcomes. The weighted value of the two outcomes will be on this line.

As the probability of each outcome is 50 percent, the expected value of A, 5, is halfway between the two outcomes. If we project to the straight line from the expected value, we get V(A), a probability-weighted average of the two possible outcomes from the bet. V(A) is negative, indicating that the gamble has a lower weighted value than remaining at the status quo.



Figure 18.1: A 50:50 gamble

Accept or reject after loss?

Suppose Alby loses his wallet containing \$100. He feels bad about it and perceives it as a loss. His reference point is unchanged at the original status quo, but the amount of money he will have after any outcome is \$100 less than otherwise. Would he be willing to take gamble A now?

After losing \$100 but not changing his reference point, he has two possible outcomes relative to his reference point: a gain of \$10 (winning \$110 minus the lost money in the wallet) and a loss of \$200 (losing \$100 and also losing his wallet).



The weighted value of gamble A is now:

$$\begin{split} V(A) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.5 \times v(110 - 100) + 0.5 \times v(-100 - 100) \\ &= 0.5 \times (10)^{\frac{1}{2}} - 0.5 \times 2 \times (200)^{\frac{1}{2}} \\ &= -12.56 \end{split}$$

The value of not playing the gamble involves remaining with a loss of \$100:

$$V(\neg A) = v(-100)$$

= $-2 \times (100)^{\frac{1}{2}}$
= -20

He will now want to play the gamble as it has a greater value than staying with his current loss. The gamble becomes attractive as it allows recovery of the loss. Alby is risk seeking in the loss domain. (He would even accept a 50:50 gamble to win or lose \$100 with an expected value of zero.)

The choice is illustrated in the following chart. The two possible outcomes, \$200 below the reference point and \$10 above the reference point, plus their values, are marked. The weighted value of those two possible outcomes is also marked, with the expected value, E of A, projected onto the line between the two outcomes to give V(A)

It is visually apparent that V(A) is above the value of a loss of \$100, the outcome if Alby does not accept the gamble. As a result, Alby wants to accept the gamble.



Figure 18.2: A 50:50 gamble after a loss

Accept or reject after adaptation to loss?

Alby has now adapted to his loss of \$100. The new reference point is the new wealth level incorporating the loss wallet. Would he take gamble A now?

We are now back to an identical situation as when he was first offered the gamble with his reference point as the status quo. He will not want to partake in the gamble.

Accept or reject after a win?

Alby wins 10,000 at the casino. He feels good about his win, so his reference point remains at his wealth excluding the win. Would he take gamble A now?

With the additional \$10,000, the value of the gamble is:

$$\begin{split} V(A) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.5 \times v(10000 + 110) + 0.5 \times v(10000 - 100) \\ &= 0.5 \times (10110)^{\frac{1}{2}} + 0.5 \times (9900)^{\frac{1}{2}} \\ &= 100.02 \end{split}$$

The value of not playing the gamble is:

$$V(\neg A) = v(10000)$$

= 10000^{1/2}
= 100

The gamble is now attractive. Alby is less risk averse at a higher wealth. Further, the gamble is entirely in the gain domain, meaning that loss aversion does not affect the decision.

The following chart, not drawn to scale, illustrates. Alby becomes increasingly risk neutral as we move further into the gain domain. You can see this through the value function curve becoming approximately straight. As a result, at a high enough wealth, the positive value bet becomes attractive. V(A) is above the value of the certain outcome, V(10000).



Figure 18.3: A 50:50 gamble after a win (not to scale)

18.1.2 A 60:40 gamble

Paddy makes decisions in accordance with prospect theory, has wealth 300 and value function:

$$v(x) = \begin{cases} x^{\frac{1}{2}} & \text{where} \quad x \geq 0 \\ -2(-x)^{\frac{1}{2}} & \text{where} \quad x < 0 \end{cases}$$

Assume Paddy weights probabilities linearly.

Paddy is offered the following bet A:

- a 60% probability to win \$150
- a 40% probability to lose \$100.

Accept or reject

Does Paddy accept bet A?

Paddy compares the value of taking versus not taking the bet:

$$\begin{split} V(\mathbf{A}) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.6 \times v(150) + 0.4 \times v(-100) \\ &= 0.6 \times (150)^{\frac{1}{2}} - 0.4 \times 2 \times (100)^{\frac{1}{2}} \\ &= -0.652 \end{split}$$

The value of not taking the bet is zero. Paddy would have no change from his reference point.

Paddy rejects the bet as V(A) is less than the V(0) = 0 that Paddy could get by simply rejecting the bet. He rejects the bet due to his loss aversion and diminishing sensitivity to gains. The loss is weighted double that of an equivalent gain, outweighing both the larger potential gain and 60% probability.

The following figure shows Paddy's value function, the bets and the value of the bets. The figure illustrates that Paddy's rejection is caused by both Paddy's loss aversion and his diminishing sensitivity in the gain domain, which has a larger effect than the diminishing sensitivity in the loss domain due to the larger magnitude of the potential gain.

Accept or reject after a loss?

Following some bad economic news, Paddy's wealth declines to \$150. Paddy cannot get over the loss, so his reference point remains his former wealth of \$300.

Paddy is offered bet A again. Does Paddy accept the bet?



Figure 18.4: A 60:40 gamble

As Paddy is now in the loss domain, the two potential outcomes from the bet are a gain of \$0 and a loss of \$250. His alternative is remaining at a point \$150 below his reference point.



Paddy compares the value of taking versus not taking the bet:

$$V(\mathbf{A}) = p_1 v(x_1) + p_2 v(x_2)$$

= 0.6 × v(-150 + 150) + 0.4 × v(-150 - 100)
= 0.6 × (0)^{1/2} - 0.4 × 2 × (250)^{1/2}
= -12.649
$$V(\neg \mathbf{A}) = v(-150)$$

= -2 × (150)^{1/2}
= -24.495

Paddy accepts the bet as V(A) is greater than the value of the certain loss of \$150.

The following figure shows Paddy's value function, the bets and the value of the bets. The figure shows that Paddy accepts the bet as he is risk seeking in the loss domain. The potential loss of another \$100 results in a smaller incremental loss of value than an equivalent win of \$100.



Figure 18.5: A 60:40 gamble after a loss

18.1.3 A gamble in the gain domain

Suppose Bill has the following reference-dependent value function:

$$v(x) = \begin{cases} x^{1/2} & \text{where } x \ge 0\\ -2(-x)^{1/2} & \text{where } x < 0 \end{cases}$$

x is the change in Bill's position relative to his reference point.

What feature of Bill's value function leads to the reflection effect?

The power of $\frac{1}{2}$ applied in both the gain and loss domain leads to diminishing sensitivity to gains and losses. The value function is concave in the gain domain and convex in the loss domain. This curvature leads to risk-averse behaviour in the gain domain and risk-seeking behaviour in the loss domain. This change in risk preference between the gain and loss domains is the reflection effect.



Figure 18.6: Bill's value function

Accept or reject?

Bill considers a choice between \$100 for certain and gamble A: (0.5, \$250; 0.5, 0).

Will Bill prefer the \$100 or gamble A?

The weighted value of gamble A is:

$$\begin{split} V(A) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.5 \times v(250) + 0.5 \times v(0) \\ &= 0.5 \times 250^{0.5} + 0.5 \times 0^{0.5} \\ &= 7.91 \end{split}$$

The value of the \$100 for certain is:

$$V(\$100) = v(100)$$

= 100^{0.5}
= 10

As V(\$100) > V(A), Bill will prefer the \$100 for certain to the gamble.

The possible outcomes from the gamble are zero and \$250. The certain outcome on offer is \$100. The expected value of the gamble is \$125. The outcomes are in the gain domain.

As he is risk averse in the gain domain, the value of the \$100 for certain exceeds the weighted value of the gamble. This can be seen through v(A) being less than v(\$100). Bill will therefore choose the \$100 for certain.

The following chart shows Bill's choices. Bill rejects the gamble because of the diminishing sensitivity to gains. This leads him to be risk averse and reject the higher expected value option of the gamble.

As all possible outcomes under our assumed reference point are in the gain domain, loss aversion does not affect his decision. Note that we do not use the value function for x < 0 in determining Bill's choice.

Accept or reject after a negative shock?

Suppose Bill were to experience a large negative shock to his wealth that does not immediately change his reference point. Could this shock cause him to change his decision concerning the \$100 and gamble A?

A large negative shock to Bill's wealth would cause him to change his decision concerning the \$100 and gamble A. The shock would move the two possible outcomes into the loss domain, where Bill is risk seeking. (For this answer, I am assuming a shock of greater than \$250. A smaller shock would change the analysis.)



Figure 18.7: Bill's consideration of gamble A and the \$100

The following diagram illustrates the outcomes relative to the reference point. Let L be a large negative number, the loss. The potential outcomes from the gamble are now L and L + 250. The certain outcome of accepting the \$100 is L + 100.



The following diagram illustrates Bill's decision after the shock. The outcomes L, L + 100 and L + 250, and their respective values, are marked. The expected value of the gamble is L + 125. The weighted value of the gamble is V(L + A).

Due to the convex curvature of the curve in the loss domain, Bill is risk seeking. As a result, the utility of the gamble is greater than the utility of the certain

outcome. This can be seen in V(L+A) being greater than v(L+100).



Figure 18.8: Bill's consideration of gamble A and the \$100 after the shock

18.2 Insurance

Summary

- Classical economic theory explains insurance purchase through risk aversion, while prospect theory offers an alternative explanation involving riskseeking behaviour in the loss domain and overweighting of small probabilities.
- A risk-neutral agent would not purchase insurance due to its negative expected value, but a risk-averse agent with a logarithmic utility function may buy insurance due to diminishing marginal utility.
- Prospect theory's reflection effect suggests people are risk-seeking in the loss domain, which alone would predict that people wouldn't buy insurance, as it involves choosing between a certain loss (premium) and a potentially larger loss.
- Probability weighting in prospect theory suggests that people overweight small probabilities and underweight probabilities near certainty, which can make insurance attractive despite being in the loss domain.
- The combination of a convex value function in the loss domain and overweighting of small probabilities can lead to insurance purchase under

prospect theory, even when the reflection effect alone would suggest avoiding insurance.

18.2.1 Introduction

Insurance has a negative expected value due to the insurer's profit and administrative costs. Why would a consumer purchase insurance?

The classical economic explanation for the purchase of insurance is based on diminishing marginal utility, which leads to risk aversion. Consumers are willing to buy insurance as the consumer prefers the certainty of the premium payment to the risk of suffering an uninsured loss.

Prospect theory provides an alternative explanation. The purchase of insurance involves a certain loss (the premium) or a gamble involving the possibility of either a large loss or the status quo. As prospect theory has people as risk seeking in the loss domain, we would not expect them to purchase insurance.

However, under prospect theory, people also overweight small probabilities. This overweighting of small probabilities can make the purchase of insurance attractive even though it is in the loss domain.

This combination of the loss domain with a small probability is the bottom-right quadrant of the fourfold pattern to risk attitudes generated by prospect theory. People tend to be risk averse in this circumstance.

	Gains	Losses
High probability	Risk aversion	Risk seeking
Low probability	Risk seeking	Risk aversion

The following numerical example is an illustration.

An agent is considering insurance against bushfire for their house valued at H = \$1 000 000. The house has a 1 in 1000 (p = 0.001) chance of burning down. An insurer is willing to offer full coverage for a premium (R) of \$1100. (Note: \$1000 is the actuarially fair price. The additional \$100 might represent profit or administrative costs.)

18.2.2 Expected value

The first question I will consider is whether a risk-neutral person would purchase the insurance. A risk-neutral utility function is:

u(x) = x

A risk-neutral agent will choose the option with the highest expected value.

If the agent purchases insurance, they pay the premium and do not suffer any loss regardless of whether there is a bushfire or not. Therefore, the expected value of purchasing insurance is the loss of the premium:

$$\mathbb{E}[\mathbf{I}] = -R$$
$$= -\$1100$$

The expected value of purchasing insurance is the guaranteed loss of the premium, \$1100.

If the agent does not purchase insurance, they face the 1 in 1000 possibility of an uninsured loss. Therefore, the expected value of not purchasing insurance is the probability of loss times the value of the house:

$$\begin{split} \mathbb{E}[\neg \mathbf{I}] &= p \times (-H) \\ &= -0.001 \times 1\,000\,000 \\ &= -\$1000 \end{split}$$

The expected value of purchasing insurance is lower than the expected value of not purchasing insurance. Therefore, a risk-neutral agent would not purchase the insurance.

18.2.3 Expected utility

Would a risk-averse agent purchase the insurance? Suppose they have a logarithmic utility function $(u(x) = \ln(x))$ and they have \$10 000 in cash in addition to their house, giving them wealth (W) of \$1 010 000.

The logarithmic utility function has diminishing marginal utility. Diminishing marginal utility is the principle that the marginal utility from each additional unit decreases. In the context of wealth, this means that each additional dollar provides less satisfaction than the previous one.

Diminishing marginal utility means that the agent will be risk averse. They will prefer a certain outcome to a gamble with the same expected value.

To understand whether this risk-averse agent will purchase insurance in this instance, we need to compare the expected utility of purchasing insurance with the expected utility of not purchasing insurance.

The expected utility of purchasing insurance is the utility of the certain outcome, which is the utility of wealth after paying the premium:

$$\mathbb{E}[U(\mathbf{I})] = u(W - R)$$

= ln(1 010 000 - 1100)
= 13.8244

The expected utility of not purchasing insurance is the probability-weighted sum of the utility of the two potential outcomes, the utility of wealth after losing the house and the utility of wealth if the house does not burn down:

$$\begin{split} \mathbb{E}[U(\neg \mathbf{I})] &= p \times u(W - H) + (1 - p) \times u(W) \\ &= 0.001 \times \ln(1\,010\,000 - 1\,000\,000) + 0.999 \times \ln(1\,010\,000) \\ &= 13.8208 \end{split}$$

The expected utility of purchasing insurance is greater than the expected utility of not purchasing insurance. This agent will purchase insurance.

The following diagrams illustrate. First I plot the agent's utility function. The utility function is concave as the agent is risk averse. Each additional unit of wealth provides less utility than the previous unit.

I then mark each possible outcome from purchasing or not purchasing insurance on the horizontal axis and the utility of each of those outcomes on the vertical axis. These are: wealth after losing the house when uninsured (W - H), wealth after paying the insurance premium (W - R), and wealth if uninsured but the house does not burn down (W). I have not drawn this to scale.

The expected utility of not purchasing insurance is the probability-weighted sum of the utility of the two potential outcomes, the utility of wealth after losing the house and the utility of wealth if the house does not burn down. I can show this on the diagram by plotting a dash-dot line between U(W - H) and U(W). The expected utility of not purchasing insurance will be on this line.

The location of the expected utility is determined by the probability p of incurring a loss. This point lies at a distance of p from U(W) along the line (or equivalently, at a distance of 1-p from U(W-H)). This point aligns with the expected value of leaving the house uninsured $\mathbb{E}[\neg I]$.

The utility of purchasing insurance, U(W - R), is greater than the expected utility of not purchasing insurance, $\mathbb{E}[U(\neg I)]$. We can see this as U(W - R) is above $\mathbb{E}[U(\neg I)]$ on the diagram. This risk-averse agent will purchase insurance.



Figure 18.9: The utility function of a risk-averse agent

18.2.4 The reflection effect

I will now move on to the elements of prospect theory, introducing them one at a time. The first element is the reflection effect.

Consider an agent who is risk seeking in the domain of losses but weights probability linearly. Their value function is:

$$v(x) = \begin{cases} x^{0.8} & \text{where} x \ge 0\\ -2(-x)^{0.8} & \text{where} x < 0 \end{cases}$$

x is the realised outcome relative to the reference point.

The reflection effect occurs as this value function is concave in the gain domain and convex in the loss domain. This leads to risk averse and risk seeking behaviour, respectively.

In this analysis, we will take the reference point as current wealth before the purchase of the insurance. Determination of the reference point can be arbitrary. What if you pay insurance every year? Would the reference point then be wealth minus the insurance payment, meaning the insurance payment is in the gain domain? In that case, the analysis changes.



Figure 18.10: The utility of the outcomes of purchasing insurance and not purchasing insurance



Figure 18.11: The expected utility line

We can see the outcomes for an agent with a reference point of wealth before purchasing insurance on the following line (not drawn to scale). The outcomes are the premium payment -R and the value of the house after a bushfire -H. If uninsured and no fire, the agent will remain at their reference point of the status quo.



Would this agent with the reflection effect and a reference point of wealth before purchasing the insurance make the purchase?

We need to compare the weighted value of purchasing insurance with the weighted value of not purchasing insurance. The agent will purchase insurance



Figure 18.12: The expected utility of purchasing insurance and not purchasing insurance

if the weighted value of purchasing insurance is greater.

The weighted value of purchasing insurance is the value of the certain outcome, which is the loss of the premium:

$$\begin{split} V({\rm I}) &= v(-R) \\ &= -(1100)^{0.8} \\ &= -271.1 \end{split}$$

The weighted value of not purchasing insurance is the probability-weighted sum of the value of the two potential outcomes, the value of wealth after losing the house and the value of wealth if the house does not burn down:

$$\begin{split} V(\neg \mathbf{I}) &= \sum_{i=1}^n p_i v(x_i) \\ &= p \times v(-H) + (1-p) \times v(0) \\ &= -0.001 \times (1000000)^{0.8} + 0.999 \times 0 \\ &= -63.1 \end{split}$$

As the value of purchasing insurance is less than the weighted value of not purchasing insurance, that is, $V(I) < V(\neg I)$, this agent does not purchase insurance. The diminishing feeling of loss leads to them weigh the certain loss of the premium relatively more heavily than the chance of losing the value of their house.

Including loss aversion in the value function does not change the decision as all possible outcomes are in the loss domain.

The following diagrams illustrate. First I plot the agent's value function. The value function is concave in the gain domain and convex in the loss domain, leading to risk averse and risk-seeking behaviour, respectively. Each additional gain or loss of a unit of wealth relative to the reference point results in a smaller change in value than the previous unit. There is diminishing sensitivity to both gains and losses.

Each outcome and the value of that outcome from purchasing or not purchasing insurance is marked on the horizontal axis and the value of those outcomes on the vertical axis. These are: the premium payment -R and the outcome after a bushfire -H. If uninsured and no fire, the agent will remain at their reference point of the status quo.


Figure 18.13: A value function with the reflection effect

The weighted value of not purchasing insurance is the probability-weighted sum of the value of the two potential outcomes, the value of wealth after losing the house and the reference point. I can show this on the diagram by plotting a dash-dot line between v(-H) and v(0). The weighted value of not purchasing insurance will be on this line.

The location of this point is determined by the probability p of incurring a loss. This point lies at a distance of p from v(0) along the line (or equivalently, at a distance of 1 - p from v(-H)). This point aligns with the expected value of leaving the house uninsured $\mathbb{E}[\neg I]$.

The value of purchasing insurance, v(-R), is less than the weighted value of not purchasing insurance, $V(\neg I)$. The agent will not purchase insurance. They are risk seeking in this loss domain.

18.2.5 Probability weighting

Would a person who is risk seeking in the domain of losses (that is, a person with a value function with reflection effect discussed earlier) who weights probability in accordance with prospect theory purchase the insurance?

Prospect theory proposes that people overweight small probabilities. They might weight probabilities as per the following table:



Figure 18.14: The value of the outcomes of purchasing insurance and not purchasing insurance $% \left({{{\mathbf{r}}_{\mathrm{s}}}_{\mathrm{s}}} \right)$







Figure 18.16: The reflection effect and insurance

Probab ili 00/1	0.01	0.1	0.25	0.5	0.75	0.9	0.99	0.999
$\begin{array}{l} (p) \\ \textbf{Weight} 0.01 \\ \pi(p) \end{array}$	0.05	0.15	0.3	0.5	0.7	0.85	0.95	0.99

That is, an outcome with a probability of p = 0.001 is given a weight $\pi(p)$ of 0.01, ten times what its probability would suggest. An outcome with a probability of p = 0.01 is given a weight $\pi(p)$ of 0.05, five times the weight.

Similarly, a probability of p = 0.99 is given a weight $\pi(p)$ of 0.95 and a probability of p = 0.999 is given a weight $\pi(p)$ of 0.99, implying that probabilities near certainty are underweighted relative to certainty itself.

To assess whether this person will purchase insurance, we need to compare the weighted value of purchasing insurance with the weighted value of not purchasing insurance. They will purchase insurance if the weighted value of purchasing insurance is greater.

The weighted value of purchasing insurance is the value of the certain outcome, which is the loss of the premium:

$$\begin{split} V({\rm I}) &= v(-R) \\ &= -(1,100)^{0.8} \\ &= -271.1 \end{split}$$

The weighted value of not purchasing insurance is the weighted sum of the value of the two potential outcomes, with the weights from the table above:

$$\begin{split} V(\neg \mathbf{I}) &= \sum_{i=1}^n \pi(p_i) v(x_i) \\ &= \pi(p) \times v(0) + \pi(1-p) \times v(-H) \\ &= \pi(0.001) \times v(-1\,000\,000) + \pi(0.999) \times v(0) \\ &= -0.01 \times (1\,000\,000)^{0.8} + 0.99 \times 0 \\ &= -631 \end{split}$$

As the value of purchasing insurance is greater than the value of not purchasing insurance, that is, $V(I) > V(\neg I)$, this person purchases insurance.

Although the diminishing feeling of loss leads to them weigh the certain loss of the premium relatively more heavily than the chance of losing the value of their house, the overweighting of the probability of fire leads them to purchase insurance. Again, if we had included loss aversion, it would not have changed the decision, as all possible outcomes are in the loss domain.

18.3 A multi-bet

A multi-bet allows a gambler to combine a series of individual bets into a single wager, with the odds of all the single bets multiplied to achieve the final payoff. The gambler only wins the wager if all of the single bets are successful. Or in other words, if a single bet is lost, the entire multi-bet is lost. A multi-bet is also known as an "accumulator" bet or "parlay".

For example, a multi-bet might combine the following bets:

- GWS Giants to defeat Adelaide Crows: \$1.65 (that is, \$1.65 is paid out for each \$1 bet)
- Fremantle Dockers to defeat Sydney Swans: \$2.10
- Essendon Bombers to defeat Geelong Cats: \$3.50
- North Melbourne Kangaroos to defeat Melbourne Demons: \$4.00
- West Coast Eagles to defeat Brisbane Lions: \$6.00

If all five bets are successful, the gambler would win \$291 for every \$1 they have bet. (That is $1.65 \ge 2.10 \ge 3.50 \ge 4 \ge 6 = 291.06$). If any of GWS, Fremantle, Essendon, North Melbourne, or the West Coast Eagles lose, the bet is lost.

Many bookmakers also offer a "cash out" option for multi-bets. If the bet has been successful up to the date of the "cash out", a gambler can "cash out" their bet before the remaining games are complete at a price offered by the bookmaker. The "cash out" offers are typically unattractive relative to the expected value of seeing out the rest of the multi-bet.

For example, Betty places a \$20 multi-bet involving all nine games of Australian Rules Football one weekend. After eight games, she has picked all eight winners. If Geelong defeats North Melbourne in the ninth game she will win \$10,000. Geelong is a heavy favourite, with a 95% probability of winning.

Betty checks the cash out price for the multi-bet and sees that she can cash out the bet now for \$7,000, a great return on her initial \$20. That return, however, is much below the expected value of seeing the multi-bet through to the end (\$9,500).

Betty makes decisions according to prospect theory. That is, she judges gains and losses relative to a reference point, is loss averse, and has diminishing sensitivity to gains and losses in both directions. She also overweights certainty (which is equivalent to overweighting small probabilities). What elements of prospect theory might lead Betty to cash out the bet before the final game?

We can consider multiple potential reference points Betty might use to make her decision.

One potential reference point is Betty's position before making the bet. In that case, Betty is comparing:

- a certain gain of \$6980 and
- a gamble with a loss of \$20 and a gain of \$9,980.

The gamble is largely in the gain domain in which Betty is risk averse. Her risk aversion may lead her to cash out rather than take the gamble. The potential \$20 loss may be overweighted due to loss aversion, but it is a relatively insignificant sum.

Another potential reference point is Betty's position immediately after making the bet. She has adapted to the payment of \$20. Here the analysis is similar. Betty is comparing:

- a certain gain of \$7000 and
- a gamble with a gain of \$10,000 or a payment of zero

The gamble is completely in the gain domain, where Betty is risk averse. Her risk aversion may lead her to cash out rather than gamble. Loss aversion is irrelevant in this instance.

Another potential reference point is that Betty is taking the \$7000 to be locked in. This means she is comparing:

- staying at the status quo with certainty and
- a gamble involving a potential loss of \$7000 and a gain of \$3,000.

In this case, the combination of risk aversion reducing the value of the gain and loss aversion increasing the relative magnitude of the pain of loss could lead her to cash out. This would be counteracted by the convex curvature in the loss domain, but the loss aversion effect would likely dominate.

Finally, Betty's weighting of probabilities will also affect her decision. Overweighting certainty means overweighting small probabilities, such as the small probability of Geelong losing. This overweighting would push her towards cashing out as the loss would have greater weight in her calculation of the weighted value of each option.

Chapter 19

Prospect theory applications

Summary

- Prospect theory can explain various real-world economic behaviours that deviate from traditional economic models, as illustrated by studies on taxi drivers, investors, and homeowners.
- Some studies on taxi driver behaviour on rainy days suggest that drivers may have a daily earnings target (reference point), and their decisions reflect loss aversion and diminishing sensitivity to gains, key elements of prospect theory.
- The disposition effect in investing, where investors tend to sell winning stocks and hold losing ones, can be explained by prospect theory's reflection effect.
- In the housing market, owners subject to nominal losses tend to set higher asking prices, potentially due to loss aversion.
- However, follow-up studies and alternative explanations demonstrate that the application of prospect theory to these scenarios is not without controversy, and results can vary based on methodology and data interpretation.

19.1 Introduction

In this part, I discuss some applied problems that have been analysed using prospect theory.

19.2 Taxi driver behaviour on rainy days

Why can't you find a taxi on a rainy day?

One possible explanation comes from Colin Camerer et al. (1997), who studied the labour supply of New York City taxi drivers.

The taxi drivers rent a cab for a 12-hour period for a fixed fee, plus petrol. Within the 12 hours, a driver can choose how long they keep the taxi out.

A taxi driver's effective wage can vary for many reasons, such as weather, subway breakdowns, day of the week and conferences. When they are busier, they have a higher effective wage. That is, they earn more fares.

In two of the three samples they examined, Camerer et al. found that drivers drove less when their effective wages were higher. This was the case for inexperienced drivers in all three samples, and they drove significantly less than experienced drivers when wages were high.

This contrasts with the basic prediction of economic theory that supply increases with price. Supply curves slope upwards.

Camerer et al. argue that this result is because taxi drivers have a daily earnings target, beyond which they derive little additional utility. This leads them to work until they reach their target, which occurs more quickly on days with a higher wage.

They argued that the drivers engage in "narrow bracketing" when they make decisions each day, isolating them as single decisions (how much should I work today?) rather than thinking about them as a stream (how much should I work each day this week?)

Aversion to falling below the reference point is consistent with loss aversion, with a result below the reference point causing stronger feelings than a result a similar amount above the reference point.

There have been numerous follow-up studies of taxi drivers. The results of these studies have varied.

- Farber (2005) studied New York cab drivers and found that the decision to stop work was primarily a function of how many hours had been worked up to that point in the day. He identified the difference between his and Camerer et al.'s result as being due to different empirical methods and measurement problems with the Camerer et al. data.
- Farber (2008) found that a labour supply model with reference-dependent targets better fits than a standard neoclassical model. However, there was substantial variation day-to-day in any given driver's reference income level and most shifts ended before that reference income was reached.
- Farber (2015) used a much larger dataset on New York taxi driver behaviour and found that, as standard economic theory would predict, taxi

drivers drive more when they can earn more. Farber also found that drivers did not earn more when it was raining.

• Finally, Martin (2017) examined taxi driver labour supply using the S-shaped reference dependence of prospect theory. That is, Martin used a model with the reflection effect, with risk-seeking behaviour in the loss domain and risk aversion in the gain domain. Martin found evidence that taxi driver behaviour was consistent with this full form of prospect theory. He differentiated from the other papers on the basis that they considered a narrower version of reference dependence focusing on loss aversion only.



Fig. 1. Reference dependence utility function examples with and without loss aversion. Note: Here the reference point is denoted by the vertical line. The inclusion of loss aversion increase the marginal utility of income in the range of income below the reference point, referred to as the "loss region".

19.3 The disposition effect

The disposition effect is the tendency for investors to sell stocks that are in the gain domain relative to the purchase price and to hold stocks that are in the loss domain.

While tax implications or portfolio rebalancing are both potential explanations for asymmetric behaviour relating to the sale of stocks, these factors are insufficient to explain the observed behaviour.

Most behavioural explanations have turned to prospect theory.

For example, Shefrin and Statman (1985) argued that the disposition effect is driven by the reflection effect, whereby investors are risk seeking in the loss domain and risk averse in the gain domain. To demonstrate how it works, they present the following scenario:

[C]onsider an investor who purchased a stock one month ago for \$50 and who finds that the stock is now selling at \$40. The investor must now decide whether to realize the loss or hold the stock for one more period. To simplify the discussion, assume that there are no taxes or transaction costs. In addition, suppose that one of two equiprobable outcomes will emerge during the coming period: either the stock will increase in price by \$10 or decrease in price by \$10. According to prospect theory, our investor frames his choice as a choice between the following two lotteries:

A. Sell the stock now, thereby realizing what had been a \$10 "paper loss".

B. Hold the stock for one more period, given 50-50 odds between losing an additional \$10 or "breaking even."

For an investor who is risk seeking in the loss domain, holding would be attractive. The following figure illustrates that the certain loss of \$10 (from the reference point of \$50) gives lower value than holding for a chance of eliminating the loss.



Figure 19.1: The disposition effect in the loss domain

If we craft an alternative scenario where the stock is now selling at \$60, selling would realise a \$10 gain, while holding the stock would be a risky prospect with the same expected value. An investor who is risk averse in the gain domain will sell.

The following figure illustrates that the certain gain of \$10 (from the reference point of \$50) gives higher value than holding for a chance of a larger gain.



Figure 19.2: The disposition effect in the gain domain

19.4 The housing market

Genesove and Mayer (2001) examined housing data from Boston. They found that owners subject to nominal losses set higher asking prices, with the increase in asking price being 25% to 35% of the difference between the expected selling price and their original purchase price.

They also found that these owners attain higher prices, covering around 3% to 18% of that difference.

This suggests sellers are averse to realising nominal losses. However, note that the aversion leads to a better outcome.

Chapter 20

Prospect theory exercises

20.1 Changes in probability

You are in a draw for \$1 million. Consider the following four scenarios where your chances of winning increases by 5%:

- A) From 0% to 5%
- B) From 5% to 10%
- C) From 50% to 55%
- D) From 95% to 100%

Most people report that scenarios A) and D) represent better news. Why?

? Answer

There is strong experimental evidence that we overweight certain events relative to near certain events. In this instance, we will tend to overweight the shift to a certain result (scenario D) relative to shifts involving intermediate probabilities (e.g. scenario B and C). This overweighting of certainty is effectively the same as overweighting low probability events. As a result, the probability shift that provides an initial chance at \$1 million is also overweighted (scenario A).

The result is that scenarios A and D tend to be seen as better news than scenarios B and C.

The following diagram illustrates. On the horizontal axis is the probability (p). On the vertical axis is the decision weight applied to each probability (π) . If people applied probability weights linearly as they do in expected utility theory, the dashed line would represent how probabilities map to weights. Under prospect theory, people overweight small probabilities and underweight probabilities short of certainty. The solid line represents one

functional mapping of these possibility and certainty effects. From this diagram, you can then see the change in decision weight from each change in probability. The changes from 0% to 5% and from 95% to 100% represent much larger changes in decision weight than the changes in intermediate probabilities.



20.2 Bitcoin

Edna, Ferdinand and Gretel each bought some Bitcoin at \$50,000. The price rose to \$80,000 and then dropped to \$60,000, at which time they sold it.

All three are loss averse and have the following reference-dependent value function:

$$v(x) = \begin{cases} x & \text{where } x \ge 0\\ 2x & \text{where } x < 0 \end{cases}$$

Edna uses the purchase price as her reference point. Ferdinand uses the peak price as his reference point. Gretel uses the sale price as her reference point.

What is the change in value for each person? Who is happiest?

? Answer

Value for Edna:

$$v(x) = v(60 - 50)$$

= 10

Value for Ferdinand:

$$v(x) = v(60 - 80)$$
$$= -2 \times 20$$
$$= -40$$

Value for Gretel:

$$v(x) = v(60 - 60)$$
$$= 0$$

Edna is happiest whereas Ferdinand is most disappointed. Both Edna and Gretel see the peak price as a foregone gain, whereas Ferdinand sees the failure to sell at the peak as a loss.

20.3 Reference points

Megan has the following reference-dependent value function:

$$v(x) = \begin{cases} x & \text{where } x \ge 0\\ 2x & \text{where } x < 0 \end{cases}$$

where x is the realised outcome relative to the reference point.

a) If you look at v(x), we expect Megan to be:

- loss averse
- have no decreased sensitivity to changes of greater magnitude.

Why?

? Answer

Megan is loss averse as the slope of the value function in the loss domain is steeper than in the gain domain. One unit of loss leads to a greater change in value than one unit of gain.

There is no decreases sensitivity as the value function is linear in both domains. Sensitivity is constant. As the size of the loss or gain increases, the change in value remains constant despite the increasing magnitude.

b) Assume Megan has received a birthday card from her aunt, which some years contains \$25 and other years contains nothing.

She opens the card and it contains \$10.

Consider two alternative reference points: Megan is pessimistic and expects no money in the card, and Megan is optimistic and expects \$25.

Compute Megan's value under each reference point. Which reference point yields higher value?

? Answer

The value of the \$10 from the reference point of expecting nothing is:

$$v(x) = v(10)$$
$$= 10$$

The value of the \$10 from the reference point of expecting \$25 is:

$$v(x) = v(10 - 25)$$

= 2 × (-15)
= -30

Megan has higher value when she does not expect to receive any money. She gets the value of a gain of \$10. If she expected the \$25, she suffered the value of a loss of \$15.

c) Megan has received the \$10 in the card and a piece of birthday cake. Her value function over money and pieces of birthday cake is:

$$v(x) = v(m - r_m) + v(4c - 4r_c)$$

Where m is the amount of money she receives, r_m is her reference point of how

much money she expects, c is how many pieces of birthday cake she receives and r_c is how many pieces of birthday cake she expects.

To illustrate how this value function works, imagine Megan expects two pieces of cake and her dog jumps onto the table and eats one of them. Her change in the value function is:

$$\begin{split} v(x) &= v(4c-4r_c) \\ &= v(4\times 1-4\times 2) \\ &= v(-4) \end{split}$$

As v(x) = 2x when x < 0:

$$v(-4) = -8$$

For this question, assume Megan does not believe that she will receive any money and she expects to eat two pieces of birthday cake. Her reference point is therefore $r_m = 0$ and $r_c = 2$. She receives \$10 from her aunt.

Her brother - who loves cake - offers to buy one of her pieces of birthday cake. What price p_s would make Megan indifferent between selling and keeping the piece of cake?

? Answer

Megan will be indifferent when the value from each option is the same:

 $\underbrace{ v(10-0+p_s)}_{\substack{\text{Value from} \\ \text{money plus} \\ \text{for cake}}} + \underbrace{ v(4 \times 1 - 4 \times 2)}_{\substack{\text{Value lost from} \\ \text{giving up cake}}} = \underbrace{ v(10-0) + v}_{\substack{\text{Value of} \\ \text{weeping cake}}} \underbrace{ v(4 \times 2 - 4 \times 2)}_{\substack{\text{Value of} \\ \text{money}}} \underbrace{ v(10 \times 2 - 4 \times 2)}_{\substack{\text{Value of} \\ \text{money}}} \underbrace{ v(10 + p_s) + v(-4) = v(10) + v(0)}_{\substack{\text{Value of} \\ \text{money}}} \\ 10 + p_s - 8 = 10 \\ p_s = 8 \end{aligned}$

d) Assume that Megan expects to receive only one piece of birthday cake. Her brother then offers to sell her his piece of cake. For $r_m = 0$ and $r_c = 1$, what price p_b would make Megan indifferent between buying the cake and eating only her own piece?

Answer

Megan will be indifferent when the value from each option is the same:



e) Why is there a difference between the price at which Megan was willing to sell cake in part c) compared to the price she was willing to pay in part d)?

? Answer

The willingness to accept in part c) is higher than the willingness to pay in part d) as in part c) the foregone cake is coded as a loss. This loss reduces value at twice the rate of a gain in cake. The payment for the cake in both parts is in the gain domain due to the birthday present of \$10, so the payment received or paid is given less weight than any loss.

20.4 Saving on a purchase

Consider the following two scenarios:

A) You are considering buying a new type of coffee bean for your home coffee machine. It costs \$50 at your local hipster cafe, but you discover that it is for sale for \$40 at the supermarket 20 minutes drive from your home. Do you make the trip?

B) You are considering buying a new laptop. It costs \$1990 at your local computer store, but you discover that it is for sale for \$1980 at another computer store 20 minutes drive from your home. Do you make the trip?

When people are presented with scenarios such as this, they tend to report that they are less likely to make the trip in Scenario B for the more expensive product.

Explain how an S-shaped value function with diminishing value in both gains and losses could result in this behaviour.

? Answer

Diminishing value means that the absolute difference between v(-40) and v(-50) is much larger than the absolute difference between v(-1980) and v(-1990). For example, suppose the value function was:

$$v(x) = \begin{cases} x^{\frac{1}{2}} & \text{where} \quad x \ge 0\\ -2(-x)^{\frac{1}{2}} & \text{where} \quad x < 0 \end{cases}$$

This would mean that the difference between v(-40) and v(-50) is:

$$v(-40) - v(-50) = -2(40^{\frac{1}{2}} - 50^{\frac{1}{2}})$$

= 1.49

The difference between v(-1980) and v(-1990) is:

$$\begin{split} v(-1980) - v(-1990) &= -2(1980^{\frac{1}{2}} - 1990^{\frac{1}{2}}) \\ &= 0.22 \end{split}$$

Much less value is gained by driving for the 20 minutes across town for the computer.



20.5 A 60:40 gamble

Suppose an agent has the following reference-dependent value function:

$$v(x) = \begin{cases} x^{3/4} & \text{where} \quad x \ge 0\\ -2(-x)^{3/4} & \text{where} \quad x < 0 \end{cases}$$

Where x is the realised outcome relative to the reference point.

Assume that the agent's reference point is the status quo and the agent is offered the gamble A:

$$(\$100, 0.6; -\$100, 0.4)$$

a) Will they want to play this gamble? Why?

? Answer

The utility from the gamble is:

$$V(A) = 0.6v(100) - 0.4v(-100)$$

= 0.6 × (100)^{0.75} - 0.4 × 2 × (100)^{0.75}
= -6.32

They will not want to play this gamble as it has a negative value for the agent. They could receive value of 0 by simply not playing.

The reason for this negative value is that the agent is loss averse. The loss of \$100 is given twice the weight of an equivalent gain, meaning that they reject the bet despite a win of \$100 being more probable.

b) Suppose the agent takes this gamble and loses \$100. They feel bad about it and perceive it as a loss. Their reference point is unchanged at the original status quo. They are offered gamble A again? Do they accept this second time? Why?

Answer

After losing \$100 but not changing their reference point, they have two possible outcomes relative to their reference point: recovery of their loss of \$100 so they come out even and a loss of \$200 (losing \$100 twice).

$$V(A) = 0.6v(-100 + 10) - 0.4v(-100 - 100)$$

= 0.6 × (0)^{0.75} - 0.4 × 2 × (200)^{0.75}
= -42.55

The utility of not playing the gamble involves remaining with a loss of \$100:

$$V(\neg A) = v(-100) = -2 \times (100)^{0.75} = -63.25$$

They will now want to play the gamble as it has a greater value than staying with their current loss. The reason the gamble becomes attractive is because it gives an opportunity to recover the loss. The agent is risk seeking over the loss domain.

c) The agent is offered a new job for which they receive a \$50,000 sign-on bonus. They adapt to their new wealth, so their reference point changes to accommodate their new situation. They are now offered gamble A again. Do they accept? Why?

? Answer

With their new reference point, this question is effectively the same as part a). They will refuse the bet.

20.6 A 50:50 gamble

Suppose Tim has the following reference-dependent value function:

$$v(x) = \begin{cases} x^{1/2} & \text{where } x \ge 0\\ -2(-x)^{1/2} & \text{where } x < 0 \end{cases}$$

x is the change in Tim's position relative to his reference point.

a) What feature of Tim's value function leads to loss aversion? Explain.

? Answer

The part of the value function that applies when x < 0 is increased by a factor of two relative to that part of the value function for when x > 0. This is done by multiplying the bottom portion by 2.

b) Tim considers the gamble A: (\$250, 0.5; -\$100, 0.5). Will Tim want to play this gamble?

Answer

$$\begin{split} V(A) &= p_1 v(x_1) + p_2 v(x_2) \\ &= 0.5 v(\$250) + 0.5 v(-\$100) \\ &= 0.5 \times 250^{1/2} - 0.5 \times 2 \times 100^{1/2} \\ &= -2.09 \end{split}$$

Tim has negative value from the gamble, when he could simply have zero value by declining. He rejects.

c) Explain what features of the value function lead him to accept or reject the gamble.

? Answer

Two features of the value function lead him to reject:

- Tim is loss averse (the -2 in the bottom equation), which leads him to give twice the weight to losses relative to an equal sized gain.
- Tim has diminishing sensitivity to losses and gains. The larger gain is reduced proportionately more by this diminishing sensitivity.

d) Suppose Tim were to experience a large positive or negative shock to his wealth that does not immediately change his reference point. Could either shock cause him to change his decision concerning gamble A? Explain.

? Answer

Both a positive and negative shock could lead Tim to change his decision. A large positive shock would place the entire bet into the gain domain. That would remove loss aversion as a factor in rejecting the bet. Given the high expected value, that would likely be sufficient for him to accept. However, a large shock would also place Tim further up the value function to a region where it is more linear (i.e. he is less risk averse). This will also increase the tendency to accept the bet.

A large negative shock would place the entire bet into the loss domain. That would remove loss aversion as a factor in rejecting the bet as there is no gain against which the loss can be given relatively greater weight. Due to diminishing sensitivity to losses, Tim is also risk seeking in the loss domain. This would lead him to accept any bet with a positive expected value.

20.7 Insurance AND gambling

Explain how probability weighting as proposed under Prospect Theory can lead a person to simultaneously gamble and purchase insurance.

? Answer

Decisions under Prospect Theory are the outcome of two factors:

- The value function that gives a value to each outcome relative to a reference point
- The probability weighting function that gives a decision weight to the probability of each outcome.

The Prospect Theory value function leads people to be risk averse in the domain of gains and risk seeking in the loss domain. This combination would tend to lead people to purchase neither insurance or lotteries. Their risk seeking behaviour could lead them to avoid the certain loss of the insurance premium. They would rather risk the loss of the insurable asset.



Figure 20.1: Purchasing insurance

Their risk averse behaviour in the domain of gains would make a lottery with a negative expected value unattractive, as the following diagram shows (not drawn to scale - the relative size of a win would typically dwarf the ticket price).



Figure 20.2: Purchasing a lottery

However, the probability weighting function can counteract that effect. By overweighting the small probability of an insurable event or a lottery win, the agent may decide to insure or purchase a lottery despite the risk attitudes inherent in the value function.

As an example, consider an agent who overweights a probability of 1 in a million lottery win by 1,000 times, acting as though they would win the \$1 million prize once every 1000 lotteries. They also overweight the probability of a one in a thousand fire that would destroy their million dollar house by 100 times, acting as though it is a 1-in-10 chance. (These numbers are extreme, but I am creating a toy example.)

Suppose the lottery ticket is \$1 and their utility function is:

$$v(x) = \begin{cases} x^{1/2} & \text{where} \quad x \ge 0\\ -(-x)^{1/2} & \text{where} \quad x < 0 \end{cases}$$

Value of purchasing the lottery:

$$\begin{split} v(x) &= \sum_{i=1}^n \pi(x_i) v(x_i) \\ &= 0.001 \times (1000000 - 1)^{1/2} - 0.999 \times (1)^{1/2} \\ &= 0.000999 \end{split}$$

The value of not purchasing the lottery is zero.

They will purchase the lottery as it has the higher value. Value of not purchasing insurance, which involves the potential loss of the house:

$$\begin{split} v(x) &= \sum_{i=1}^n \pi(x_i) v(x_i) \\ &= -0.1 \times (1000000)^{1/2} + 0.9 \times (0)^{1/2} \\ &= -100 \end{split}$$

Value of purchasing insurance, which is the certain loss of the premium:

$$\begin{split} v(x) &= \sum_{i=1}^n \pi(x_i) v(x_i) \\ &= -(1000)^{1/2} \\ &= -31.6 \end{split}$$

They purchase insurance as it has higher value than not purchasing it.

20.8 Locking in a win

Amber sued a tabloid newspaper for defamation. The trial has completed and the judge has retired to make her decision.

Amber's lawyer tells her that she has a 95% chance of winning and receiving \$1,000,000 in damages, meaning she has a 5% chance of leaving with nothing. (Ignore any potential costs.) She then receives an offer of settlement from the newspaper for \$800,000. Amber accepts.

Use the fourfold pattern of attitudes to risk under prospect theory to explore why Amber accepted the settlement offer.

Answer

Rejecting the offer has a higher expected value than the value of the settlement:

 $E[\text{reject}] = p_{\text{win}} x_{\text{win}}$ $= 0.95 \times 1000000$ = 950000

Two forces under prospect theory would lead Amber to take the option with the lower expected value:

- People are risk averse in the gain domain. Amber would prefer certainty to a gamble with the same expected value, and depending on the level of risk aversion, would be willing to accept an amount for certain lower than the expected value of the bet. That is, the certainty equivalent of the gamble would be less than its expected value.
- People overweight small probabilities. The 5% probability of losing is likely overweighted by Amber.

20.9 Lotteries as incentives

A common intervention to encourage behaviour change is to offer a lottery ticket as an incentive. For example, a lottery ticket might be offered to people who complete a survey.

The government is considering whether it should incentivise vaccination with either:

- i) a payment of \$10 or
- ii) a lottery ticket with a 1 in a million chance of winning \$10 million.

Use concepts from prospect theory to explain why either option might be more successful.

💡 Answer

There are two opposing forces that could lead to either option being preferred.

Both the lottery and payment of \$10 are in the gain domain. Under prospect theory, people tend to be risk averse in the gain domain. This will make the payment, which has the same expected value as the lottery, more attractive.

Conversely, under prospect theory, people do not weight outcomes directly by their probability. They apply decision weights that tend to overweight small probabilities and underweight probabilities just short of certainty. This means that the small chance of \$10 million will likely be overweighted. This could lead to the value of the lottery exceeding that of a certain payment. Which effect dominates depends on how risk-averse people are and how

much they overweight the probability.

Part IV

Intertemporal choice

Summary

- Intertemporal choice involves decisions with costs and benefits occurring at different times, which can be significant due to delayed feedback, irreversibility, and high stakes.
- A core principle is that people tend to discount future costs and benefits, preferring earlier benefits and later costs.
- In this subject, discrete time is used, where time occurs in a series of steps, as opposed to continuous time.
- Streams of payoffs in discrete time can be represented using notation like $S = (t_1, x_1; t_2, x_2; ...; t_n, x_n)$, where t_i is the time period and x_i is the payoff in that period.

Intertemporal choice refers to decisions involving costs and benefits occurring at different times.

Almost every decision is an intertemporal choice. Intertemporal choices can be important because:

- First, feedback may not be immediate. For example, when will you realise that your retirement savings are inadequate?)
- Second, your choices may be irreversible. What can you do if you reach retirement age with little savings?
- And finally, stakes can be large, affecting your health, wealth, family or career.

One of the core principles of intertemporal choice is that people tend to discount future costs and benefits. They prefer to receive benefits earlier, rather than later, and prefer to incur costs later rather than earlier.

Discrete versus continuous time

In this subject, we will consider what is called "discrete time". In discrete time, time occurs in a series of steps. For example, we might consider the following sequence of time periods:

$$t_0, t_1, t_2, t_3, t_4, t_5, \dots$$

At each discrete moment in time, the agent might make a decision or receive a payoff. We assume that there is no moment between those two steps. For example, if t = 0 is today and t = 1 is in one week, we consider only those two moments, not any time between.

Discrete time contrasts with "continuous time", where time is a continuous variable. Time is divisible into an infinite number of steps. For example, if t = 0 is today and t = 1 is in one week, there is an infinite number of other points in time between.

Notation

One way of representing a stream of payoffs in discrete time is the following form:

$$S = (t_1, x_1; t_2, x_2; ...; t_n, x_n)$$

 t_i is the period in which the payoff is received. \boldsymbol{x}_i is the payoff received in period $t_i.$

Consider these two simple streams of payoffs.

For both streams, period $t_1=0,$ which is now. Period $t_2=1,$ which is one year from today.

The first stream is \$100 now and nothing in a year.

$$S_1 = (0, \$100; 1, 0)$$

The second stream is nothing today and \$107 in one year.

$$S_2 = (0, 0; 1, \$107)$$

Would you prefer S_1 or S_2 ?

Chapter 21

Exponential discounting

Summary

- Exponential discounting is a model where agents discount future costs and benefits at a consistent rate through time, with each additional period of delay resulting in a discount by a factor of δ .
- Under this model, agents seek to maximise the discounted utility of future consumption, resulting in a smooth decline in the present value of future payoffs over time.
- One implication of exponential discounting is time consistency, meaning an agent's relative preference between two future options remains unchanged as time passes and the options draw closer.
- The model assumes consumption independence, where utility in one period is independent of consumption in other periods, allowing goods to be split and moved across periods without changing their value beyond the discount effect.
- Additional assumptions include stationary preferences (the utility function remains constant across periods) and utility independence (decisionmakers only aim to maximise the sum of discounted utilities without preference for their distribution over time).

21.1 Introduction

Exponential discounting occurs when an agent discounts future costs and benefits at a consistent rate through time.

Under exponential discounting, each additional period of delay results in a discount of a future cost or benefit by a factor of δ . The discount factor δ is a

number between 0 and 1. The higher the discount factor, the less the agent discounts future costs and benefits.

You will often see discussion of the "discount rate", r. In discrete time, the relationship between δ and r is as follows:

$$\delta = \frac{1}{1+r}$$

A larger discount factor implies less discounting. A larger discount rate implies more discounting.

Under the standard model of exponential discounting, an agent with a choice between alternative streams of payoffs will seek to maximize the discounted utility of the future path of consumption.

The following equation is an example of exponential discounting, with a stream of costs or benefits x_0 through to x_T incurred at periods 0 through to T. U_0 is utility of the stream of payoffs at time t = 0. x_t is the payoff in period t.

$$\begin{split} U_0 &= u(x_0) + \delta u(x_1) + \delta^2 u(x_2) + \delta^3 u(x_3) + \ldots + \delta^T u(x_T) \\ &= \sum_{t=0}^{t=T} \delta^t u(x_t) \\ 0 &\leq \delta \leq 1 \end{split}$$

Each period of delay results in a discount of the future cost or benefit by a factor of δ . One period of delay results in a discount of δ . Two periods of delay results in a discount of δ^2 . Three periods of delay results in a discount of δ^3 and so on.

The degree of discounting in this equation evolves each period as 1, δ , δ^2 , δ^3 , δ^4 and so on. This results in a smooth decline in present value of a future payoff over time.

21.2 Visualising exponential discounting

Figure 21.1 illustrates the effect of exponential discounting. The figure plots the size of the discount as a function of t for an exponential discounter with $\delta = 0.9$ and $\delta = 0.75$.

21.3 Time-consistency

One implication of exponential discounting is time consistency.



Figure 21.1: Exponential discount curves

Once the agent starts moving along the consumption path, they are timeconsistent with their initial plan. For example, consider an agent facing the following two choices:

Would you like \$100 today or \$110 next week?

Would you like \$100 next week or \$110 in two weeks?

An exponential discounter will choose \$100 in both choices or \$110 in both choices. The reason is that after one week the second choice effectively becomes the same as the first choice. Time consistency implies that they will continue to want to make the same choice regardless of when they are making it.

21.4 Exponential discounted utility model assumptions

The standard model of exponential discounting is underpinned by several assumptions.

21.4.1 Consumption independence

The first assumption is consumption independence.

Consumption independence means that utility in period t + k is independent of consumption in any other period. An outcome's utility is unaffected by outcomes in prior or future periods.

Imagine a world where an exponential discounter intends to consume a behavioural economics subject.

Suppose that exponential discounter wants to consume lecture 1 at t + 3 and lecture 2 at t+4. Under consumption independence, if the agent does not attend lecture 1, they still expect to benefit from lecture 2 at t+4 consistent with the plan they decided at t.

This assumption allows us to write $x = x_1 + x_2 + x_3 + \ldots + x_n$. That is, good x can be split, allocated and moved across periods without changing the value of that good beyond the effect of the discount.

21.4.2 Stationary preferences

A second assumption is stationary preferences.

That is, $U_t = U_{t=K}$.

The utility function is stationary across periods. The functional form of U_t is the same as the functional form of U_{t+k} .

That means that if someone likes ice cream today, they will get the same utility from ice cream at a future time of consumption. Any preference for ice cream today versus tomorrow comes from the discounting of future consumption, not from changes in taste.

Similarly, with stationary preferences, you would not learn to appreciate the taste of wine over time.

21.4.3 Utility independence

A third assumption is utility independence.

Under utility independence, all that matters is maximizing the sum of discounted utilities. Decision makers have no preference for the distribution of utilities. They don't seek to delay gratification or get unpleasant things out of the way.

Chapter 22

Exponential discounting examples

Summary

- Exponential discounters make decisions by comparing the discounted utility of different options, choosing the one with the highest value.
- The discount factor δ determines how much future payoffs are discounted, with a higher δ resulting in less discounting of future outcomes.
- Visualisations, such as graphs showing discounted utility over time, can help illustrate how exponential discounters evaluate and compare different options.
- The following examples illustrate that:
 - Time consistency is demonstrated through consistent choices across different time frames.
 - The amount needed to compensate for a delay increases exponentially with time.
 - Different discount factors can lead to different choices between the same options.

22.1 Introduction

In this part, I will work through several numerical examples of decisions by an exponential discounter.

22.2 Example 1: \$100 today or \$110 next week

Alison is an exponential discounter with discount factor $\delta = 0.95$ and utility each period of $u(x_n) = x_n$. She is offered two choices.

Choice 1: Would she prefer \$100 today (t = 0) or \$110 next week (t = 1)?

To determine this, we calculate the discounted utility of each option. Alison will prefer the option with the highest discounted utility.

We write the discounted utility of the \$100 today as $U_0(0, $100)$. The subscript 0 indicates that the utility is calculated at time t = 0. The first number in the brackets indicates the time the payment is received. In this case, the \$100 is received at t = 0, today. The second number in the brackets is the value of the payment.

The discounted utility of the \$100 today, $U_0(0, $100)$, is:

$$U_0(0,\$100) = u(\$100)$$

= 100

There is no discount δ applied to the \$100 as it is received immediately.

We write the discounted utility of the \$110 next week as $U_0(1, $110)$. The subscript 0 again indicates that the utility is calculated at time t = 0. The first number in the brackets, 1, indicates the time the payment is received, t = 1. The second number in the brackets is the value of the payment.

The discounted utility of the \$110 next week is:

$$\begin{split} U_0(1,\$110) &= \delta u(\$110) \\ &= 0.95 \times 110 \\ &= 104.5 \end{split}$$

In this case, a discount of $\delta = 0.95$ is applied to the \$110 as it is received in one week.

We can now compare the discounted utility of each option. $U_0(0,\$100) = 100 < 104.5 = U_0(1,\$110)$. Alison would prefer to receive \$110 next week as it leads to higher discounted utility.

Figure 22.1 visualises the effect of discounting in Choice 1.

The two bars represent the options: \$100 at t = 0 and \$110 at t = 1. The line from the \$110 option represents the discounted utility of that option at each time. At t = 0 the discounted utility of the \$110 received at t = 1 is 104.5. At t = 0 we can see that the \$110 is preferred as the line indicating discounted utility is higher than the utility of the \$100 received immediately.



Figure 22.1: Alison's consideration of Choice 1

Choice 2: Would Alison prefer \$100 next week (t = 1) or \$110 in two weeks (t = 2)?

Again, we calculate the discounted utility of each option. Alison will prefer the option with the highest discounted utility.

The discounted utility of the \$100 next week is:
$$\begin{split} U_0(1,\$100) &= \delta u(\$100) \\ &= 0.95 \times 100 \\ &= 95 \end{split}$$

The discounted utility of the \$110 in two weeks is:

$$\begin{split} U_0(2,\$110) &= \delta^2 u(\$110) \\ &= 0.95^2 \times 110 \\ &= 99.275 \end{split}$$

We can now compare the discounted utility of each option. $U_0(1,\$100) = 95 < 99.275 = U_0(2,\$110)$. Alison would prefer to receive \$110 in two weeks.

The set of decisions across Choice 1 and Choice 2 are time consistent. If Alison selected \$110 in two weeks for Choice 2 and was given a chance to change her choice after one week (which is effectively Choice 1), she would not change her decision.

Figure 22.2 visualises the effect of discounting in Choice 2.

The two bars represent the options: \$100 at t = 1 and \$110 at t = 2. The line from each represents the discounted utility of that option at each time.

For example, at t = 1 the discounted utility of the \$100 received at t = 1 is 100 and the discounted utility of the \$110 received at t = 2 is 104.50. We can read those values from the line. For any time t we can determine which option would be preferred by seeing which line is higher.

You will note that the two lines do not cross. For an exponential discounter, if one line is higher at any particular time t, it is higher at all times.

Figure 22.4 visualises Choice 2 reconsidered at t = 1, which as noted earlier, is effectively Choice 1. The discounted utility of the \$100 received immediately is less than the discounted utility of \$110 in one week. The preference for the higher-value option remains.

22.3 Example 2: How much for a one-year delay?

Brenda is an exponential discounter with discount factor $\delta=0.95$ per week and utility each period of $u(x_n)=x_n$



Figure 22.2: Alison's consideration of Choice 2



Figure 22.3: Alison's consideration of Choice 2 - discounted utilities



Figure 22.4: Alison's reconsideration of Choice 2

She is offered \$100 today. What sum would she need to be offered in one year (52 weeks) to prefer that later payment to the \$100 today?

To determine this, we calculate what sum received in one year would give Brenda the same discounted utility as receiving \$100 today.

The discounted utility of the \$100 today is:

$$U_0(0,\$100) = u(\$100)$$

= 100

The discounted utility of the sum y received in 52 weeks is:

$$\begin{split} U_0(52,\$y) &= \delta^{52} u(\$y) \\ &= 0.95^{52} \times y \end{split}$$

As there is a 52-period delay, the discount factor δ is applied 52 times.

She will prefer y in 52 weeks if U(52, y) is greater than 100.

$$\begin{array}{l} U_0(52,\$ y) > 100 \\ 0.95^{52} \times y > 100 \\ \\ y > \frac{100}{0.95^{52}} \\ \\ y > \$1440.03 \end{array}$$

Brenda would be willing to wait a year for payment if she was paid more than \$1440.03.

Figure 22.5 visualises this problem. The bar at t = 52 represents the \$1440.03 Brenda would need to be paid (at minimum) to prefer that payment to \$100 today. The line extended from that bar back to t = 0 indicates the discounted utility of that payment at any time t. At t = 0 the discounted utility of the \$1440.03 is equal to the utility of \$100.

22.4 Example 3: \$10 in five days or \$20 in 10 days?

Chelsea is an exponential discounter with discount factor $\delta = 0.75$ and utility each period of $u(x_n) = x_n$.



Figure 22.5: Brenda's choice

Would Chelsea prefer \$10 in five days (t = 5) or \$20 in 10 days (t = 10)?

We calculate the discounted utility of each option. Chelsea will prefer the option with the highest discounted utility.

The discounted utility of the \$10 in five days is:

$$\begin{split} U_0(5,\$10) &= \delta^5 u(\$10) \\ &= 0.75^5 \times 10 \\ &= 2.37 \end{split}$$

The discount factor δ is applied five times as the payment is received in five days.

The discounted utility of the \$20 in 10 days is:

$$\begin{split} U_0(10,\$20) &= \delta^{10} u(\$20) \\ &= 0.75^{10} \times 20 \\ &= 1.13 \end{split}$$

The discount factor δ is applied 10 times as the payment is received in 10 days.

We can now compare the discounted utility of each option. $U_0(5,\$10) = 2.37 > 1.13 = U_0(10,\$20)$. The discounted utility is higher for the \$10 in five days. As a result, Chelsea will prefer to receive \$10 in five days.

Dorothy is an exponential discounter with discount factor $\delta = 0.95$ and utility each period of $u(x_n) = x_n$. Dorothy has a larger discount factor than Chelsea, meaning that she applies a smaller discount to future outcomes.

Would Dorothy prefer \$10 in five days (t = 5) or \$20 in 10 days (t = 10)?

The discounted utility of the \$10 in five days is:

$$\begin{split} U_0(5,\$10) &= \delta^5 u(\$10) \\ &= 0.95^5 \times 10 \\ &= 7.74 \end{split}$$

The discounted utility of the \$20 in 10 days is:

$$\begin{split} U_0(10,\$20) &= \delta^{10} u(\$20) \\ &= 0.95^{10} \times 20 \\ &= 11.97 \end{split}$$

The discounted utility is higher for the \$20 in 10 days. Dorothy will prefer to receive \$20 in 10 days.

Figure 22.6 visualises the choices and Chelsea and Dorothy's discounting of the payoffs.

In both charts, vertical bars represent the \$10 in five days and \$20 in 10 days. The lines projecting back to t = 0 represent the discounted utility of those payoffs at each time.

When $\delta = 0.75$, the heavy discount to the more distant payoff means that it has a lower discounted utility than the smaller, sooner payment of \$10. When $\delta = 0.95$, the discount is less severe and the \$20 in 10 days has a higher discounted utility than the \$10 in five days.



Figure 22.6: Exponential discounting

Chapter 23

Exponential discounting anomalies

Summary

- Empirical investigations have revealed several anomalies inconsistent with the exponential discounting model.
- Estimates of the discount factor δ show high variability across studies, with evidence suggesting it increases with the time horizon, contrary to the constant rate assumed in exponential discounting.
- Preference reversals have been observed in experiments, where subjects change their choices as the time to receive rewards draws closer, violating the time consistency assumption of exponential discounting.
- Real-world scenarios, such as choices between healthy and unhealthy snacks, provide further evidence of time inconsistency in decision-making, challenging the exponential discounting model's assumptions.

23.1 Introduction

Empirical investigation of people's choices over time has revealed several anomalies that are inconsistent with the exponential discounting model.

23.2 Estimation of δ

The first anomaly is that estimates of the discount factor δ are highly variable.

Frederick et al. (2002) plotted estimates of δ from a set of published papers.

Figure 1a shows that the estimated discount factor increases with the time horizon. If studies with horizons of one year or less are excluded, there is no relationship between the discount factor and time horizon (Figure 1b).



Figure 1a. Discount Factor as a Function of Time Figure Horizon (all studies) Figure

Figure 1b. Discount Factor as a Function of Time Horizon (studies with avg. horizons > 1 year)

This suggests that the discount factor may vary with time.

Similar evidence comes from R. Thaler (1981). He estimated the discount rate of experimental subjects by presenting them with a choice between a prize today or a larger prize later. In each case, the subjects were asked, given the size of the prize that could be received today, how large would the future prize would need to be such that waiting would be as attractive as receiving the money now.

For example, some subjects were asked how large a future prize would need to be such that they would be happy to wait three months, one year or three years rather than receiving \$15 today. The median answers were \$30, \$60 and \$100 respectively. If you use these answers to calculate an implied discount rate, the result is 277%, 139% and 63%. It can be seen that this discount rate is decreasing with the length of delay (which matches a pattern of an increasing discount factor with the length of delay).

23.3 Preference reversal

A second anomaly relates to the consistency of choices over time. An exponential discounter is time consistent in that if they prefer one of two future payoffs, they will continue to prefer that same payoff as they get closer to the time of receipt. They do not experience what is called a preference reversal.

Green et al. (1994) offered experimental subjects choices between a smaller reward and a delayed larger reward while varying the delay. For example, they offered the choice between:

- first, \$20 now or \$50 in three months
- second, \$20 in one week or \$50 in three months and one week.

Across the choices offered to the experimental subjects, there was a consistent effect whereby incrementing the delay for both rewards equally would result in a switch from the small sooner reward to the latter larger reward. Adding one week to both rewards in the first choice can result in people changing their preference between the sooner and later reward.

A similar result was found in a study by Kirby and Herrnstein (1995). In that case, 34 of 36 experimental subjects reversed preference from a larger later reward to a smaller earlier reward as the delays to both decreased.

Here is one intuitive set of choices from R. Thaler (1981).

Choice A: Choose between one apple today and two apples tomorrow.

Choice B: Choose between one apple in one year and two apples in one year plus one day.

Some people might select one apple today in the first choice, but no-one would select one apple in one year in the second choice.

23.3.1 Healthy choices

We can also see evidence for preference reversal when choosing between a healthy and unhealthy option.

Read and Leeuwen (1998) asked 200 study participants to choose between healthy or unhealthy snacks that they would receive in one week. For example, for their snack next week, they might choose between a banana, apple, Mars bar or Snickers bar. 48% of men and 51% of women chose the healthy choice.

At the scheduled time one week later the experimenters asked the study participants to choose again. Although no reference was made to their previous choice, this effectively allowed them to change their mind. This time, only 25% of men and 11% of women chose the healthy snack. While many changed from the healthy to the unhealthy snack, almost no one changed from the unhealthy to the healthy snack.

This result is evidence of time inconsistency. The participants made different decisions depending upon when they made the decision.

Chapter 24

Present bias

Summary

- Present bias is a concept developed to account for anomalies in the exponential discounting model, where people place additional weight on costs and benefits in the present.
- The $\beta\delta$ model, also known as quasi-hyperbolic discounting, uses two discount factors: β for short-term discounting and δ for long-term discounting.
- In the $\beta\delta$ model, all future payoffs are discounted by $\beta\delta$ once, and then by δ for each additional period of delay, resulting in a larger discount for the first period than subsequent periods.
- Present bias can lead to time inconsistency, where decisions made at one point may be reversed later if given the opportunity to change.
- The $\beta\delta$ model maintains assumptions of consumption independence, stationary preferences, and utility independence from the exponential discounting model.

24.1 Introduction

One concept developed to account for anomalies in the exponential discounting model is present bias.

Present bias occurs when we place additional weight on costs and benefits at the present time.

24.2 The $\beta\delta$ (beta-delta) model

One simple model of present bias is the quasi-hyperbolic discounting model, otherwise known as the $\beta\delta$ model.¹

Under the quasi-hyperbolic discounting model, two discount factors are applied to future costs and benefits.

The first is β (beta), the short-term discount factor. All future payoffs are discounted by a single application of β , a number between 0 and 1. The discount β is applied simply because the payoff is not immediate. The higher the short-term discount factor, the less the agent discounts payoffs that are not received now.

The second is the discount factor δ that is also present in the exponential discounting model. Each additional period of delay results in a discount of a future cost or benefit by a factor of δ . The discount factor δ is also a number between 0 and 1. The higher the discount factor, the less the agent discounts future costs and benefits.

As for exponential discounting, an agent with a choice between alternative streams of payoffs under the $\beta\delta$ model will seek to maximize the discounted utility of the future path of consumption.

The following equation provides a mathematical representation of the $\beta\delta$ model, with a stream of costs or benefits x_0 through to x_T incurred at periods 0 through to T. U_0 is the discounted utility of the stream of payoffs at time t = 0. x_t is the payoff in period t.

$$\begin{split} U_0 &= u(x_0) + \beta \delta u(x_1) + \beta \delta^2 u(x_2) + \ldots + \beta \delta^T u(x_T) \\ &= u(x_0) + \beta \sum_{t=1}^{t=T} \delta^t u(x_t) \\ &0 \leq \delta \leq 1 \\ &0 \leq \beta \leq 1 \end{split}$$

The first period of delay results in a discount of the cost or benefit by a factor of $\beta\delta$. Each further period of delay results in a discount of δ .

As a result, the degree of discounting evolves over time as 1, $\beta\delta$, $\beta\delta^2$, $\beta\delta^3$, $\beta\delta^4$ and so on. This progression results in a larger discount for the first period of delay ($\beta\delta$) than the degree of discount for each subsequent period of delay (δ). There is a relative weighting toward the present.

¹This model is a discrete-time version of hyperbolic discounting.

Present bias of this nature can result in time inconsistency, with decisions at one point reversed at another if the agent is given the opportunity to change their mind.

24.3 Visualising present bias

The following figures illustrate the effect of present bias.

Figure 24.1 plots the size of the discount as a function of t for a present-biased agent with $\beta = 0.75$ and $\delta = 0.9$. The discount curve for an exponential discounter with $\delta = 0.9$ is also plotted. The curve for the present-biased agent has a large drop for the first period of delay. From then on, the discount is proportionally the same as for the exponential discounter.

We can read off the total discount factor at any time t from this chart. For example, the total discount factor for the exponential discounter is 0.9 at t = 1, 0.81 at t = 2 and 0.43 at t = 8. The total discount factor for the present-biased agent is 0.675 at t = 1, 0.61 at t = 2 and 0.32 at t = 8.



Figure 24.1: Exponential versus present bias discount curves

Figure 24.2 shows the discount curve for present-biased and exponential discounting agents with different parameters. The present-biased agent and exponential discounter have the same discount factor $\delta = 0.75$. The present-biased

agent also has the short-term discount factor $\beta = 0.75$. Again, the presentbiased agent discounts the first period of delay more than the exponential discounter.



Figure 24.2: Exponential versus present bias discount curves - higher discount rate

Figure 24.3 shows a scenario where the present-biased agent has a higher discount factor $\delta = 0.9$ than the exponential discounter with $\delta = 0.75$. The present-biased agent also has the short-term discount factor $\beta = 0.75$. The present-biased agent discounts the first period of delay more than the exponential discounter. However, due to their higher δ , the present-biased agent discounter additional periods of delay less than the exponential discounter and ultimately has a lower total discount for periods further in the future.

24.4 Assumptions

The exponential discounting model is underpinned by many assumptions. These include:

- Time consistency, whereby once the agent starts moving along the consumption path, they are time-consistent with their initial plan.
- Consumption independence, whereby utility in period t + k is independent of consumption in any other period. An outcome's utility is unaffected by



Figure 24.3: Exponential versus present bias discount curves - different delta

outcomes in prior or future periods.

- Stationary preferences, whereby $U_t = U_{t+k}.$ The utility function is stationary across periods.
- Utility independence, whereby all that matters is maximising the sum of discounted utilities. Decision makers are assumed to have no preference for the distribution of utilities.

Under the $\beta\delta$ model, we are loosening the assumption of time consistency. An agent may change their initial plan over time.

However, we maintain the assumptions of consumption independence, stationary preferences and utility independence.

Chapter 25

Present bias examples

25.1 Summary

- The following examples illustrate the contrast between exponential discounters' time-consistent choices and present-biased agents' time-inconsistent decisions when faced with immediate versus future rewards.
- The use of graphical representations to visualise how present bias affects the relative attractiveness of rewards over time, particularly when the short-term discount factor no longer applies, can be used to illustrate choice and preference reversal.
- The following examples illustrate that:
 - The $\beta\delta$ model creates a "kink" in the discounting curve, with a steeper discount for the immediate future due to the short-term discount factor β .
 - Present-biased agents may choose smaller-sooner rewards when available immediately, but larger-later rewards when both options are in the future.
 - Present bias has practical implications in financial decision-making, demonstrated through scenarios such as loan repayment choices.

25.2 Introduction

In the section, I provide some simple examples of the $\beta\delta$ model.

25.3 Exponential discounting versus present bias

For the first example, we will consider the following pair of choices presented to an exponential discounting agent and a present-biased agent and contrast their decisions.

Choice 1: Would you like \$100 today or \$110 next week?

Choice 2: Would you like \$100 next week or \$110 in two weeks?

25.3.1 The exponential discounter

The exponential discounter has $\delta = 0.95$ and utility each period of $u(x_n) = x_n$.

Would the exponential discounter prefer \$100 today (t = 0) or \$110 next week (t = 1)?

When we worked through this problem in Section 22.2, we calculated that:

$$U_0(0,\$100)=100<104.5=U_0(1,\$110)$$

The exponential discounter will prefer to receive \$110 next week as it leads to higher discounted utility.

Choice 2: Would the exponential discounter prefer \$100 next week (t = 1) or \$110 in two weeks (t = 2)?

When we worked through this problem in Section 22.2, we calculated that:

$$U_0(1,\$100)=95<99.275=U_0(2,\$110)$$

The exponential discounter will prefer to receive \$110 in two weeks.

The set of decisions across Choice 1 and Choice 2 are time consistent. If the exponential-discounting agent selected \$110 in two weeks for Choice 2 and was given a chance to change their choice after one week (which is effectively an offer of Choice 1), they would not change their decision.

25.3.2 The present-biased agent

The present biased agent has $\delta = 0.95$, $\beta = 0.95$ and utility each period of $u(x_n) = x_n$.

Choice 1: Would this agent prefer \$100 today (t = 0) or \$110 next week (t = 1)? The discounted utility of the \$100 today is:

$$U_0(0,\$100) = u(\$100)$$

= 100

The discounted utility of the \$110 next week is:

$$\begin{split} U_0(1,\$110) &= u(x_0) + \beta u(x_1) \\ &= \beta \delta u(\$110) \\ &= 0.95 \times 0.95 \times 110 \\ &= 99.275 \end{split}$$

As $U_0(0,\$100)>U_0(1,\$110),$ the present-biased agent will prefer to receive \$100 this week.

Choice 2: Would this present-biased agent prefer \$100 next week (t = 1) or \$110 in two weeks (t = 2)?

The discounted utility of the \$100 next week is:

$$U_0(1,\$100) = u(x_0) + \beta u(x_1) + \beta \delta^2 u(x_2)$$

= $\beta \delta u(\$100)$
= $0.95 \times 0.95 \times 100$
= 90.25

The discounted utility of the \$110 in two weeks is:

$$\begin{split} U_0(2,\$110) &= u(x_0) + \beta u(x_1) + \beta \delta^2 u(x_2) \\ &= \beta \delta^2 u(\$110) \\ &= 0.95 \times 0.95^2 \times 110 \\ &= 94.31 \end{split}$$

As $U_0(1,\$100)=90.25<94.31=U_0(2,\$110),$ the present-biased agent will prefer to receive \$110 in two weeks.

If we consider those two choices by the present-biased agent together, we see the following pattern.

For choice 1, the present-biased agent will prefer \$100 now to \$110 in one week. Their preference for benefits now due to the short-term discount factor β leads them to prefer the immediate payoff.

For choice 2, the present-biased agent will prefer \$110 in two weeks to \$100 in one week. They are willing to wait longer for a larger reward, with both outcomes in the future and subject to the short-term discount factor β .

Consider what would happen if this present-biased agent selected the \$110 in two weeks in Choice 2, but after one week we asked if they would like to reconsider their choice. They are effectively being offered Choice 1. This would then lead them to change their mind and take the immediate \$100.

This combination of decisions is time inconsistent. The present-biased agent's actions are not consistent with their initial plan.

We can see this change in preference in the following diagram.

The vertical bars represent the payments of \$100 and \$110. The lines projecting back from the bars to the y-axis represent the discounted utility of each payment at each time.

There is a kink in the line projecting from the \$110 in two weeks, representing the effect of the short-term discount factor β . Between t = 1 and t = 2 both the short-term and usual discount factors are applied. This leads to that part of the curve having a steeper slope than between t = 0 and t = 1 where only the usual discount factor is applied.

At t = 0 the discounted utility of the \$110 at t = 2 is higher and that payment is therefore preferred. At t = 1 when the \$100 is no longer discounted by the short-term discount factor β , it suddenly becomes more attractive. If offered on that day, would be chosen in substitute of the \$110 due in another week.

25.4 A longer delay

Assume there is a present-biased agent with $\beta = 0.75$, $\delta = 0.9$ and utility each period of $u(x_n) = x_n$.

Would this agent prefer \$10 in five days (t = 5) or \$20 in 10 days (t = 10)?

The discounted utility of the \$10 in five days is:



Figure 25.1: 100 next week or 110 in two weeks?

$$\begin{split} U_0(5,\$10) &= \beta \delta^5 u(\$10) \\ &= 0.75 \times 0.9^5 \times 10 \\ &= 4.43 \end{split}$$

The discounted utility of the \$20 in 10 days is:

$$\begin{split} U_0(10,\$20) &= \beta \delta^{10} u(\$20) \\ &= 0.75 \times 0.9^{10} \times 20 \\ &= 5.23 \end{split}$$

As $U_0(5,\$10) = 4.43 < 5.23 = U_0(10,\$20)$, this present-biased agent will prefer to receive \$20 in 10 days.

Five days pass so we are now at t = 5. We ask the agent if they would like to change their mind.

The discounted utility of the \$10 today is:

$$U_5(5,\$10) = u(\$10)$$

= 10

The discounted utility of the \$20 in five days is:

$$\begin{split} U_5(10,\$20) &= \beta \delta^5 u(\$20) \\ &= 0.75 \times 0.9^5 \times 20 \\ &= 8.86 \end{split}$$

As $U_5(5,\$10) = 10 > 8.86 = U_5(10,\$20)$, this present-biased agent will prefer to receive \$10 today. They have changed their preference between the two payments relative to their decision at t = 0.

We can see this change in preference in the following diagram.

The vertical bars represent the payments of \$10 and \$20. The lines projecting back from the bars to the y-axis represent the discounted utility of each payment

at each time. There is a kink in the line in the period immediately before each payment, representing the effect of the short-term discount factor β .

At t = 0 (and through to t = 4) the discounted utility of the \$20 at t = 10 is higher and that payment is therefore preferred. At t = 5 when the \$10 is no longer discounted by the short-term discount factor β , it suddenly becomes more attractive. If offered on that day, would be chosen in substitute of the \$20 due in another five days.



Figure 25.2: \$10 in five days or \$20 in 10 days?

25.5 Paying a loan

Charlie is a naive present-biased agent with $\beta = 0.5$, $\delta = 0.95$ and u(x) = x. Charlie loaned \$100 to Carol. Carol is due to pay Charlie back in 7 days (at t=7). Carol tells Charlie that she would prefer to pay him back later, and offers 200 in 14 days (at t = 14) if he is willing to wait. Charlie is considering whether to accept Carol's offer.

(a) What does Charlie choose at t = 0?

To determine what Charlie chooses at t = 0, we need to compare the discounted utility of the two options.

The discounted utility of \$100 at t = 7 is:

$$U_0(7, 100) = 0.5 \times 0.95^7 \times 100$$

= 34.92

The discounted utility of \$200 at t = 14 is:

$$\begin{split} U_0(14,200) &= 0.5 \times 0.95^{14} \times 200 \\ &= 48.77 \end{split}$$

Charlie chooses the option with the highest discounted utility, which is \$200 at t = 14.

(b) At t = 7 Charlie considers whether he should now demand payment of \$100 at t = 7 rather than wait for payment of \$200 at t = 14. What does Charlie choose at t = 7?

To determine what Charlie chooses at t = 7, we need to compare the discounted utility of the two options.

The discounted utility of \$100 at t = 7 is:

$$U_7(7,100) = 0.95^0 \times 100$$

= 100

The discounted utility of \$200 at t = 14 is:

$$U_7(14,200) = 0.5 \times 0.95^7 \times 200$$

= 69.83

Charlie chooses the option with the highest discounted utility, which is \$100 at t = 7. He has changed his mind. This is because the sum at t = 7 is no longer subject to the short-term discount factor β .

(c) Draw a graph illustrating Charlie's choices.

The following chart shows each of the two options presented to Charlie, \$100 at t = 7 and \$200 at t = 14. The line extended from each back to t = 0 represents the the discounted utility of each option at time t.

It can be seen that from t = 0 to t = 6, the discounted utility of \$200 at t = 14 is higher than the discounted utility of \$100 at t = 7. However, at t = 7, the discounted utility of \$100 at t = 7 is higher than the discounted utility of \$200 at t = 14. Hence Charlie changes his mind.



Figure 25.3: Paying a loan

Chapter 26

Sophisticated present bias

Summary

- Present-biased agents can be categorised as naive or sophisticated, based on their awareness of future preference reversals.
- Naive agents incorrectly believe their future selves will discount rewards consistently, leading to time-inconsistent choices.
- Sophisticated agents accurately predict their future preference reversals and plan accordingly.
- Naive agents use forward reasoning in decision-making, while sophisticated agents use backward reasoning.
- The following examples show that the decision-making processes of naive and sophisticated agents can lead to different choices, even with identical discount factors.
- Sophisticated agents may sometimes choose worse immediate outcomes to avoid even worse future outcomes due to anticipated preference reversals. They tend to take action earlier than naive agents, a behaviour known as "prepoperation".

26.1 Introduction

Suppose a present-biased agent considers whether they prefer a smaller, sooner payoff or a larger, later payoff. They decide they will wait for the larger, later payoff. Time passes until the smaller pay-off becomes immediately available. They change their mind and take the smaller payoff.

Why did the agent initially choose the larger payoff when they would take the smaller payoff? Were they aware of the likely outcome of their preferences? If

they were aware, they might anticipate changing their mind and plan accordingly.

For many of us, we are aware that we make time-inconsistent decisions. We know that we often cave when faced with immediate temptation. Even if we don't plan to eat a whole packet of cookies for dinner, have more than one beer, or stay up late watching just one more episode, we often do.

26.2 Naïve and sophisticated present-biased agents

To bring the idea that we are aware of our present bias into the analysis, I will distinguish between two types of present-biased agents: naive and sophisticated.

Consider how a present-biased agent discounts a payoff for each successive period of delay under the $\beta\delta$ model. The first period of delay results in the application of a total discount factor of $\beta\delta$. Each additional period of delay results in the application of an additional discount factor of δ . The discount factor applied for the first period of delay is relatively smaller - or the magnitude of the discount relatively greater - than the additional discount applied for any additional period of delay.



A naive present-biased agent believes that the relative difference in discount between periods that they can see today will persist as time passes. That is, they believe that the discount factor of δ applied between, say, periods 2 and 3, will still be the relative discount when they reach period 2.



However, when period 2 does arrive, the outcome in period 3 will be discounted by both the short-term discount factor β and the usual discount factor δ , compared to no discount for the outcome in period 2. As a result, the relative

discount of outcomes at different times will change when one of those outcomes becomes due today.



A sophisticated present-biased agent correctly believes that they will apply both the short-term and usual discount factors in the future. As a result, they understand that the relative discount of different outcomes will change when one of those outcomes becomes due today. They understand that if faced with the temptation to take benefits or to delay costs, they will do so.

A sophisticated and naive person can make different decisions despite having the same level of impatience, δ , and the same level of present bias, β . Any difference emerges in the way that they reason through an intertemporal choice.

The naive agent makes their plans by forward reasoning, starting from today.

1. First, they decide their preferred option for today, period 0, believing that they will stick to their plan in the future.



2. When they move to the next period they recompute their plan, again believing they will stick to the plan once they move to the next period.



3. They repeat this process as they move through time.



In contrast, the sophisticated present-biased agent makes their plans by backward reasoning, starting from the final period.

1. For that final period, they solve for the preferred action.



2. They then consider one period earlier and solve for the preferred action, accounting for the previous decision. This means, that if they calculate that they will cave into temptation at a certain time, they will remove the option to resist from their choice set. Absent the ability to commit to an action, they will resign themselves to giving into temptation in advance.



3. They repeat this process as they move back to today.



26.3 Examples

The difference between naive and sophisticated present-biased agents is best illustrated through examples.

26.3.1 \$100 next week or \$110 in two weeks

For the first example, suppose we have a naive and a sophisticated present-biased agent, each with $\beta = 0.95$, $\delta = 0.95$ and utility each period of $u(x_n) = x_n$.

We offer them both the following choice.

Would you like \$100 next week or \$110 in two weeks?

We also tell them that we will allow them to reconsider their decision next week.

When the naive present-biased agent considers this problem, they simply compare the discounted utility of each payoff from the perspective of today.

The discounted utility of the \$100 next week is:

$$\begin{split} U_0(1,\$100) &= \beta u(x_1) \\ &= \beta \delta u(\$100) \\ &= 0.95 \times 0.95 \times 100 \\ &= 90.25 \end{split}$$

The discounted utility of the \$110 in two weeks is:

$$\begin{split} U_0(2,\$110) &= \beta \delta^2 u(x_2) \\ &= \beta \delta^2 u(\$110) \\ &= 0.95 \times 0.95^2 \times 110 \\ &= 94.31 \end{split}$$

As $U_0(1, \$100) = 90.25 < 94.31 = U_0(2, \$110)$, at t = 0 the present-biased agent will prefer to receive \$110 in two weeks.

But this choice does accord with the naive agent's preferences next week. Next week (at t = 1) they will calculate utility of the \$100 as:

$$U_1(1,\$100) = u(x_1)$$
$$= u(\$100)$$
$$= 100$$

The discounted utility of the \$110 a week later is:

$$\begin{split} U_1(2,\$110) &= \beta \delta u(x_2) \\ &= \beta \delta u(\$110) \\ &= 0.95 \times 0.95 \times 110 \\ &= 99.275 \end{split}$$

As $U_1(1,\$100)>U_1(2,\$110),$ the present-biased agent will prefer to receive \$100 immediately.

At t = 0 their preference is inconsistent with what it will be next week at t = 1.

The sophisticated present-biased agent will approach this decision differently. They consider it using backward induction.

First, they will look at the choice they will face next week and calculate the discounted utility of each option as they will calculate it at that time.

That is, they calculate the discounted utility of the \$100 at t = 1 from the perspective of t = 1.

$$U_1(1,\$100) = u(x_1)$$

= $u(\$100)$
= 100

They then calculate the discounted utility of the \$110 at t = 2 from the perspective of t = 1.

$$\begin{split} U_1(2,\$110) &= \beta \delta u(x_2) \\ &= \beta \delta u(\$110) \\ &= 0.95 \times 0.95 \times 110 \\ &= 99.275 \end{split}$$

As $U_1(1, \$100) > U_1(2, \$110)$, the sophisticated agent sees that next week, at t = 1 they will take the \$100.

This is the same pair of calculations that the naive agent made at t = 1. The difference is that the sophisticated agent makes this calculation at t = 0 from the perspective of t = 1. The naive agent does not consider this perspective until t = 1.

After considering their preference at t = 1, the sophisticated agent then considers their choice at t = 0. They see their future decision at t = 1 and know that \$110 in two weeks is not available to them. The only option they have is to choose \$100 in one week. They effectively accept their future present bias now and choose the \$100 in one week.

In this example, being naive or sophisticated does not change their final choice. It only changes their beliefs about their final decision over time. The sophisticated agent knows at t = 0 what they will do at t = 1. The naive agent is unaware that at t = 1 they will make a decision inconsistent with their choice at t = 0. We can summarise their decisions at each time as follows:

	Naive agent	Sophisticated agent
$\begin{aligned} t &= 0\\ t &= 1 \end{aligned}$	\$110 at $t = 2$ \$100 at $t = 1$	\$100 at $t = 1$ \$100 at $t = 1$

26.3.2 Watching a movie

Suppose we have a naive and a sophisticated present-biased agent, each with $\beta = 0.5$ and $\delta = 1$. They are present-biased, but beyond that present bias demonstrate no impatience.

We offer them the following choice.

- An OK movie today (t = 0) that gives utility of 6
- A good movie next week (t = 1) that gives utility of 10
- A great movie in two weeks (t = 2) that gives utility of 16

We also tell the agents that next week they will be offered an opportunity to change their minds.

First, we consider the naive agent. They calculate utility from the perspective of today.

$$U_0(0, \text{OK}) = u(\text{OK})$$

= 6
$$U_0(1, \text{good}) = \beta \delta u(\text{good})$$

= 0.5 × 1 × 10
= 5
$$U_0(2, \text{great}) = \beta \delta^2 u(\text{great})$$

= 0.5 × 1² × 16
= 8

As $U_0(2, \text{great}) > U_0(0, \text{OK}) > U_0(1, \text{good})$, the naive agent will choose the great movie in two weeks.

But what then happens when the naive agent is given the chance to change their mind after one week?

$$\begin{split} U_1(1, \mathrm{good}) &= u(\mathrm{good}) \\ &= 10 \\ U_1(2, \mathrm{great}) &= \beta \delta u(\mathrm{great}) \\ &= 0.5 \times 1 \times 16 \\ &= 8 \end{split}$$

As $U_1(1, \text{good}) > U_1(2, \text{great})$, the naive agent will change their mind and watch the good movie immediately.

What of our sophisticated agent?

They will make their decision today based on correct beliefs about their future preferences. To do this, they solve by backward induction. First, what will their decision be next week?

$$\begin{split} U_1(1, \mathrm{good}) &= u(\mathrm{good}) \\ &= 10 \\ U_1(2, \mathrm{great}) &= \beta \delta u(\mathrm{great}) \\ &= 0.5 \times 1 \times 16 \\ &= 8 \end{split}$$

As $U_1(1, \text{good}) > U_1(2, \text{great})$, the sophisticated agent can see that they will choose to watch the good movie immediately.

Knowing this is the case, the sophisticated agent now decides whether they prefer the OK movie today or the good movie next week.

$$\begin{split} U_0(0,\mathrm{OK}) &= u(\mathrm{OK}) \\ &= 6 \\ U_0(1,\mathrm{good}) &= \beta \delta u(\mathrm{good}) \\ &= 0.5 \times 1 \times 10 \\ &= 5 \end{split}$$

As $U_0(0,{\rm OK})>U_0(1,{\rm good}),$ the sophisticated agent prefers the OK movie today.

From today's perspective, the sophisticated agent would prefer the great movie in two weeks, but as they know they will cave in to their present bias next week and watch the good movie, they make today's decision on that basis. They know that if they select the great movie today they won't watch it.

26.3.3 A library fine

A naive present-biased agent has failed to return their library books and is fined at t = 0. They can select one of the following payment schemes:

(0, -\$10), (1, -\$15) or (2, -\$25)

That is, they can pay \$10 today, \$15 at t = 1 or \$25 at t = 2.

They are free to change the scheme over time as they see fit.

The agent's utility is linear in money - that is, $u(x_n)=x_n$ - with discount factors $\beta=0.5$ and $\delta=1.$

The naive agent calculates the utility of each option today.

$$\begin{split} U_0(0,-\$10) &= u(-\$10) \\ &= -10 \\ \\ U_0(1,-\$15) &= \beta \delta u(-\$15) \\ &= 0.5 \times 1 \times (-15) \\ &= -7.5 \\ \\ U_0(2,-\$25) &= \beta \delta^2 u(-\$25) \\ &= 0.5 \times 1^2 \times (-25) \\ &= -12.5 \end{split}$$

As $U_0(1, -\$15) > U_0(0, -\$10) > U_0(2, -\$25)$, the naive agent will choose to pay \$15 at t = 1.

A week passes and the naive agent is now at t = 1, the time when they were planning to pay the fine. The naive agent reconsiders their decision.

$$\begin{split} U_1(1,-\$15) &= u(-\$15) \\ &= -15 \\ U_1(2,-\$25) &= \beta \delta u(-\$25) \\ &= 0.5 \times 1 \times (-25) \\ &= -12.5 \end{split}$$

As $U_1(2, -\$25) = -12.5 \ge -15 = U_1(1, -\$15)$, the naive agent changes their decision and further delays their payment. They now choose to pay \$25 at t = 2.

When they reach t = 2, they have no choice but to pay the \$25.

A sophisticated present-biased agent has also failed to return their library books and is fined at t = 0. They can select one of the following payment schemes:

(0, -\$10), (1, -\$15) or (2, -\$25)

They are free to change the scheme over time as they see fit.

The sophisticated agent's utility is linear in money $u(x_n) = x_n$, with discount factors $\beta = 0.5$ and $\delta = 1$.
For the sophisticated agent, we start calculating utility from the final period.

At t = 2, they have no choice but to pay the \$25.

What of t = 1?

$$\begin{split} U_1(1,-\$15) &= u(-\$15) \\ &= -15 \\ U_1(2,-\$25) &= \beta \delta u(-\$25) \\ &= 0.5 \times 1 \times (-25) \\ &= -12.5 \end{split}$$

The sophisticated agent can see that if they choose at t = 1, they will choose to pay \$25 at t = 2.

Now we iterate at t = 0. The sophisticated agent only compares \$10 at t = 0 with \$25 at t = 2 because they know that at t = 1 they will select \$25 at t = 2. They know that if they delay once they will delay again and end up paying the largest fine. Hence they limit their comparison to those outcomes that might occur:

$$\begin{split} U_0(0,-\$10) &= u(-\$10) \\ &= -10 \\ U_0(2,-\$25) &= \beta \delta^2 u(-\$25) \\ &= 0.5 \times 1^2 \times (-25) \\ &= -12.5 \end{split}$$

As $U_0(0, -\$10) > U_0(2, -\$25)$, the sophisticated agent will choose to pay \$10 at t = 0.

In the examples, we have seen two contrasting outcomes for the sophisticated agent.

In the movie example, they watch an OK movie today, rather than a good movie in one week or a great movie in two, because they knew that they would watch the good movie in one week if they delayed to watch the great movie. As a result, they watched an earlier, worse movie than the naive agent. In the library fine example, they paid the lowest possible fine as they saw they would further delay paying the fine in the future, leading it to increase even more.

The sophisticated agent's behaviour arises from two tensions:

- 1. They understand that they will not be able to wait for the optimal time.
- 2. They are present-biased so they prefer benefits today and costs delayed to the future.

In combination, these imply a sophisticated agent will generally take action earlier than the naive agent. They can "prepoperate", which is doing something now when it would be better to wait.

Chapter 27

Commitment

Summary

- Sophisticated present-biased agents can foresee their future actions. This foresight allows them to use commitment devices mechanisms that lock them into a course of action by changing the value or availability of future options.
- Commitment devices can work through three main channels:
 - Forcing the agent to maintain the optimal course of action.
 - Depressing the value of the undesired action
 - Increasing the value of the optimal action
- The following examples illustrate how these commitment devices work in practice:
 - Forcing optimal action: Buying a non-refundable movie ticket, Odysseus tying himself to the mast, and using lay-by schemes.
 - Depressing the value of bad actions: Using platforms like stickK to impose penalties for failing to meet goals.
 - Increasing the value of optimal actions: Temptation bundling, such as pairing exercise with audiobooks or massages.

27.1 Introduction

A sophisticated present-biased agent can foresee their future actions and adjust their decisions today based on their foresight.

This foresight provides an opportunity. By seeing their future selves fail, they can commit themselves to a course of action today that they would not otherwise be able to choose.

27.2 Commitment device

People often implement the opportunity to commit themselves to a course of action by using a commitment device.

A commitment device is a mechanism that locks you into a course of action by changing the value or availability of future options.

Commitment devices may work through the following channels:

- They may depress the value of the bad course of action.
- They may increase the value of the optimal course of action
- They may force the agent to maintain the optimal course of action.

27.3 Commitment examples

I will now illustrate those channels with examples.

27.3.1 Forcing the optimal course of action

A sophisticated present-biased agent with $\beta = 0.5$ and $\delta = 1$ is choosing between three movies, an OK movie, a good movie, and a great movie. The OK movie would give utility of 6, the good movie would give utility of 10, and the great movie would give utility of 16.

The problem is that only the OK movie is showing today (t = 0). The good movie is showing next week (t = 1), and the great movie is showing in two weeks (t = 2).

The agent has enough money and time to watch only one movie. Should they watch the OK movie today or wait for the good or great movie?

To determine their action, they solve by backward induction.

First, what will their decision be next week?

$$\begin{array}{l} U_1(1, \mathrm{good}) = u(\mathrm{good}) \\ = 10 \end{array}$$

$$U_1(2, \text{great}) = \beta \delta u(\text{great})$$
$$= 0.5 \times 1 \times 16$$
$$= 8$$

As $U_1(1, \text{good}) > U_1(2, \text{great})$, the sophisticated agent can see that they will choose to watch the good movie immediately.

Knowing this is the case, the sophisticated agent now decides whether they prefer the OK movie today or the good movie next week.

$$\begin{split} U_0(0,\mathrm{OK}) &= u(\mathrm{OK}) \\ &= 6 \\ U_0(1,\mathrm{good}) &= \beta \delta u(\mathrm{good}) \\ &= 0.5 \times 1 \times 10 \\ &= 5 \end{split}$$

As $U_0(0, \text{OK}) > U_0(1, \text{good})$, the sophisticated agent prefers the OK movie today.

But when they compare all three options at t = 0, they would prefer the great movie in two weeks.

$$\begin{split} U_0(2, \text{great}) &= \beta \delta^2 u(\text{great}) \\ &= 0.5 \times 1^2 \times 16 \\ &= 8 \\ &\geq 6 = U_0(0, \text{OK}) \end{split}$$

It is only because they can foresee their future failing after one week that they don't wait for the great movie.

But what if they could commit themselves today? For example, suppose they could purchase a non-refundable, non-resalable ticket to the great movie in two weeks. The result is that a sophisticated present-biased agent would buy a ticket to the great movie in two weeks and prevent their future self from changing their action.

27.3.2 Forcing the optimal course of action: Odysseus

Another example of forcing the agent to maintain the optimal course of action is the story of Odysseus. Odysseus was required to sail past the sirens at t = 1. At that time he can either jump off his ship to join the sirens (and die) or sail on past and live, having a great life at t = 2.

As Odysseus is a sophisticated present-biased agent, he considers what he is likely to do at t = 1. He realises that he will jump off the ship for the immediate benefit of joining the sirens at the loss of the longer-term discounted benefit of living.

But from the perspective of Odysseus today, at t = 0, with both the benefits of the sirens and living in the future and therefore discounted, Odysseus would prefer to live. As a result, he decides to commit himself to that course of action by instructing his crew to the him to the mast (plus leaving his ears unplugged so that he also gets some benefits from the siren song).

27.3.3 Forcing the optimal course of action: lay-by

Suppose a quasi-hyperbolic discounting agent with discount factors $\beta = 1/2$ and $\delta = 1$ wants a new jacket for work. They need to save for three months to purchase the jacket. But each month they save they forgo consumption that would boost their utility.

They receive the following payoffs for each action:

	t=0	t=1	t=2
Save	0	0	45
Spend	10	10	10

First, consider what a naive quasi-hyperbolic discounting agent chooses. At t = 0, they calculate the discounted utility of each option.

$$\begin{split} U_0(\text{Save}) &= 0 + \beta\delta \times 0 + \beta\delta^2 \times 45 \\ &= 0.5 \times 45 \\ &= 22.5 \end{split}$$

$$U_0(\text{Spend}) = 10 + \beta \delta \times 10 + \beta \delta^2 \times 10$$
$$= 10 + 0.5 \times 10 + 0.5 \times 10$$
$$= 20$$

At t = 0, $U_0(\text{Save}) > U_0(\text{Spend})$. The agent plans to save for the jacket.

One month now passes. The agent has saved for a month. They could now spend their savings from last month and this month, giving a short-term boost, or keep saving for their jacket. The payoffs for each action are:

	t=1	t=2
Save	0	45
Start spending at $t = 1$	20	10

Again, the naive agent calculates the discounted utility of each option.

$$U_1(\text{Save}) = 0 + \beta \delta \times 45$$
$$= 0.5 \times 45$$
$$= 22.5$$

$$U_1(\text{Start spending at t=1}) = 20 + \beta \delta \times 10$$
$$= 20 + 0.5 \times 10$$
$$= 25$$

At t = 1 spending now has the highest discounted utility. After saving for the first period, the agent spends despite initially wanting to save.

Now let's consider this problem from the point of view of a sophisticated presentbiased agent. They see their full choice set as:

	t=0	t=1	t=2
Save	0	0	45
Start spending at t=1	0	20	10
Spend	10	10	10

The sophisticated agent works backward through time. At t = 2, if they have saved, the agent will buy the jacket. Otherwise, the agent will spend.

The discounted utility of the options at t = 1 is as follows:

 $U_1(\text{Save}) = 0 + \beta \delta \times 45$ $= 0.5 \times 45$ = 22.5

 $U_1({\rm Start\ spending\ at\ }t=1)=20+\beta\delta\times 10$ $=20+0.5\times 10$ =25

As $U_1(\text{Start spending at } t = 1) > U_1(\text{Save})$, the sophisticated agent would spend.

The sophisticated agent now knows that saving for the jacket is not available to them as they will spend at t = 1 regardless of their initial action. There was no need to consider the option to start spending at t = 0 as if they had spent then there is no other choice at t = 1.

The sophisticated agent now chooses between the two feasible options at t = 0:

$$\begin{split} U_0(\text{Start spending at } t=1) &= 0 + \beta\delta \times 20 + \beta\delta^2 \times 10 \\ &= 0.5 \times 20 + 0.5 \times 10 \\ &= 15 \end{split}$$

$$U_0(\text{Spend}) = 10 + \beta \delta \times 10 + \beta \delta^2 \times 10$$
$$= 10 + 0.5 \times 10 + 0.5 \times 10$$
$$= 20$$

As $U_0(\text{Spend}) < U_0(\text{Start spending at } t = 1)$, they start to spend at t = 0. Contrast this with the naïve agent who chooses Save at t = 0.

Now consider what the sophisticated agent may do in the presence of lay-by. Lay-by involves paying a deposit and administrative fee toward the purchase of a product. You receive your purchase when you make payment in full later.

For payment of an administrative fee equivalent to 1 unit of utility and an initial deposit (the agent's savings in t = 0), the agent can reserve the jacket, preventing them from spending that money at t = 2. The new set of options is:

	t=0	t=1	t=2
Save	0	0	45
Start spending at t=1	0	20	10
Spend	10	10	10
Lay-by	-1	0	45

The lay-by option is strictly worse than Save. But what happens if lay-by is available to the sophisticated present-biased agent?

Again, working backward, at t = 2, the agent buys the jacket if they have saved and otherwise spends.

At t = 1, by the same logic we looked at previously, they spend if they can do so. That eliminates Save from their choice set.

At t = 0, they now compare the feasible options:

$$\begin{split} U_0(\text{Start spending at } t = 1) &= 0 + \beta\delta \times 20 + \beta\delta^2 \times 10 \\ &= 0.5 \times 20 + 0.5 \times 10 \\ &= 15 \end{split}$$

$$\begin{split} U_0(\text{Spend}) &= 10 + \beta\delta \times 10 + \beta\delta^2 \times 10 \\ &= 10 + 0.5 \times 10 + 0.5 \times 10 \\ &= 20 \end{split}$$

$$\begin{split} U_0(\text{Lay-by}) &= -1 + \beta\delta \times 0 + \beta\delta^2 \times 45 \\ &= -1 + 0.5 \times 0 + 0.5 \times 45 \\ &= 21.5 \end{split}$$

As $U_0(\text{Lay-by}) > U_0(\text{Spend}) > U_0(\text{Start spending at } t = 1)$, lay-by is the preferred option at t = 0. As it binds the agent in the future, they can stick to this plan.

Note that lay-by is strictly worse than Save as the agent must pay the administrative fee. But the sophisticated quasi-hyperbolic discounting agent chooses it as the only feasible way to get their jacket. Without lay-by they know they will start spending at t = 1 and end up with lower utility from the perspective of their t = 0 self.

27.3.4 Depressing the value of the bad course of action

stickK is a online platform that enables people to commit to future courses of action. stickK works as follows:

First, you state a time-based goal, such as not smoking during the next month, losing 5kg over the next 90 days, or writing the next chapter of your PhD thesis by Christmas.

Second, you commit a stake that will be paid to a charity (or an anti-charity) if you fail to meet your goal.

At the stated time, you then report (or a referee appointed by you reports) whether you have met your goal. If you fail to report or report that you failed to meet your goal, your credit card is debited by the staked amount.



I used stickK during my PhD. I regularly set deadlines for tasks and staked, say, a payment of \$500 to the National Rifle Association. If I didn't complete the task, I lost the money. Throughout the PhD, I met my benchmarks every time except for one.

To understand how stick works in the context of the $\beta\delta$ model, consider a sophisticated present-biased agent who is weighing the enjoyment they get from smoking versus the long-term health effects.

This sophisticated present-biased agent has $\beta = 1/2$ and $\delta = 1$. They enjoy smoking, which gives them utility of 5. However, the agent also likes being healthy. Higher health gives utility of 8.

At t = 0 the agent is deciding whether to smoke over the next month (t = 1). If the agent doesn't smoke, they will have better health at t = 2.

The sophisticated agent works backward through time. At t = 1 its payoffs are:

```
U_1(\text{smoking}) = 5U_1(\text{healthy}) = \beta \delta \times 8= 4
```

The agent decides to smoke.

As a result, at t = 0, knowing that it will cave at t = 1, the agent doesn't bother committing to not smoking, even though from the perspective of t = 0 refraining from smoking is the better option:

$$U_0(\text{smoking}) = \beta \delta \times 5$$

= 2.5
$$U_0(\text{healthy}) = \beta \delta^2 \times 8$$

= 4

But now suppose the agent learns about stickK. The agent has the option of staking a sum at t = 0 to prevent it from smoking. The agent decides to stake an amount equivalent to utility 5 that would be incurred at t = 2.

Working backward through time, the agent knows that at t = 1 if it has not staked any money with stickK, it will smoke. But what if it has?

$$\begin{split} U_1(\text{stickK+smoking}) &= u(\text{smoking}) + \beta\delta \times u(\text{lost stake}) \\ &= 5 + \beta\delta \times (-5) \\ &= 2.5 \end{split}$$

$$\begin{split} U_1(\text{stickK+healthy}) &= \beta\delta \times u(\text{healthy}) \\ &= \beta\delta \times 8 \\ &= 4 \end{split}$$

The agent would refrain from smoking at t = 1 due to the penalty they would have to pay.

This means the agent's options at t = 0 are effectively a comparison between smoking and using stickK to commit to not smoking. The discounted utility of each option is as follows.

$$\begin{split} U_0(\text{smoking}) &= \beta\delta \times u(\text{smoking}) \\ &= \beta\delta \times 5 \\ &= 2.5 \end{split}$$

$$\begin{split} U_0(\text{stickK+healthy}) &= \beta \delta^2 \times u(\text{healthy}) \\ &= \beta \delta^2 \times 8 \\ &= 4 \end{split}$$

As $U_0(\text{stickK+healthy}) > U_0(\text{smoking})$, the agent chooses to commit using stickK.

One interesting feature of these two options at t = 0 is that the penalty does not appear in either calculation of the discounted utility. We have already calculated that if the penalty is present at t = 1, the agent will not smoke. So when they consider their options at t = 1, there is no cost to committing.

For any problem of this form, the agent could always successfully use stickK to commit to any action. The agent just needs to make the stake high enough.

27.3.5 Increasing the value of the optimal action: temptation bundling

"Temptation bundling" involves increasing the value of the optimal course of action by adding a temptation to that course.

Consider the following example.

Beth has the choice between exercising and watching television today at t = 0.

- Beth does not enjoy exercise, which gives utility of 0. However, exercise leads her to be healthy, giving utility of 12 in the future at t = 1.
- Beth enjoys watching television, which gives utility of 6. However, watching television leads her to be unhealthy with utility of 0 in the future at t = 1.

Beth is sophisticated and discounts the future quasi-hyperbolically. Beth's $\beta = 1/2$ and her $\delta = 2/3$. Does Beth exercise?

Beth works through the options by backward induction. At t = 1 there is no choice to be made, as Beth simply enjoys the utility of the action she chose at t = 0.

For t = 0 she calculates the discounted utility of each option as follows:

$$\begin{split} U_0(\text{exercise}) &= u(\text{exercise}) + \beta \delta u(\text{healthy}) \\ &= 0 + \frac{1}{2} \times \frac{2}{3} \times 12 \\ &= 4 \\ \\ U_0(\text{television}) &= u(\text{television}) + \beta \delta u(\text{unhealthy}) \end{split}$$

$$= 6 + \frac{1}{2} \times \frac{2}{3} \times 0$$
$$= 6$$

Beth chooses to watch television.

But what if Beth loved massages and remembers she has a massage voucher she has been saving? What if she decides that if she exercises today, she will go for a massage straight after? Let us assume that the utility of a massage is 3.

The discounted utility of exercising is now:

 $U_0(\text{exercise}+\text{massage}) = u(\text{exercise}) + u(\text{massage}) + \beta \delta u(\text{healthy})$

$$= 3 + \frac{1}{2} \times \frac{2}{3} \times 12$$
$$= 7$$

 $U_0(\text{exercise}+\text{massage}) > U_0(\text{television})$, so Beth now chooses to exercise.

This example is not strictly a commitment device as Beth could cheat. Beth could watch television and get the massage. But people are often good at creating "mental accounts" by which they make certain money or activities out-of-bounds unless certain conditions are met.

27.3.6 Gym attendance

One empirical illustration of temptation bundling comes from an experiment to increase gym attendance.

Kirgios et al. (2020) found that teaching gym-goers how to temptation bundle with a free audiobook boosts gym visits. Simply receiving a free audiobook with no explicit instruction boosts exercise, implying that people who are given audiobooks by gyms can infer they should temptation bundle. However, teaching temptation bundling modestly outperforms simply giving gym-goers free audiobooks.

Chapter 28

Delayed gratification, spread and variation

Recall the assumption of utility independence:

All that matters is maximising the sum of discounted utilities. Decision-makers are assumed to have no preference for the distribution of utilities.

However, there is evidence that people care about the shape of the utility stream over time. There is evidence that people delay gratification, and prefer spread and variation. They don't care solely about maximising discounted utility.

This evidence suggests that the assumption of utility independence does not hold. I will now discuss these three bodies of evidence.

28.1 Delayed gratification

The first concerns delayed gratification.

Consider the following example:

It is a sunny weekend. You can either study today and go to the beach tomorrow, or you can go to the beach today and study tomorrow.

Studying gives you a utility of 10. Going to the beach gives you a utility of 20.

What would an exponential discounter with $\delta = 0.8$ do?

What would a present-biased agent with $\beta = 0.5$ and $\delta = 0.8$ do?

Both agents would go to the beach today and study tomorrow. They will always schedule pleasant tasks before unpleasant tasks.

Does this match people's observed behaviour?

There is considerable evidence that people will schedule unpleasant tasks first and pleasant ones later. This might be thought of as a preference for an increasing utility profile.

How could this be possible for someone who discounts the future?

One way is to ease the requirement that δ be less than one. This provides a solution to the weekend problem but also leads to the potential of endlessly postponing pleasant experiences.

Easing this requirement also clashes with other evidence that people often postpone unpleasant tasks and that people have a δ much less than one for many decisions.

28.2 Spread of utility

Another body of evidence suggests that we prefer a spread of utility. We like to distribute pleasant experiences over time.

In part, this emerges from diminishing marginal utility. Additional units of a good or service on a day when we already have ample will provide less utility than on a day when we have little.

However, some of the evidence cannot be accounted for by diminishing marginal utility.

For example, suppose someone plans to catch up with one friend over lunch and another at dinner. Some people prefer these two events on different days, giving them a spread of utility over time.

28.3 Variation

We also have a preference for variation. Consider the following:

Your favourite meal is lasagna. Your second favourite meal is spaghetti bolognese. Your third favourite is fish and chips.

You are offered the following two options:

- 1. Lasagna every night for the next week.
- 2. Alternating meals of lasagna, spaghetti and fish and chips.

We don't choose to have the same good or service over and over.

Chapter 29

Intertemporal choice applications

Summary

- The following examples demonstrate how understanding intertemporal choice can inform policy design and product development to help individuals overcome present bias and achieve long-term goals.
- People willingly choose illiquid savings accounts with equal or lower interest rates, suggesting demand for commitment devices among sophisticated, present-biased agents.
- A voluntary commitment product involving forfeitable deposits increased the likelihood of quitting smoking, demonstrating the effectiveness of commitment devices in behaviour change.
- Default effects significantly influence organ donor registration rates, potentially due to present bias or rational cost-benefit calculations.
- The Save More Tomorrow program successfully increases retirement savings by mitigating loss aversion and present bias.
- ANZ Bank's Progress Saver and Term Deposit accounts illustrate how present bias and the demand for commitment devices influence consumer financial decisions.

29.1 Introduction

In this part, I discuss several applications of the intertemporal choice concepts we have covered.

29.2 Savings

The first example relates to the use of commitment devices to increase savings.

Beshears et al. (2020) offered experimental participants the opportunity to save in two accounts, one liquid and the other with liquidity constraints such as withdrawal penalties. They found that the experimental participants put nearly half of their money in the illiquid account even though it paid the same interest rate. This behaviour contrasts with the standard economic prediction that all money should go to the liquid account, which dominates the illiquid account's features. Even when the interest rate on the illiquid account was lower, it still attracted around a quarter of the money.

This extract from Table 3 in the paper shows the proportion of funds allocated to each type of "commitment account" when experimental participants had a choice between an account with no liquidity constraints paying 22% interest and the commitment account.

Withdrawal restrictions on commitment account prior to commitment date	Commitment account interest rate		
	21%	22%	23%
10% early withdrawal penalty	27.6	38.9	58.2
	(2.8)	(3.4)	(3.4)
20% early withdrawal popular	-	44.8	61.1
20% earry withdrawar penaity		(3.4)	(3.4)
No oppy withdrawale	-	56.0	59.9
NO Early Withdrawais		(4.1)	(3.6)

You can see that where the interest rates between the liquid and illiquid accounts were equal, the accounts with harsher constraints attracted more money. The account with a higher withdrawal penalty (20% compared to 10%) attracted more money, and the account that barred withdrawals attracted even more.

This result suggests a demand among sophisticated, present-biased agents for products that will enable them to control their future behaviours.

This behaviour has also been observed outside the lab.

Ashraf et al. (2006) offered a commitment savings product called SEED (Save, Earn, Enjoy Deposits) to randomly chosen clients of a Philippine bank. SEED restricted access to savings for one year.

Other than providing a possible commitment savings device, no further benefit accrued to individuals with this account.

Despite this, 28% of participants took up a commitment savings product. The average savings balance increased by 42% after six months and 82% after one year.

29.3 Smoking

The following example relates to quitting smoking.

Giné et al. (2010) tested a voluntary commitment product for smoking cessation.

Smokers were offered a product that comprised a savings account in which they deposit funds. After six months, they took urine tests for nicotine and cotinine.

If they passed, the money was returned. Otherwise, it was forfeited.

The result was that 11% of smokers offered the product took it up. Smokers offered the product were 3 percentage points more likely to pass the 6-month test than the control group

The effect persisted in a surprise test at 12 months.

29.4 Organ donation

The next example concerns organ donation.

In European countries there are registers of people who will donate their organs in case of death. There is a considerable gap in the percentage of registered organ donors between countries. Why?



Effective consent rates, by country. Explicit consent (opt-in, gold) and presumed consent (optout, blue).

Figure from E. J. Johnson and Goldstein (2003).

Yellow countries have an opt-in policy. People are required to register as an organ donor.

Green countries have an opt-out policy known as presumed consent. Citizens are presumed to consent unless they opt out (often through submission of a form). This outcome is a classic example of a default effect. Defaults are sticky. The stickiness of defaults is typically assumed to come from loss aversion or present bias.

Present bias might influence the decision as follows. There is an immediate cost of changing from the default, be that time, effort or money. That cost is not discounted. In contrast, the future benefit of their action is discounted.

If you asked these people about their future plans, they might say they intend to change their organ donation registration. In the hypothetical, both the costs and benefits and in the future. They believe they will switch later.

However, it is possible to argue that the stickiness of the defaults is due to a rational cost-benefit calculation rather than loss aversion or present bias. The cost of opting out is real, and people may not have a strong preference about whether they are an organ donor.

Further, registration does not mean that your organs will be donated. Other factors, such as family preference, affect donation. Among other things, the absence of any active consent in situations of presumed consent means that the family cannot take the organ donation register as an indication of the deceased's wishes. There is little benefit in changing the registration if it will have little effect on the outcome you care about (actual organ donation).

E. J. Johnson and Goldstein (2003) argued that there is a positive relationship between an opt-out policy and organ donations. But, it is much weaker than the registration numbers would suggest and based on a simple regression that likely does not capture all relevant variables.

There are alternatives to presumed consent that may increase organ donation rates.

One alternative is using defaults more transparently with easy opt out. For example, when obtaining your driver's licence, there could be a section stating: "Please tick this box if you do not wish to be registered as an organ donor". This measure would likely increase registration over the alternative of asking people to tick the box if they wish to be an organ donor.

Another alternative is "active choice", where citizens are required to indicate whether or not they wish to be registered. This choice could also be built into a form such as a driver's licence application or renewal.

Your organ donation preferences	
 I want to donate everything below I only want to donate the following: Bone Tissue Corneas (Eye Tissue) Heart Heart Valves Kidneys 	Liver Lungs Pancreas Skin Tissue
 I want to record my decision not to be an organ donor Remove me from the Australian Organ Donor Register 	
 Selecting Save donation preferences means you consent to: The use by and disclosure of your personal information to organ donation. 	Services Australia to search for existing registrations. Authorised Medical Personnel for the purposes of
Save donation preferences Cancel	

29.5 Save More Tomorrow

The next example of intertemporal choice relates to retirement savings.

The Save More Tomorrow program, designed by Richard H. Thaler and Benartzi (2004), combines prospect theory and time preference principles to increase retirement savings.

Under Save More Tomorrow, customers are asked to commit in advance to allocating a fraction of their future salary increases toward their retirement savings accounts.

Save More Tomorrow is designed to reduce loss aversion when deciding contribution amounts. A commitment of a proportion of a pay rise means that the contribution can increase over time, but pay never decreases.

The program is designed to reduce the effect of present bias. The cost of the savings is in the future, meaning that the costs are subject to the short-term discount factor and not disproportionately overweighted relative to the benefits.

The program capitalises on participants' propensity to stick with the status quo, as people are unlikely to unwind their future commitments despite being able to opt out at any time. That ability to opt out also reduces regret and disappointment aversion.

The first tests of the Save More Tomorrow program by Richard H. Thaler and Benartzi (2004) resulted in 78 per cent of those offered the plan joining, 80 per cent of those remaining in the plan through the fourth pay rise, and average savings rates increasing from 3.5 per cent to 13.6 per cent over 40 months. This compares to much lower savings rates by those who declined advice, accepted a recommended savings rate or took advice but declined to enrol in Save More Tomorrow.

Figure 29.1 illustrates the results. Along the horizontal axis are the four groups: those who declined to enrol in the program, those who declined to receive advice, those who accepted a recommended savings rate, and those who accepted the Save More Tomorrow program. The vertical axis shows the percentage of income saved at each of five measurement points; before they received advice and after the following four raises.



Figure 29.1: Savings rates for SMarT

Note the savings rate is higher than the default rate in Australia. Could the default in Australia create a low anchor for some people?

29.6 The Progress Saver Account

Another applied example of how we can use intertemporal choice in an applied setting concerns ANZ bank's progress saver account.

ANZ bank's progress saver account pays bonus monthly interest on the condition that a customer deposits at least \$10 into the account and makes no withdrawals. The bonus interest was 3.74% per year at the time of writing. If you fail to make the minimum deposit or withdraw from the account, you are paid a nominal interest rate of 0.01% per year on your savings for the month.



Many account holders do not receive bonus interest each month. Most notably, while they often make deposits, they later withdraw funds, leading to the loss of the bonus interest.

This behaviour may be evidence of a preference reversal due to present bias. Today, the customer may have a preference for saving money for a long-term goal rather than short-term spending at a closer date. But when that opportunity for spending arrives, they prefer withdrawing from the account and spending the money now. That spending is no longer subject to the short-term discount factor β , whereas the long-term savings goal and any interest toward achieving it are.

For example, suppose a customer is deciding whether to save money for their house deposit far in the future or spend the money on a new pair of shoes when they go shopping next week. Both are in the future and are subject to a short-term discount factor. In that circumstance, saving for the deposit might be preferred. However, when the shopping day comes, the shoes can be bought **now**. The shoes are not subject to the short-term discount factor and may give higher discounted utility than the long-term savings goal. The customer then withdraws the funds for the shoes, having changed their mind.

The behaviour could also be for rational reasons, such as a change in circumstances.

The bank has another product, a term deposit savings account, that pays, at the time of writing, 0.15% interest if you commit your money for 12 months. Many customers still use the term deposit despite paying much lower interest than the progress saver account and constraining access to funds.



Why do people use this apparently sub-optimal product?

Some customers are what we call "sophisticated" present-biased agents. They are present biased, but they know they are present biased. They can see their future failings as they think through problems using backward induction. As a result, they can implement strategies to restrain their future self, such as a commitment device. A commitment device is a mechanism that locks you into a course of action by changing the value or availability of future options.

If they foresaw spending their money on shoes, they would know that depositing in the progress saver account would not lead to them saving for their house deposit long term. As a result, that person may decide to forgo the possibility of higher interest (that they won't receive) to constrain their future self from buying shoes when shopping. The term deposit provides that constraint, acting as a commitment device.

Chapter 30

Intertemporal choice exercises

30.1 Exercise or television?

Olga and Paul can choose one of the following options:

- Exercising at t = 0 (utility = 0) and being healthy at t = 1 (utility = 30).
- Watching television at t = 0 (utility = 15) and being unhealthy at t = 1 (utility = 0).

a) Olga discounts the future exponentially with $\delta = 2/3$. At t = 0, what is Olga's utility of exercising and watching television? What does Olga do?

Answer $U_0(exercise) = u(x_0) + \delta u(x_1)$ $= 0 + \frac{2}{3} \times 30$ = 20 $U_0(television) = u(x_0) + \delta u(x_1)$ $= 15 + \frac{2}{3} \times 0$ = 15Olga gets higher discounted utility from exercising, so chooses to exercise. b) Paul discounts the future quasi-hyperbolically with $\beta = 3/4$ and $\delta = 2/3$. At t = 0, what is Paul's utility of exercising and watching television? What does Paul do?

Answer

$$\begin{split} U_0(exercise) &= u(x_0) + \beta \delta u(x_1) \\ &= 0 + \frac{3}{4} \times \frac{2}{3} \times 30 \\ &= 15 \\ U_0(television) &= u(x_0) + \beta \delta u(x_1) \\ &= 6 + \frac{3}{4} \times \frac{2}{3} \times 0 \\ &= 15 \end{split}$$

Paul is indifferent between the two options, so could choose either.

30.2 Today or tomorrow?

Terry and Andy are given the choice between the following three options:

- A (utility of 3 at t = 0)
- B (utility of 4 at t = 1)
- C (utility of 5 at t = 2).

a) Suppose that Terry discounts the future exponentially with $0 < \delta < 1$. He is indifferent between A and B at t = 0. What does this tell you about Terry's δ ?

? Answer

The utility of A and B must be equal.

$$U_0(A) = U_0(B)$$
$$3 = \delta 4$$
$$\delta = \frac{3}{4}$$

b) Andy discounts the future quasi-hyperbolically with $0 < \beta < 1$ and $0 < \delta < 1$. At t = 0, Andy is indifferent between A and B. What does this tell you about Andy's β and $\delta?$

? Answer

The utility of A and B must be equal.

$$U_0(A) = U_0(B)$$
$$3 = \beta \delta 4$$
$$\beta \delta = \frac{3}{4}$$

We cannot determine anything else about the two as the discounting between t = 0 and t = 1 is a function of both the short-term and exponential discount factor.

c) At t = 0, Andy is indifferent between B and C. What does this tell you about Andy's β and δ ?

? Answer

The utility of B and C must be equal.

$$U_0(B) = U_0(C)$$
$$\beta \delta 4 = \beta \delta^2 5$$
$$4 = \delta 5$$
$$\delta = \frac{4}{5}$$

d) Combining the results of (b) and (c), compute Andy's β and δ ?

? Answer

We have already computed: (1) $\delta = \frac{4}{5}$. We also know: (2) $\beta \delta = \frac{3}{4}$ Accordingly, substituting (1) into (2):

$$\beta \frac{4}{5} = \frac{3}{4}$$
$$\beta = \frac{15}{16}$$

30.3 Today or tonight?

Kate and Jack have utility function u(x) = x and can choose one of \$3 now (t = 0), \$4 this afternoon (at t = 1), or \$7 tonight (t = 2).

a) Kate is an exponential discounter with $\delta = \frac{1}{2}$. What does she choose?

? Answer

We calculate discounted utility of each option and choose the highest.

$$\begin{split} U_0(\$3) &= 3 \\ U_0(\$4) &= \delta \times 4 \\ &= 2 \\ U_0(\$7) &= \delta^2 \times 7 \\ &= \frac{7}{4} \end{split}$$

Kate chooses the \$3 now.

b) Jack is a hyperbolic discounter with $\beta = \frac{1}{2}$ and delta = 1 what does Jack choose?

? Answer

You calculate discounted utility of each option and choose the highest.

$$\begin{split} U_0(\$3) &= 3\\ U_0(\$4) &= \beta\delta \times 4\\ &= \frac{1}{2} \times 1 \times 4\\ &= 2\\ U_0(\$7) &= \beta\delta^2 \times 7\\ &= \frac{1}{2} \times 1^2 \times 7\\ &= 3.5 \end{split}$$

Jack chooses the \$7 tonight.

Although Jack is a hyperbolic discounter, the short-term discount factor is only applied once.

c) This afternoon (t = 1) comes and Jack reconsiders his decision. Should he take \$4 this afternoon (t = 1), or \$7 tonight (t = 2). What does Jack decide?

? Answer

We calculate discounted utility of each option and choose the highest.

$$U_1(\$4) = 4$$
$$U_1(\$7) = \beta \delta \times 7$$
$$= \frac{1}{2} \times 1 \times$$
$$= 3.5$$

7

Jack changes his mind and takes the \$4 immediately.

d) Why did or did not Jack change his mind?

? Answer

At t = 0 the difference in discount between the \$4 at t = 1 and \$7 at t = 2 is δ . Both are discounted by β as neither is immediately available. However, at t = 1 the difference in discount between the \$4 and \$7 is $\beta\delta$. The \$4 is no longer subject to the immediate discount of β , making it relatively more attractive.

e) Show the answers to parts a) to d) using a diagram.

💡 Answer

The below figure shows the each of the payoffs that Kate could receive at t = 0, 1, 2. The lines represent the discounted utility of each option at each time.

At t = 0 we can see that the \$3 immediately gives higher discounter utility than both of the larger, later options.



The below figure shows the each of the payoffs that Jack could receive at t = 0, 1, 2. The lines represent the discounted utility of each option at each time.

At t = 0 we can see that the \$4 paid at t = 1 gives higher discounter utility. We can also see that the lines represented the discounted utility at each time cross, giving the potential for a time inconsistent decision.



In the next figure we move forward to t = 1. Jack's preferred option changes. He now prefers the \$4 immediately. It is no longer discounted by β .



30.4 Small reward now or large reward later?

Alfred and Blake are exponential discounters. They can choose between a small reward now or a larger reward later. Alfred discounts the future heavily (low δ). Blake does not (δ close to one).

a) How might Alfred and Blake's choices be affected by their discount factor?

? Answer

Each will discount the two rewards, with the size of the discount for the larger reward reflecting the size of the delay relative to the delay for the small reward.

Alfred and Blake will prefer to wait if the discounted utility of the larger reward is higher than that of the small reward.

As Alfred has a higher discount rate, Alfred is more likely than Blake to prefer the small reward as Alfred will discount the larger reward by more.

b) Catherine is a quasi-hyperbolic discounter. She can choose between the same two rewards. Before either award is available, she prefers the large reward. However, she changes her mind and chooses when the small reward at the time it becomes available.

Why did Catherine change her mind?

? Answer

If Catherine prefers the large reward, this means the discounted utility of the larger reward is higher than that for the small reward. As both rewards are experienced with delay, the small and large reward are discounted by the short-term discount factor, plus the exponential discount factor proportional to the delay.

When the small reward becomes available, Catherine no longer applies any discount to the small reward. The exponential discount factor applied to the larger reward also decreases for the time that has passed. Despite the size of the exponential discount being applied to each reward decreasing by the same amount, the removal of the short-term discount factor from the small reward must have been of sufficient that it now has the larger expected utility.

30.5 Credit score

A credit score is a score developed by credit agencies and lenders as a measure of how risky a borrower is. The score is derived using data on past behaviour. A person with a higher credit score is considered more likely to repay a loan on time.

Researchers found a correlation between people's discount factors, β and $\delta,$ and their credit score.

What would this correlation be? Explain why you might see this relationship.

? Answer

You would expect to see a positive correlation between the credit score and the discount factors.

 β relates to present-bias, which results in greater weight being given to any costs or benefits today relative to costs and benefits experienced with delay.

 δ relates to impatience, with each successive period of delay being subject to a discount relative to the previous.

Both discount factors could affect the credit score.

To the extent the future is discounted due to either discount rate, this makes borrowing more attractive provided the interest costs are not too high.

If either discount factor were low enough, an agent might be willing to borrow more than they could feasibly pay off in the future (or at extortionate interest rates). The benefits of consumption today might exceed the costs of failure to repay, with those costs sufficiently discounted that them agent is willing to incur them. An agent with low β might be particularly vulnerable to this as an impulsive purchase today might be attractive even though the debt (e.g. credit card) might be payable soon. Failure to pay would hurt their credit score.

Unintended payment problems are more likely to be caused by low β .

If someone with low δ borrowed money with the intention to pay it off in

the future, they would stick to that plan. They are time consistent. Someone with low β may borrow with the intention to pay it off. However, when the day of payment arrives, they may change their mind and prefer not to incur the immediate cost of payment as that exceeds the discounted costs in the future (such as penalty interest and a low credit score). This would also hurt their credit score.

30.6 Time inconsistent preferences

Recall the question in Section 30.3.

We found that at t = 0 Jack planned to wait until tonight (t = 2) for the \$7, but in the afternoon (t = 1) he changed his mind and took the \$4 that was available immediately.

Jack's friend Allan is a sophisticated quasi-hyperbolic discounter with $\beta = \frac{1}{2}$ and $\delta = 1$.

Does your answer change for Allan? Why?

Answer

The behaviour of Jack that we observed in Tutorial 5 was that of a naive hyperbolic discounter. In each period he calculated his preferred option and acted as though he would stick to that decision in the future. He did not anticipate changing his mind at t = 1.

Allan considers the choice by using backward induction from the final period.

If Allan waits until tonight, he takes the \$7 with certainty. Now considering his choice at t = 1.

$$U_1(\$4 \text{ at } t = 1) = 4$$

 $U_1(\$7 \text{ at } t = 2) = \beta \delta \times 7$
= 3.5

At t = 1 Allan can see that he will cave in take the \$4. Allan now considers t = 0 with the knowledge he will take the \$4 at t = 1. He knows that he will not make it to the \$7 at t = 2, so he removes it from his choice set.

$$U_0(\$3 \text{ at } t = 0) = 3$$
$$U_0(\$4 \text{ at } t = 1) = \beta\delta \times 4$$
$$= 2$$

At t = 0 Allan chooses the \$3. His sophistication leads him to preproperate. He takes a lower amount earlier than the naive Jack.

30.7 Eating cake

Kelvin and Linda both like chocolate cake. There are two periods in which they can eat cake, t = 1, 2. They receive an immediate benefit of 12 for eating cake at t = 1 and 6 for eating cake at t = 2

At t = 3 they incur the costs of their diet and pay a cost of 8. The total cost depends on how much cake they eat.

Kelvin and Linda have preferences of $\beta = 0.5$ and $\delta = 1$.

To illustrate, if Kelvin or Linda eat at t = 1 they receive a benefit of 12 at t = 1 and a cost of 8 at t = 3. If Kelvin or Linda eat at both at t = 1 and at t = 2, they receive a benefit of 12 at t = 1 and 6 at t = 2 and a cost of 16 at t = 3.

Kelvin is naive, Linda is sophisticated.

a) When will Kelvin eat?

? Answer

We can calculate utility of eating in each round separately as the utilities are additive.

At t = 0:

$$\begin{split} U_0(\text{eat at } t=1) &= \beta\delta\times 12 - \beta\delta^3\times 8 \\ &= 0.5\times 1\times 12 - 0.5\times 1^3\times 8 \\ &= 2 \end{split}$$

$$\begin{split} U_0(\text{eat at }t=2) &= \beta \delta^2 \times 6 - \beta \delta^3 \times 8 \\ &= 0.5 \times 1^2 \times 6 - 0.5 \times 1^3 \times 8 \\ &= -1 \end{split}$$

Kelvin plans to eat at t = 1 but not t = 2. At t = 1:

$$\begin{split} U_1(\text{eat at } t=1) &= 12 - \beta \delta^2 \times 8 \\ &= 12 - 0.5 \times 1^2 \times 8 \\ &= 8 \end{split}$$

$$\begin{split} U_1(\text{eat at } t=2) &= \beta\delta\times 6 - \beta\delta^2\times 8 \\ &= 0.5\times 1\times 6 - 0.5\times 1\times 8 \\ &= -1 \end{split}$$

Kelvin eats at t = 1 but does not plan to eat at t = 2. At t = 2:

$$U_2(\text{eat at } t = 2) = 6 - \beta \delta^2 \times 8$$
$$= 6 - 0.5 \times 1^2 \times 8$$
$$= 2$$

Kelvin eats at t = 2.

Kelvin, being naive, thinks he will stick to his initial plan of eating at t = 1 but not at t = 2, but ends up eating in both periods.

b) When will Linda eat?

? Answer

Linda is sophisticated and solves their problem backward. There is no decision to make at t = 3. She simply incurs the cost of their past behaviour.

At t = 2:

$$U_2(\text{eat at } t = 2) = 6 - \beta \delta^2 \times 8$$
$$= 6 - 0.5 \times 1^2 \times 8$$
$$= 2$$

Linda anticipates that she will eat at t = 2. Linda knows she will eat at t = 2, so she don't need to reconsider her plan for then at t = 1. At t = 1:

$$\begin{split} U_1(\text{eat at }t=1) &= 12 - \beta \delta^2 \times 8 \\ &= 12 - 0.5 \times 1^2 \times 8 \\ &= 8 \end{split}$$

Linda also anticipates that she will eat at t = 1. At t = 0, Linda can now see that she will eat at t = 1 and t = 2 no matter what she decides now. She simply accepts her future course of action.

c) Assume that Kelvin and Linda can pay price p at t = 0 for a binding commitment device that prevents them from eating more than they initially plan. Assuming costs and benefits are measured in dollars, what is the maximum price p that Linda would pay to use the commitment device?

? Answer

From the perspective of period t = 0 all costs and benefits are discounted. At t = 0:

$$\begin{split} U_0(\text{eat at }t=1) &= \beta\delta\times 12 - \beta\delta^3\times 8\\ &= 0.5\times 1\times 12 - 0.5\times 1^3\times 8\\ &= 2\\ U_0(\text{eat at }t=2) &= \beta\delta^2\times 12 - \beta\delta^3\times 8 \end{split}$$

$$= 0.5 \times 1^{2} \times 6 - 0.5 \times 1^{3} \times 8$$
$$= -1$$

From the perspective of t = 0, eating at t = 2 has negative discounted utility. As a result, if a commitment device available, Linda would use it to commit to not eating at t = 2. The commitment device would prevent a loss of 1. Therefore, Linda would pay up to p =\$1.

d) What happens to Kelvin, a naive agent, in the presence of the commitment device?

? Answer

Being naive, Kelvin does not perceive the necessity of using a commitment device as he trusts he will comply with his plans.
30.8 Completing an assignment

Your assignment is due today at t = 0.

You can complete the assignment within a day, but on that day you will incur a utility cost of 10.

For every day you submit late, you lose one mark. You experience a utility cost of 1 for every mark lost **on the day it is lost**. If handed in more than 6 days late, you will fail and experience utility cost of 1000. (In other words, if you haven't yet submitted, you will definitely submit at t = 6).

This leaves you with a decision to hand in at t = 0, t = 1, t = 2, ..., t = 5, or t = 6.

For instance:

- If you submit today, at t = 0, you experience utility cost of 10.
- If you submit tomorrow, at t = 1, you will experience utility cost of 10 plus a utility cost of 1 for the mark lost on that day.
- If you submit at t = 2, you will experience a utility cost of 1 at t = 1 for the mark lost, a utility cost of 1 at t = 2 for another mark lost, and a utility cost of 10 for completing the assignment.
- If you submit at t = 6, you will experience a utility cost of 1 on each day from t = 1 to t = 6 for the marks lost, plus a utility cost of 10 at t = 6 for completing the assignment.

You are a hyperbolic discounter with $\beta = 0.75$ and $\delta = 1$.

a) When do you finish if you are naive?

? Answer

At t = 0 you have 7 possible plans to consider: finishing your assignment at t = 0 through to t = 6. You compare the discounted utilities from the various plans as follows.

- At t = 0:
 - Finish today (t = 0): -10
 - Finish tomorrow (t = 1): 0.75(-10 1) = -8.50
 - Finish at t = 2: 0.75(-10 2) = -9.25
 - Finish at t = 3: 0.75(-10 3) = -10
 - •
 - Finish at t = 6: 0.75(-10 6) = -12

At t = 0, you plan to finish at t = 1 as that yields the highest discounted utility. You put off finishing the assignment until tomorrow.

At t = 1 you then reconsider your decision. The penalty received at t = 1 is now sunk and won't affect the decision:

• Finish today (t = 1): -10 = -10

- Finish tomorrow (t = 2): 0.75(-10 1) = -8.25
- Finish at t = 3: 0.75(-10 2) = -9
 - ...
- Finish at t = 6: 0.75(-10 5) = -11.25

At t = 1, you change your plan and now intend to finish at t = 2, which yields the highest discounted utility.

This pattern continues day after day, always intending to complete tomorrow, until t = 6, with an ultimate outcome of -10 utility cost on that day, plus -1 utility cost each day for the last 6 days.

- Finish at t = 5: -10
- Finish at t = 6: 0.75(-10 1) = -8.25

b) When do you finish if you are sophisticated?

? Answer

If you are sophisticated, you will start at the end and work backwards. At t = 6, you know that if you haven't finished the assignment you must, with -10 utility.

At t = 5, you choose between:

- Finish at t = 5: -10
- Finish at t = 6: 0.75(-10 1) = -8.25

You do not plan to do the assignment at t = 5 and you remove t = 5 from your plans.

At t = 4, you choose between:

- Finish at t = 4: -10
- Finish at t = 6: 0.75(-10 2) = -9

You do not plan to do the assignment at t = 4 and you remove t = 4 from your plans.

At t = 3, you choose between:

• Finish at t = 3: -10

• Finish at t = 6: 0.75(-10 - 3) = -9.75

You do not plan to do the assignment at t = 3 and you remove t = 3 from your plans.

At t = 2, you choose between:

- Finish at t = 2: -10
- Finish at t = 6: 0.75(-10 4) = -10.5

Utility is higher completing the assignment earlier. You plan to do the assignment at t = 2 and you remove t = 6 from your plans.

At t = 1, you choose between:

• Finish at t = 1: -10

• Finish at t = 2: 0.75(-10 - 1) = -8.25

You plan to do the assignment at t = 2 and you remove t = 1 from your plans.

At t = 0, you choose between:

- Finish at t = 0: -10
- Finish at t = 2: 0.75(-10 2) = -9

You will not change your intentions further and will complete the assignment at t = 2.

30.9 Saving for retirement

Citizens of Perthia have the choice between the following options:

- Saving for retirement at $t=1\ (u_1=0)$ and having a comfortable retirement at $t=2\ (u_2=20)$
- Spending at t = 1 ($u_1 = 10$) and having a difficult retirement at t = 2 ($u_2 = 0$).

a) Ellie is an exponential discounter with $\delta = 3/4$. What does Ellie choose at t = 0 and t = 1? Why?

? Answer

$$\begin{split} U_0(save) &= \delta u_1 + \delta^2 u_2 \\ &= 0 + \left(\frac{3}{4}\right)^2 \times 20 \\ &= 11.25 \\ U_0(spend) &= \delta u_1 + \delta^2 u_2 \\ &= \frac{3}{4} \times 10 + 0 \\ &= 7.5 \\ U_1(save) &= u_1 + \delta u_2 \\ &= 0 + \frac{3}{4} \times 20 \\ &= 15 \\ U_1(spend) &= u_1 + \delta u_2 \\ &= 10 + 0 \\ &= 10 \end{split}$$

Ellie intends to save in both periods and does so. Ellie is time consistent.

Knowing that Ellie is an exponential discounter, we did not need to calculate her decision at both t = 0 and t = 1. As exponential discounters are time consistent, we could have simply determined her decision for one period and known that would also be her decision at other times.

b) Freddie is a naive quasi-hyperbolic discounter with $\beta = 1/4$ and $\delta = 1$. What does Freddie choose at t = 0 and t = 1? Why?

? Answer

$$\begin{split} U_0(save) &= \beta \delta u_1 + \beta \delta^2 u_2 \\ &= 0 + \frac{1}{4} \times (1)^2 \times 20 \\ &= 5 \end{split}$$

$$U_0(spend) = \beta \delta u_1 + \beta \delta^2 u_2$$
$$= \frac{1}{4} \times 1 \times 10 + 0$$
$$= 2.5$$

At t = 0 Freddie plans to save.

$$U_1(save) = u_1 + \beta \delta u_2$$

= 0 + $\frac{1}{4} \times 1 \times 20$
= 5

$$U_1(spend) = u_1 + \beta \delta u_2$$

= 10 + 0
= 10

At t = 1 Freddie chooses to spend Freddie has changed his mind. He is time inconsistent. At t = 0 both saving and spending are subject to the short-term discount factor β , with the relative discount between saving and spending being only δ . However, at t = 1 the benefit of spending is available immediately and no longer discounted by β .

c) Grant is a sophisticated quasi-hyperbolic discounter with $\beta = 1/4$ and $\delta = 1$. What does Grant do? Why?

? Answer

Grant works through his options using backward induction. At t = 2 Grant has no decision to make. At t = 1:

$$\begin{split} U_1(save) &= u_1 + \beta \delta u_2 \\ &= 0 + \frac{1}{4} \times 1 \times 20 \\ &= 5 \end{split}$$

$$\begin{split} U_1(spend) &= u_1 + \beta \delta u_2 \\ &= 10 + 0 \\ &= 10 \end{split}$$

At t = 1 Grant will spend. (This is the same calculation we made for Freddie at t = 1.)

As Grant knows he will spend at t = 1, that is the only feasible option available at t = 0. Grant will choose to spend.

Ultimately, Grant takes the same action as Freddie. However, he can see his future decisions and is aware of that coming failure to stick with what would be his preferences at t = 0.

d) Grant is offered the opportunity to bind himself to a course of action at t = 0 at the cost of 1 point of utility at t = 2. What does Grant do?

? Answer

In the presence of the commitment device, Grant now has two feasible options to consider at t = 0.

$$U_0(commit) = \beta \delta u_1 + \beta \delta^2 u_2$$

= 0 + $\frac{1}{4} \times (1)^2 \times (20 - 1)$
= 4.75
$$U_0(spend) = \beta \delta u_1 + \beta \delta^2 u_2$$

$$= \frac{1}{4} \times 1 \times 10 + 0$$
$$= 2.5$$

Grant now chooses to commit at a price of one unit of utility at t = 2.

30.10 Buy-now pay-later

Buy-now pay-later works as follows: a person purchases an item with an initial payment of one-quarter of the purchase price. They get access to the purchased item immediately. They then pay three equal instalments each fortnight until they have paid for the purchase in full. If they fail to make a payment on time, they are required to pay a fee of \$10 and are barred from using the buy-now pay-later facility in the future.

Vernon used a buy-now pay-later provider to purchase a new jacket for \$200. He paid \$50 on the day of the purchase and is now required to pay the next \$50 instalment in two weeks. That is, Vernon's schedule of costs and benefits is:

- Purchase date: Gains jacket and pays \$50
- In two weeks: Pays \$50
- In four weeks: Pays \$50
- In six weeks: Pays \$50

At that time of the purchase Vernon intends to pay for the jacket as required by the buy-now pay-later provider in two, four and six weeks.

Two weeks after the purchase when his payment became due Vernon changed his mind and did not make the payment. He purchased a carton of beer for a party that night with the money instead. Vernon's options and the cost and benefits of those options had not changed since the purchase date.

Is Vernon an exponential discounter or present-biased? Why? Explain why Vernon changed his mind.

💡 Answer

As Vernon has exhibited time inconsistent behaviour, he must be presentbiased. An exponential discounter would be time consistent.

A present-based person subjects costs and benefits experienced with any delay to a short-term discount factor.

When Vernon purchased the jacket, the benefit of the jacket and the cost of \$50 would not have been subject to any discount. The cost of the three future payments (or the benefit of alternative uses of that money such as purchasing beer) would be subject to the short-term discount. Any fee for failing to make a payment and the loss of the buy-now pay-later facility for that failure would also be subject to that short-term discount factor. When the next payment is due, the cost of that payment (or the benefit of

the beer) is no longer subject to the short-term discount factor, whereas any fee from failing to make the payment and the loss of the facility are still subject to that short-term discount factor. As a result, cost of the payment / benefit of the beer now has relatively greater weight than the future costs, leading Vernon to change his mind.

$\mathbf{Part}~\mathbf{V}$

Probability foundations

Imagine you're a surgeon deciding whether to operate on a patient. The procedure could save their life, but it also carries risks. Or picture yourself as an investor choosing between stocks for your retirement savings. Each investment offers different potential returns, but also different chances of loss. Or envision yourself at a poker table, trying to judge whether your friend is bluffing with a weak hand or sitting on a royal flush. Calling their hand could have markedly different outcomes.

These decisions share a common thread: we must make choices without perfect information. Sometimes we know the probabilities, like the odds of rolling a six on a fair die. We call this risk. Other times we can only make educated guesses without knowing the probabilities, like whether an investment will pay off. We call this uncertainty. And to analyse decision-making under risk and uncertainty, we need to consider how people form beliefs and compute probabilities.

Consider that poker game. Your friend just raised the stakes significantly. Are they bluffing? You might consider how often they've bluffed before (your prior belief), combine it with how they're acting now (new evidence), and update your assessment of their hand. This process of updating beliefs based on new evidence - which we'll explore as Bayes' Rule - lies at the heart of probability theory.

In this part, we'll proceed in three steps. First, I will lay out the foundations of probability theory - the basic rules and principles that underpin all probabilistic thinking. Then I will discuss Bayes' Rule, a powerful tool that shows us how to update our beliefs when we get new information. Finally, I will extend these ideas into subjective expected utility theory, which helps us understand how people make decisions when probabilities aren't known.

Chapter 31

Probability theory

Summary

•

- The probability function $P(\cdot)$ assigns a real number between 0 and 1 to each event, with $\sum P(\cdot) = 1$ for all possible outcomes.
- For mutually exclusive events A and B, the probability of either event occurring is:

$$P(A \cup B) = P(A) + P(B)$$

• For non-mutually exclusive events A and B, the probability of either event occurring is:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

• For independent events, the probability of both events occurring is the product of their individual probabilities:

$$P(A \cap B) = P(A)P(B)$$

• Conditional probability P(A|B) is the probability of A given that B has occurred. It's calculated as:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

31.1 The Monty Hall problem illustrates conditional probability: given the host opens a specific door, what is the probability of the car being behind the unopened door?

31.2 Introduction

In this part, I introduce some basic concepts in probability theory.

31.3 The probability function

The probability of an outcome is the chance with which it occurs. We denote the probability of outcome A as P(A), where $P(\cdot)$ represents a probability function that assigns a real number to each event.

The probability function has the following features.

First, the probability of outcome A lies between 0 and 1. That is:

$$0 \le P(A) \le 1$$

For example, the probability of drawing the Ace of Spades from a full deck of 52 cards is 1 in 52 or ~ 0.02 .

The probability of flipping a head with a fair coin is 1 in 2 or 0.5.

Second, the probability of the entire outcome space equals 1.

For example, suppose we have 52 possible cards we can draw from the deck, each with 1 in 52 probability. If we draw a single card, the probability that we draw one of those cards is:

$$\frac{1}{52} + \frac{1}{52} + \frac{1}{52} + \dots + \frac{1}{52} = \sum_{n=1}^{n=52} \frac{1}{52}$$
$$= 1$$

Third, suppose outcomes A and B are mutually exclusive. In that case, the probability of A or B is the sum of the probability of A and the probability of B. That is:

$$P(A \text{ or } B) = P(A \cup B)$$

= $P(A) + P(B)$

For example, if we have a deck with 52 cards, the probability of pulling out an Ace with a single draw is as follows.

$$\begin{split} P(A \bigstar \cup A \heartsuit \cup A \diamondsuit \cup A \bigstar) &= P(A \bigstar) + P(A \heartsuit) + P(A \diamondsuit) + P(A \diamondsuit) \\ &= \frac{1}{52} + \frac{1}{52} + \frac{1}{52} + \frac{1}{52} \\ &= \frac{4}{52} \end{split}$$

Alternatively, suppose outcomes A and B are not mutually exclusive. In that case, the probability of one or the other is the sum of the probability of A and the probability of B minus the probability of both occurring. That is:

$$P(A\cup B)=P(A)+P(B)-P(A\cap B)$$

where $P(A \cap B)$ is the probability of both outcome A and B.

For example, if we have a deck with 52 cards, the probability of pulling out an Ace or a Diamond with a single draw is as follows.

$$\begin{split} P(A\cup \diamondsuit) &= P(A) + P(\diamondsuit) - P(A\cap \diamondsuit) \\ &= \frac{4}{52} + \frac{1}{4} - \frac{1}{52} \\ &= \frac{16}{52} \end{split}$$

Finally, if outcomes A and B are independent, the conjunction of the two independent outcomes is the product of their probabilities. That is

$$P(A \cap B) = P(A)P(B)$$

For example, suppose we draw a single card from a deck of cards, place that card back in the deck, and then make another draw. The probability of drawing the Ace of Spades in either draw is 1/52. The probability of drawing the Ace of Spades twice is:

$$\begin{aligned} P(A \spadesuit \cap A \spadesuit) &= P(A \spadesuit) \times P(A \spadesuit) \\ &= \frac{1}{52} \times \frac{1}{52} \\ &= \frac{1}{2704} \end{aligned}$$

Note that if A and B are mutually exclusive, they are not independent and $P(A \cap B) = 0$.

31.4 Conditional probability

Conditional probability concerns the probability of an outcome given another outcome.

For example, drawing a card from a deck of cards with replacement - that is, putting back each card after it is drawn - means that whatever card was drawn in the first draw does not affect the probability of the outcome of the second draw. Each draw is independent of the other.

But what if you draw two cards from the same deck without replacement?

In that case, the two draws are not independent of each other. For instance, if you pull out the Ace of Spades first, the second card cannot be the Ace of Spades.

We say here that the probability of drawing an Ace of Spades on the second draw is conditional on the result of the first draw.

When one outcome is conditional on another, such as the probability of outcome A conditional on outcome B occurring, we write this conditional probability as P(A|B).

Suppose I draw two cards from a deck without replacement. What is the probability of drawing an Ace for both draws?

We know that the first draw affects the probability of drawing an Ace on the second draw. If the first card is an Ace, one less Ace is in the deck for the second draw.

The probability of drawing an Ace on the first draw is 4 in 52. If I draw an Ace in the first draw, the probability of drawing an Ace on the second is 3 in 51. There is one less Ace and one less card than for the first draw. By multiplying the probability of these two events together, we can get the probability of an Ace on both draws.

 $P(\text{Ace } 1\text{st} \cap \text{Ace } 2\text{nd}) = P(\text{Ace } 1\text{st}) \times P(\text{Ace } 2\text{nd}|\text{Ace } 1\text{st})$

$$= \frac{4}{52} \times \frac{3}{51}$$
$$= \frac{1}{221}$$

31.4.1 Formula for conditional probability

We can see that the solution to this problem has taken the form:

$$P(A \cap B) = P(A|B)P(B)$$

The joint probability of two outcomes equals the probability of A conditional on B multiplied by the probability of B.

We can rearrange this formula to determine the probability of A given outcome B.

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

If A and B are independent, P(A|B) = P(A). In that case, the formula simplifies to that for calculating the probability of the conjunction of independent outcomes we saw earlier, $P(A \cap B) = P(A)P(B)$. The equation $P(A \cap B) = P(A|B)P(B)$ is a more general version of how to calculate the conjunction of two events.

Due to symmetry, we can also write the conditional probability as:

$$P(A \cap B) = P(A|B)P(B) = P(B|A)P(A)$$

31.4.2 Example

As a test of this formula, let's take our previous example of drawing two Aces from the same deck. What is the probability of drawing an Ace on the second draw if you drew an Ace on the first?

$$P(\text{Ace 2nd}|\text{Ace 1st}) = \frac{P(\text{Ace 1st} \cap \text{Ace 2nd})}{P(\text{Ace 1st})}$$
$$= \frac{\frac{1}{221}}{\frac{4}{52}}$$
$$= \frac{3}{51}$$

31.4.3 The Monty Hall problem

Consider the following problem as answered by Marilyn vos Savant in her column Ask Marilyn in Parade magazine (Savant, 1990):

Suppose you're on a game show and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice?

This problem is known as the Monty Hall problem as it is loosely based on the American game show Let's Make a Deal. Monty Hall was the original host of the show.

Assume that the rules of this game show are that:

- The host must always open a door that you did not choose.
- The host must always open a door to reveal a goat and never the car.
- The host must always offer you the choice to switch between the chosen door and the remaining closed door.

For this question, you are effectively being asked: what is the probability that the car is behind Door 2 conditional on the host opening door 3.

To help us think about this problem, consider the following tree that maps the possible outcomes after you select Door 1. The first split of the tree represents the 1/3 probability that the car is behind each of the three doors. Given the car's location, the next split represents the probability that the host opens each door. The final column indicates the probability of each combination of car location and door opened.



If the car is behind Door 1, which you have selected, the host could open either Door 2 or Door 3 with equal probability. If the car is behind Door 2, the host must open Door 3. If the car is behind Door 3, the host must open Door 2.

Given the host opened Door 3, we can calculate the conditional probability that the car is behind door 2 as follows:

$$P(C2|D3) = \frac{P(C2 \cap D3)}{P(D3)}$$
$$= \frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{6}}$$
$$= \frac{2}{3}$$

You should switch to door 2.

Chapter 32

Bayes' rule

Summary

- Bayes' rule is a method for calculating the conditional probability of an event.
- Bayes' rule gives the probability of outcome A given B, P(A|B), using the unconditional probability of outcome A, the probability of observing B given A and the total probability of outcome B:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

• The total probability P(B) can be calculated as:

$$P(B) = P(B|A)P(A) + P(B|\neg A)P(\neg A)$$

• Bayes' rule can be thought of as a way to update beliefs about a hypothesis *H* given evidence *E*:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

32.1 Introduction

Bayes' rule is a method for estimating the conditional probability of an event.

Specifically, Bayes' rule allows us to use the following information to estimate the conditional probability of outcome A given outcome B:

- The unconditional probability of outcome A
- The probability of observing outcome B given outcome A
- The total probability of outcome B.

The formula for Bayes' rule is:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$
$$= \frac{P(B|A)P(A)}{P(B)}$$

The denominator P(B) is the total probability of event B. If the total probability of event B is not directly available, we can often calculate it with information concerning the conditional probabilities of B given the occurrence (or not) of A.

$$P(B) = P(B|A)P(A) + P(B|\neg A)P(\neg A)$$

The symbol \neg represents "not".

We can therefore write Bayes' rule as follows:

$$\begin{split} P(A|B) &= \frac{P(B|A)P(A)}{P(B)} \\ &= \frac{P(B|A)P(A)}{P(B|A)P(A) + P(B|\neg A)P(\neg A)} \end{split}$$

32.2 Updating beliefs

We can think of Bayes' rule as how we should update our beliefs in light of a new event.

Rational agents should update their beliefs using Bayes' rule.

In this case, the following elements are involved:

- A hypothesis, H. For example, "the coin is fair" or "the coin is rigged".
- The prior probability of the hypothesis H being true, P(H). For example, "the coin is fair" has a prior probability of 0.5.
- The probability of observing event E given a hypothesis H, P(E|H). For example, "the coin shows a head" has a probability of 0.5 given that the coin is fair.

• The posterior probability of the belief H given the event E, P(H|E). For example, we would have an updated probability in our hypothesis that "the coin is fair" based on the coin showing a head.

Under this framing, Bayes' rule is formulated as follows:

$$\underbrace{\frac{P(H|E)}{\text{Posterior belief}}}_{\text{Posterior belief}} = \frac{\frac{P(E|H) \quad \widetilde{P(H)}}{P(E)}}{P(E)}$$
$$= \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|\neg H)P(\neg H)}$$

32.3 A rigged coin

Suppose your friend has two coins. One is a fair coin with a head on one side and a tail on the other. The second coin is a rigged coin with a head on both sides.

Your friend takes one of the coins and flips it. The coin shows a head. What is the probability that this coin is the rigged coin?

We will assume that he randomly selected either coin with a probability of 50%. We take that as our prior belief:

P(rigged) = 0.5

The probability of a head if it is the rigged coin is 1.

P(head|rigged) = 1

To use Bayes' rule, we need the total probability that a head comes up, P(head). Here we will use the formula for total probability.

$$P(\text{head}) = P(\text{head}|\text{rigged})P(\text{rigged}) + P(\text{head}|\text{fair})P(\text{fair})$$
$$= 1 \times 0.5 + 0.5 \times 0.5$$
$$= 0.75$$

Putting this into Bayes' rule:

$$P(\text{rigged}|\text{head}) = \frac{P(\text{head}|\text{rigged})P(\text{rigged})}{P(\text{head})}$$
$$= \frac{1 \times 0.5}{0.75}$$
$$= \frac{2}{3}$$

Your friend flips the coin again and gets another head. What is the updated probability that the coin is rigged?

The prior belief is now $P(\text{rigged}) = \frac{2}{3}$. The total probability of flipping a head is:

$$\begin{split} P(\text{head}) &= P(\text{head}|\text{rigged})P(\text{rigged}) + P(\text{head}|\text{fair})P(\text{fair}) \\ &= 1 \times \frac{2}{3} + 0.5 \times \frac{1}{3} \\ &= \frac{5}{6} \end{split}$$

Putting this into Bayes rule:

$$P(\text{rigged}|\text{head}) = \frac{P(\text{head}|\text{rigged})P(\text{rigged})}{P(\text{head})}$$
$$= \frac{1 \times \frac{2}{3}}{\frac{5}{6}}$$
$$= \frac{4}{5}$$

Your belief that the coin is rigged has now increased to 80%.

Your friend flips the coin 10 more times and gets 10 more heads. What is the updated probability that the coin is rigged?

We use our prior belief of $P(\text{rigged}) = \frac{4}{5}$.

The total probability of flipping 10 heads is:

P(10 heads) = P(10 heads|rigged)P(rigged) + P(10 heads|fair)P(fair)

$$= 1 \times \frac{4}{5} + \left(\frac{1}{2}\right)^{10} \times \frac{1}{5}$$

= 0.8002

Putting this into Bayes' rule:

$$P(\text{rigged}|10 \text{ heads}) = \frac{P(10 \text{ heads}|\text{rigged})P(\text{rigged})}{P(10 \text{ heads})}$$
$$= \frac{1 \times \frac{4}{5}}{0.8002}$$
$$= 0.99976$$

We now believe the coin is rigged with greater than 99.9% probability.

32.4 Balls from an urn

You have two urns filled with balls. Urn 1 has 30% black balls and 70% yellow balls. Urn 2 has 70% black balls and 30% yellow balls. The labels have fallen off the urns, so you do not know which urn is which.

You reach into one of the urns and pull out a yellow ball. What is the probability that you have drawn the ball from urn 1?

The Bayes' rule formula to solve this problem is:

$$P(\text{urn 1}|\text{yellow}) = \frac{P(\text{yellow}|\text{urn 1})P(\text{urn 1})}{P(\text{yellow})}$$

We take the prior probability of the ball coming from urn 1 to be 50%. The probability of drawing a yellow ball from urn 1 is 70%.

The total probability of drawing a yellow ball is:

$$\begin{split} P(\text{yellow}) &= P(\text{yellow}|\text{urn 1})P(\text{urn 1}) + P(\text{yellow}|\text{urn 2})P(\text{urn 2}) \\ &= 0.3 \times 0.5 + 0.7 \times 0.5 \\ &= 0.5 \end{split}$$

Putting this into Bayes' rule:

$$P(\text{urn 1}|\text{yellow}) = \frac{P(\text{yellow}|\text{urn 1})P(\text{urn 1})}{P(\text{yellow})}$$
$$= \frac{P(\text{yellow}|\text{urn 1})P(\text{urn 1})}{P(\text{yellow}|\text{urn 1})P(\text{urn 1}) + P(\text{yellow}|\text{urn 2})P(\text{urn 2})}$$
$$= \frac{0.7 \times 0.5}{0.7 \times 0.5 + 0.3 \times 0.5}$$
$$= 0.7$$

You put the first ball back in the urn, reach in again and pull out a black ball. What is the probability that you have drawn the ball from urn 1?

Given we have already drawn one ball and updated our probability, we will use the prior probability of P(urn 1) = 0.7.

The total probability of drawing a black ball is:

$$\begin{split} P(\mathrm{black}) &= P(\mathrm{black}|\mathrm{urn}\ 1)P(\mathrm{urn}\ 1) + P(\mathrm{black}|\mathrm{urn}\ 2)P(\mathrm{urn}\ 2) \\ &= 0.3\times 0.7 + 0.7\times 0.3 \\ &= 0.42 \end{split}$$

Putting this into Bayes rule:

$$P(\text{urn 1}|\text{black}) = \frac{P(\text{black}|\text{urn 1})P(\text{urn 1})}{P(\text{black})}$$
$$= \frac{0.3 \times 0.7}{0.42}$$
$$= 0.5$$

The answer of 0.5 should seem intuitive. We have now drawn one black and one yellow ball. In combination, this is uninformative and we are back at our initial prior of 0.5.

32.5 The Monty Hall problem

Recall the Monty Hall problem:

Suppose you're on a game show and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice?

Assume that the rules of this game show are that:

- The host must always open a door that you did not choose.
- The host must always open a door to reveal a goat and never the car.
- The host must always offer you the choice to switch between the chosen door and the remaining closed door.

We want to know the probability that the car is behind Door 2 given the host opened Door 3. We want to know P(C2|D3).¹

To determine this using Bayes' rule, we would use the following formula:

$$P(C2|D3) = \frac{P(D3|C2)P(C2)}{P(D3)}$$

P(D3) is the probability that the host opens Door 3. It is calculated using the formula for total probability:

$$P(D3) = P(D3|C1)P(C1) + P(D3|C2)P(C2) + P(D3|C3)P(C3)$$

Each of those elements are as follows.

P(C1), P(C2) and P(C3) are our prior probability of the car being behind each door, which is $\frac{1}{3}$.

P(D3|C1) is the probability that the host opens door 3, given the car is behind door 1. The host could open either of Door 2 or Door 3 as neither has the car behind it, so the probability of Door 3 is $\frac{1}{2}$.

P(D3|C2) is the probability that the host opens Door 3, given the car is behind Door 2. The host must open that door, so the probability is one. They cannot open the door you have chosen or the door that the car is behind.

¹Technically, we want $P(C2|D3 \cap X1)$ where X1 is our selection of Door 1. However, adding this complication to the calculation does not change the answer.

P(D3|C3) is the probability that the host opens door 3, given the car is behind door 3. The host cannot open a door to show the car, so the probability is zero. Returning to our equations, the total probability of the host opening Door 3 is:

$$P(D3) = P(D3|C1)P(C1) + P(D3|C2)P(C2) + P(D3|C3)P(C3)$$
$$= \frac{1}{2} \times \frac{1}{3} + 1 \times \frac{1}{3} + 0 \times \frac{1}{3}$$
$$= \frac{1}{2}$$

Now we can calculate the probability that the car is behind Door 2, given the host opened Door 3:

$$P(C2|D3) = \frac{P(D3|C2)P(C2)}{P(D3)}$$
$$= \frac{1 \times \frac{1}{3}}{\frac{1}{2}}$$
$$= \frac{2}{3}$$

The contestant should switch doors.

Chapter 33

Probability theory exercises

33.1 Judging a fund manager

You want to know if a fund manager is skilled as you believe skilled management can lead to outperformance. The following is known:

- 20% of fund managers are skilled.
- If a fund manager is skilled, the probability that they outperform the market is 80%.
- If a fund manager is unskilled, the probability that they outperform the market is 40%.

You observe an outperforming fund manager. What is the probability that the fund manager is skilled?

? Answer

We calculate the solution using Bayes' rule.

$$P(\text{skilled}|\text{outperform}) = \frac{P(\text{outperform}|\text{skilled})P(\text{skilled})}{P(\text{outperform})}$$

We can calculate P(outperform) using the law of total probability.

P(outperform) = P(outperform|skilled)P(skilled) + P(outperform|unskilled)P(unskilled)

 $= 0.8 \times 0.2 + 0.4 \times 0.8$

= 0.48

Inputting this into Bayes' rule, we get:

$$\begin{split} P(\text{skilled}|\text{outperform}) &= \frac{P(\text{outperform}|\text{skilled})P(\text{skilled})}{P(\text{outperform})} \\ &= \frac{0.8 \times 0.2}{0.48} \\ &= 0.333 \end{split}$$

A fund manager who outperforms is skilled with 0.333 probability. You cannot neglect that low base rate of skilled managers.

33.2 Detecting a terrorist

Every month 100 million people fly on commercial airlines. Imagine 10 of them are terrorists.

Airport security are able to correctly identify that a person is a terrorist in 99% of cases and a non-terrorist in 99.9% of cases.

A person is identified by airport security as a terrorist. Using Bayes' rule, what is the probability that they are a terrorist?

Answer
We can use Bayes' Theorem to calculate the conditional probability: $P(\text{terrorist} \mid \text{identified}) = \frac{P(\text{identified} \mid \text{terrorist})P(\text{terrorist})}{P(\text{identified})}$ $= \frac{0.99 \times 0.0000001}{0.99 \times 0.0000001 + 0.001 \times 0.999999}$ = 0.000099Or approximately 1 in 10,000.

33.3 Rolling a die

You have two six-sided dice:

- one die is fair with the numbers 1 through to 6 occurring with equal probability
- the other die is loaded and always rolls a 5 or a 6 with equal probability.

You pull one die out of your pocket and roll it. You did not check which die it was before you rolled. (Assume you could have pulled either die out of your pocket with equal probability.)

a) The die shows a six. What is the probability that it is the loaded die?

Answer We are asked to calculate P(D1 loaded|6). To calculate this, I will use Bayes' rule. $P(D1 \text{ loaded}|6) = \frac{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}{P(6)}$ $= \frac{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}$ $= \frac{0.5 \times 0.5}{0.5 \times 0.5 + \frac{1}{6} \times 0.5}$ = 0.75The first die is the loaded die with 75% probability.

b) You pull the other die out of your pocket and roll it. It shows a 5. What is the updated probability that the first die is the loaded die?

💡 Answer

We are asked to calculate P(D1 loaded|D2 5). We have a prior probability of the first die being loaded of 0.75, which comes from the answer in the previous question.

Putting this into Bayes' rule:

$$P(\text{D1 loaded}|\text{D2 5}) = \frac{P(\text{D2 5}|\text{D1 loaded})P(\text{D1 loaded})}{P(\text{D2 5})}$$

$$= \frac{P(\text{D2 5}|\text{D1 loaded})P(\text{D1 loaded})}{P(\text{D2 5}|\text{D1 loaded})P(\text{D1 loaded}) + P(\text{D2 5}|\text{D1 fair})P(\text{D1 fair})}$$
$$= \frac{\frac{1}{6} \times 0.75}{\frac{1}{6} \times 0.75 + 0.5 \times 0.25}$$
$$= 0.5$$

Each die is loaded with a 50% probability.

You could reach the same result by calculating the probability that the second die is fair given a 5. This is the same as the probability that the first die is loaded given a 5 on the second die. This approach gives the same result as the above calculation, but you need only think about one die.

$$P(\text{D2 fair}|5) = \frac{P(5|\text{D2 fair})P(\text{D2 fair})}{P(5)}$$
$$= \frac{P(5|\text{D2 fair})P(\text{D2 fair})}{P(5|\text{D2 fair})P(\text{D2 fair}) + P(5|\text{D2 loaded})P(\text{D2 loaded})}$$
$$= \frac{\frac{1}{6} \times 0.75}{\frac{1}{6} \times 0.75 + 0.5 \times 0.25}$$
$$= 0.5$$

33.4 Male or female

These two related questions come from magazine columnist Marilyn vos Savant.

a) A shopkeeper says she has two new baby beagles to show you, but she doesn't know whether they're male, female, or a pair. You tell her that you want only a male, and she telephones the fellow who's giving them a bath. "Is at least one a male?" she asks him. "Yes!" she informs you with a smile. What is the probability that the other one is a male?

? Answer

The prior probabilities are:

$$\begin{split} P(MM) &= 0.25 \\ P(MF) &= P(FM) = 0.25 \\ P(FF) &= 0.25 \end{split}$$

Using Bayes' rule:

$$P(MM|M) = \frac{P(M|MM)P(MM)}{P(M)}$$

= $\frac{P(M|MM)P(MM)}{P(M|MM)P(MM) + P(M|MF)P(MF)}$
+ $P(M|FM)P(FM) + P(M|FF)P(FF)$
= $\frac{1 \times \frac{1}{4}}{1 \times \frac{1}{4} + 1 \times \frac{1}{4} + 1 \times \frac{1}{4} + 0 \times \frac{1}{4}}$
= $\frac{\frac{1}{4}}{\frac{3}{4}}$
= $\frac{1}{3}$

b) Say that a woman and a man (who are unrelated) each have two children. We know that at least one of the woman's children is a boy and that the man's oldest child is a boy. Can you explain why the chances that the woman has two boys do not equal the chances that the man has two boys?

? Answer

For the woman, the prior probabilities before learning she has a boy are:

$$\begin{split} P(BB) &= 0.25 \\ P(BG) &= P(GB) = 0.25 \\ P(GG) &= 0.25 \end{split}$$

Using Bayes' rule:

$$P(BB|B) = \frac{P(B|BB)P(BB)}{P(B)}$$

$$= \frac{P(B|BB)P(BB)}{P(B|BB)P(BB) + P(B|BG)P(BG)}$$

$$+ P(B|GB)P(GB) + P(B|GG)P(GG)$$

$$= \frac{1 \times \frac{1}{4}}{1 \times \frac{1}{4} + 1 \times \frac{1}{4} + 1 \times \frac{1}{4} + 0 \times \frac{1}{4}}$$

$$= \frac{\frac{1}{\frac{4}{3}}}{\frac{3}{4}}$$

$$= \frac{1}{3}$$

For the man, the prior probabilities before learning his eldest is a boy are:

$$P(BB) = 0.25$$
$$P(BG) = 0.25$$
$$P(GB) = 0.25$$
$$P(GG) = 0.25$$

Using Bayes' rule:

$$P(BB|B) = \frac{P(B|BB)P(BB)}{P(B)}$$

$$= \frac{P(B|BB)P(BB)}{P(B|BB)P(BB) + P(B|BG)P(BG)}$$

$$+ P(B|GB)P(GB) + P(B|GG)P(GG)$$

$$= \frac{1 \times \frac{1}{4}}{1 \times \frac{1}{4} + 1 \times \frac{1}{4} + 0 \times \frac{1}{4} + 0 \times \frac{1}{4}}$$

$$= \frac{\frac{1}{4}}{\frac{2}{4}}$$

$$= \frac{1}{2}$$

Note: questions such are these are typically sensitive to unstated assumptions, particularly around the procedure used to elicit the information around the sex of the child. Due to this, part a) is probably less vulnerable to alternative assumptions than b) as it contains information about the elicitation procedure in the question.

33.5 Luggage

You have taken a flight and are worried that your luggage will not be on the flight. You know that 20% of bags have not been arriving with the passenger.

The following table of conditional probabilities (from Pearl & Mackenzie (2018)) gives the probability that your bag will be on the luggage carousel conditional on the bag being on the flight and how long you have been waiting at the carousel.

Bag on plane	Time elapsed	Carousel=true
False	0	0
False	1	0
False	2	0
False	3	0
False	4	0
False	5	0
False	6	0
False	7	0
False	8	0

Bag on plane	Time elapsed	Carousel=true
False	9	0
False	10	0
True	0	0
True	1	10
True	2	20
True	3	30
True	4	40
True	5	50
True	6	60
True	7	70
True	8	80
True	9	90
True	10	100

a) You have been waiting 5 minutes for your bag and it has not arrived. What is the probability that your bag was not on the flight?

b) You have been waiting 9 minutes for your bag and it has not arrived. What is the probability that your bag was not on the flight?

Answer

 $P(\text{false}|\text{not arrived after 9}) = \frac{P(\text{not arrived after 9}|\text{false})P(\text{false})}{P(\text{not arrived after 9})}$ $= \frac{P(\text{not arrived after 9}|\text{false})P(\text{false})}{P(\text{not arrived after 9}|\text{false})P(\text{false})}$ + P(not arrived after 9|false)P(false)

$$= \frac{1 \times 0.2}{1 \times 0.2 + 0.1 \times 0.8}$$

= 0.714

The probability that the bag was not on the flight is 71%.

c) You have been waiting 10 minutes for your bag and it has not arrived. What is the probability that your bag was not on the flight?

💡 Answer

 $P(\text{false}|\text{not arrived after 10}) = \frac{P(\text{not arrived after 10}|\text{false})P(\text{false})}{P(\text{not arrived after 10})}$ $= \frac{P(\text{not arrived after 10}|\text{false})P(\text{false})}{P(\text{not arrived after 10}|\text{false})P(\text{false})}$ + P(not arrived after 10|true)P(true) $= \frac{1 \times 0.2}{1 \times 0.2 + 0 \times 0.8}$ = 1

The probability that the bag was not on the flight is 100%.

33.6 Murder?

In a now famous case, Sally Clark was convicted of murder after the death of her two sons. The defence argued both children had died of sudden infant death syndrome. The prosecution called statistical evidence that the chance of a SIDS death was 1 in 8543, so the chance of two children dying of SIDS was 1 in 73 million (8543×8543).

The conviction was overturned on a second appeal. For what reasons could the following simple calculation be misleading?

 $P(\text{SIDS death} \land \text{SIDS death}) = P(\text{SIDS death}) \times P(\text{SIDS death})$

$$= \frac{1}{8543} \times \frac{1}{8543}$$
$$= \frac{1}{7.3 \times 10^7}$$

? Answer

Reason 1: The probability of a 2nd SIDS death given first SIDS death is not independent. e.g. SIDS deaths are related due to genetics, family environment. That is, the relevant probability is:

P(SIDS death|genetics, family environment, etc)

This would mean that:

 $P(2 \text{ SIDS deaths in same family}) > (P(1 \text{ SIDS death}))^2$

The appropriate calculation is:

 $P(\text{SIDS death} \land \text{SIDS death}) = P(1^{st} \text{ SIDS death}) \times P(2^{nd} \text{ SIDS death}|1^{st} \text{ SIDS death})$

Reason 2: We also need to consider the probability of alternative possibility - i.e. murder. We want calculate:

$$P(\text{murder}|2 \text{ deaths}) = \frac{P(2 \text{ deaths}|\text{murder})P(\text{murder})}{P(2 \text{ deaths})}$$

If murder itself is also unlikely, then it is not correct to simply attribute the alternative to SIDS the residual probability.

33.7 A fire alarm

You know the following statistics about fire:

- The probability of your house catching fire on any particular day is 1 in 10,000

- Your fire a larm correctly detects a house fire 95% of the time
- The probability that your fire alarm sounds on a day when there is no fire (a false alarm) is 1 in 100.

Your alarm goes off. What is the probability that your house is on fire?

Answer

$$P(\text{fire}|\text{alarm}) = \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm})}$$

$$= \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm}|\text{fire})P(\text{fire}) + P(\text{alarm}|\neg\text{fire})P(\neg\text{fire})}$$

$$= \frac{0.95 \times 0.0001}{0.95 \times 0.0001 + 0.01 \times 0.9999}$$

$$= 0.0094$$
The probability of a fire if the alarm goes off is 0.94%.

Chapter 34

Subjective expected utility theory

Summary

- Subjective expected utility theory applies to situations of uncertainty where probabilities are unknown.
- Agents maximise subjective expected utility using subjective probabilities. Subjective expected utility is defined as:

$$\mathbb{E}[U(x)] = \sum_{i=1}^n \pi(x_i) u(x_i)$$

- We assume that decision-makers maintain coherent subjective probabilities consistent with Bayesian probability theory.
- Probabilistic coherence protects decision-makers from exploitation through Dutch Books.
- An anomaly in subjective expected utility theory is the Ellsberg paradox, where people's preference for known over unknown probabilities violates the theory's predictions.

34.1 Introduction

When deciding whether to invest in an untested technology or choose between conflicting medical treatments, we rarely have access to exact well-defined probabilities. Subjective expected utility theory provides us with a framework for examining decisions in such circumstances, where probabilities are unknown.
In standard expected utility theory, probabilities are derived from known frequencies or distributions. In contrast, subjective expected utility theory uses subjective probabilities based on personal judgment.

This subjective approach maintains the mathematical framework of expected utility theory. Agents maximise subjective expected utility, $\mathbb{E}[U(X)]$, by assigning subjective probabilities, $\pi(x_i)$, to outcomes x_i . The equation for subjective expected utility theory, except for these subjective probabilities. Subjective expected utility equals the subjective probabilities of each outcome x_i :

$$\begin{split} \mathbb{E}[U(X)] &= \pi(x_1) u(x_1) + \pi(x_2) u(x_2) + \ldots + \pi(x_n) u(x_n) \\ &= \sum_{i=1}^n \pi(x_i) u(x_i) \end{split}$$

This process involves:

- 1. Defining utility $u(x_i)$ over final outcomes x_1, \ldots, x_n .
- 2. Defining subjective probabilities $\pi(x_i)$ for each outcome.
- 3. Weighting each outcome's utility by its subjective probability.
- 4. Summing the weighted utilities.

34.2 Axioms for subjective expected utility

The axioms for subjective expected utility theory were first set out by Leonard Savage in 1954. These axioms have a close alignment with the axioms for standard expected utility theory, although they are adapted to the subjective nature of the probabilities.

One of these axioms is the **sure-thing principle**. Informally, we could state it as follows:

Suppose there are two possible states of the world. If you prefer one option (A) over another (B) in one possible state, and you also prefer A over B under the alternative state, then you should prefer A over B even if you do not know which state will occur.

34.3 Coherent probabilities

The Sure-Thing Principle, in combination with other axioms such as continuity, effectively forces subjective probabilities to be coherent.

Each potential event must have a subjective probability of between 0 and 1 (non-negativity). Formally:

$$0 \le \pi(A) \le 1$$

Subjective probabilities must be additive. That is, for mutually exclusive events A and B, the subjective probability of either event occurring is:

$$\pi(A \cup B) = \pi(A) + \pi(B)$$

Further, the sum of the subjective probabilities of all all mutually exclusive, exhaustive events events is one (normalisation):

$$\sum \pi(\cdot) = 1$$

Once subjective probabilities satisfy the usual rules of probability, **Bayes' Rule** becomes the only way to revise these probabilities consistently when new information arrives. If you learn some event B has happened, you must update the probability of A using:

$$\pi(A \mid B) = \frac{\pi(A \cap B)}{\pi(B)},$$

Otherwise, you could violate the Sure-Thing Principle by preferring one option over another in each conditional scenario ("if B" vs. "if not B") yet reversing that preference when you do not know whether B has occurred. Thus, coherent subjective probabilities naturally imply Bayesian updating.

34.3.1 Why coherence matters

Coherence ensures that subjective probabilities avoids paradoxes and do not cause errors in decision-making.

For example, coherent subjective probabilities can allow someone to avoid a **Dutch Book**. A Dutch book is a set of bets that guarantees a loss to one party regardless of the outcome. If probabilities violate coherence requirements, such a set of bets can be constructed against the decision-maker.

For example, consider someone who assigns probabilities that violate additivity:

- Probability of rain tomorrow: 60%
- Probability of no rain tomorrow: 50%

These probabilities sum to 110%, violating normalisation. A bookmaker could:

- Sell a \$60 bet paying \$100 if it rains
- Sell a \$50 bet paying \$100 if it doesn't rain

The decision-maker pays \$110 in total. If it rains, they win \$100 from bet 1 and nothing from bet 2. If it doesn't rain, they win \$100 from bet 2 and nothing from bet 1. The decision-maker pays \$110 in total but can only win \$100, guaranteeing a loss of \$10.

The Dutch Book argument thus demonstrates that incoherent probabilities lead to exploitable inconsistencies. Decision-makers who follow the coherence requirements are protected from such guaranteed losses.

34.4 An anomaly: the Ellsberg Paradox

The Ellsberg paradox demonstrates that many people's choices violate subjective expected utility theory.

Consider two urns, each containing 100 balls. In the "risky" urn, there are 50 red and 50 black balls. In the "ambiguous" urn is a mix of red and black balls of unknown proportions. You are offered the following two gambles:

- The first, which we will call gamble A, will involve a draw from the risky urn. If you correctly predict the colour that is drawn, you will win \$100.
- The second, gamble B, involves a draw from the ambiguous urn. Again, if you correctly predict the colour that is drawn, you will win \$100.



Most people strictly prefer gamble A to gamble B. However, this robust pattern violates subjective expected utility theory and the sure-thing principle.

To understand why, consider the probability of drawing a ball from the risky urn. The probability of drawing a red ball equals the probability of drawing a black ball which equals 0.5. Whatever colour you predict for gamble A, the expected utility of the draw from the risky urn is 0.5 times the utility of \$100.

$$\mathbb{E}[U(A)] = 0.5u(\$100)$$

If you draw from the ambiguous urn, we need to know the subjective probabilities you give for each ball colour. These don't have to be 50%. Maybe the experimenter likes red. Whatever those subjective probabilities, we can calculate subjective expected utility. Let the subjective probability that the ball is red be $\pi(r)$, with the subjective probability of the ball being black being $\pi(b) = 1 - \pi(r)$ (by normalisation).

The subjective expected utility of predicting red for gamble B is the subjective probability of the ball being red times the utility of winning \$100:

$$\mathbb{E}[U(B_r)] = \pi(r) u(\$100)$$

The subjective expected utility of predicting black for gamble B is the subjective probability of the ball being black times the utility of winning \$100:

$$\mathbb{E}[U(B_b)] = (1 - \pi(r))u(\$100)$$

As you get to predict the colour, choosing gamble B will deliver you the higher of the subjective expected utility of red or black. Your subjective expected utility is the higher of the probability of red or black times the utility of winning \$100.

$$\mathbb{E}[U(B)] = \max\{\pi(r)u(\$100), (1 - \pi(r))u(\$100\}$$
$$= \max\{\pi(r), 1 - \pi(r)\}u(\$100)$$

If $\pi(r) = 0.5$, then gambles A and B give the same subjective expected utility. But for any $\pi(r) \neq 0.5$, gamble B should be strictly preferred. This is because when your subjective probability differs from 0.5, you can increase your subjective expected utility by betting on whichever colour you believe is more likely in the ambiguous urn. Therefore, no matter what your subjective belief, gamble B should be at least weakly preferred to gamble A.

To put this into the language of the sure-thing principle, there are 101 possible states of the world (no red balls in the ambiguous urn, one red ball in the ambiguous urn, etc.). For 100 of those states - whenever there are not 50 of each colour in the urn - you should strictly prefer gamble B. In the case of equal numbers of each ball, you should weakly prefer gamble B (indifference). Therefore, given you should prefer gamble B to gamble A whatever the state, you should prefer gamble B not knowing precisely how many balls of each colour are in the ambiguous urn.

This paradox suggests people distinguish between risk (known probabilities) and ambiguity (unknown probabilities) in ways that subjective expected utility theory cannot capture. This result - known as the Ellsberg paradox - illustrates "ambiguity aversion". People tend to prefer gambles with known probabilities over gambles with ambiguous probabilities.

Part VI

Heuristics and biases

Imagine you are trying to catch a ball.

One strategy would be to calculate the ball's trajectory - that is, it's speed, angle, and spin - make some adjustments for wind resistance. You then move to the spot where the ball will land.

Another would be to use a simple rule: keep your eye on the ball and maintain a constant angle of gaze as you move. This mental shortcut - known as the gaze heuristic - works remarkably well. Yet this same strategy can lead you to run in a curved path rather than straight to where the ball will land. It might lead to first run away and then run toward the landing point.



This approach to catching a ball is an example of a heuristic - a mental shortcut. Heuristics are essential tools that help us navigate a complex world with constrained cognitive bandwidth. But heuristics can also lead to systematic biases.

In this part, I will explore how we use heuristics to form judgments and make decisions under uncertainty. I will examine how these mental shortcuts can lead to biases and errors in our reasoning.

First, I'll examine fundamental heuristics that people commonly use when assessing probabilities and making predictions. I'll cover the representativeness heuristic (judging probability by similarity to prototypes), the availability heuristic (estimating frequency based on how easily examples come to mind), and anchoring and adjustment (using initial values as reference points for estimates).

Building on this foundation, I'll then analyse specific biases that emerge in probability judgments. We'll explore the conjunction fallacy (judging combinations of events as more likely than individual events) and base-rate neglect (failing to consider background probabilities). We will also consider two famous biases that we encounter in the world of sports and gambling: the hot-hand fallacy and the gambler's fallacy.

We will then look at how heuristics can actually serve as effective decisionmaking tools rather than just sources of error. Through the lens of the biasvariance trade-off, we'll see how simpler heuristic approaches can sometimes outperform more complex strategies.

Finally, I'll examine overconfidence, an inability to calibrate the accuracy of our judgments. I will look at three dimensions of overconfidence, being overprecision, overestimation, and overplacement. We'll explore how these manifest in real-world decision-making contexts.

This exploration will give us a richer understanding of how people actually make probabilistic judgments, rather than just how they theoretically should. This provides crucial insights for analysing decision-making under uncertainty.

Chapter 35

Heuristics

Summary

- Heuristics are mental shortcuts or rules of thumb people use to make decisions. They often involve limited information and simpler computation than full optimization.
- The representativeness heuristic involves judging the probability of an event or category membership based on how similar it is to a prototype or stereotype, often neglecting relevant base rates.
- The availability heuristic leads people to estimate frequency or probability based on how easily examples come to mind, which can be influenced by factors like recency or vividness.
- Anchoring and adjustment is a process where people start with an initial value (anchor) and adjust insufficiently from it, even when the anchor is arbitrary or irrelevant, leading to biased estimates.

35.1 Introduction

Heuristics are mental shortcuts or rules of thumb that people use to make decisions. They differ from optimisation in that they typically involve a limited information set and a more computationally tractable decision method.

An example of a heuristic is the recognition heuristic. Under this heuristic, if one of two objects is recognised, infer that the recognised object is more likely to be the target object. For example, if you want to predict which of two players will win a tennis match, and you know of only one of the two players, infer that the player you know will win. Similarly, if judging the relative size of two cities, of which you have heard of only one, infer that the city you have heard of is larger.

There is substantial evidence that people use heuristics. People don't normally calculate conditional probabilities using Bayes' rule. Instead, the heuristics might approximate Bayes rule under certain conditions.

Heuristics are often accurate and tractable, but in some environments can lead to error.

Tversky and Kahneman (1974) defined three now classic heuristics: representativeness, availability and anchoring.

I will illustrate these three and then discuss a series of biases in probability judgment for which these heuristics may provide an explanation.

35.2 Representativeness heuristic

Suppose you wish to estimate the probability that an event or person belongs to a certain class.

- "What is the probability that event A belongs to class B?"
- "What is the probability that process B will generate event A?"

Under the representativeness heuristic, people evaluate probabilities by the degree to which A is representative of (similar to) B.

Tversky and Kahneman (1974) provide the following example:

[C]onsider an individual who has been described by a former neighbor as follows:

"Steve is very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail."

How do people assess the probability that Steve is engaged in a particular occupation from a list of possibilities (for example, farmer, salesman, airline pilot, librarian, or physician)? How do people order these occupations from most to least likely? In the representativeness heuristic, the probability that Steve is a Iibrarian, for example, is assessed by the degree to which he is representative of, or similar to, the stereotype of a librarian. Indeed, research with problems of this type has shown that people order the occupations by probability and by similarity in exactly the same way.

35.3 Availability heuristic

Under the availability heuristic, people assess the frequency of a class or the probability of an event by the ease with which they can recall instances or occurrences. If an event is more "available", it is judged to have a higher frequency or probability.

For example, if assessing the probability of a heart attack, you might recall occurrences among people you know. If you are assessing the probability of shark attack, you might recall how often you hear of attacks on the news.

In one experiment, Tversky and Kahneman (1974) gave experimental subjects lists of names. In some lists, the men were more famous than the women, and in other lists, vice versa. After viewing the list, they were asked whether the list had more men or women.

For each list, the subjects judged that the sex with more famous names was more common. Those names were more available in their minds.

35.4 Anchoring and adjustment

When using anchoring and adjustment, people estimate by starting from an initial value and adjust from that value to obtain the final estimate.

Suppose you know the odds of outcome A, and want to estimate the odds of outcome B. Anchoring and adjustment implies that you will start with the odds of outcome A and adjust to obtain the odds of outcome B.

The accuracy of anchoring and adjustment depends on the anchor's quality and the size of the adjustment from the anchor.

The quality of the anchor relates to the strength of the correlation between the anchor and the quantity being estimated. Empirically, people tend to use weak or irrelevant anchors.

The size of the adjustment then needs to account for the relationship between the anchor and the quantity being estimated. Empirically, observed adjustments from the anchor are too small.

As an example, Tversky and Kahneman (1974) asked subjects to estimate the percentage of African countries in the United Nations.

A number between 0 and 100 was determined by spinning a wheel in the subjects' presence. The subjects were instructed to indicate first whether that number was higher or lower than the estimated percentage, and then to estimate the value of the quantity by moving upward or downward from the given number.

Different groups were given different numbers from the wheel. These arbitrary numbers had a marked effect on estimates. The median estimates of the percentage of African countries in the United Nations were 25 and 45 for groups that received numbers 10 and 65 from the wheel, respectively.

Payoffs for accuracy did not reduce the effect of the anchor.

35.5 Heuristics examples

35.5.1 A used car

You are shopping for a used car. You see a car you like and ask the salesperson how much it costs.

She says "one hundred thousand dollars". You know this number is too high, and after some negotiation, purchase the car for \$20,000. You feel pleased with your negotiation skills.

You later see a similar car for sale for \$15,000.

What heuristic could lead to your pattern of behaviour?

This pattern of behaviour could be caused by anchoring and adjustment.

When using anchoring and adjustment, people estimate by starting from an initial value and adjust from that value to obtain the final estimate. Empirically, people tend to use weak or irrelevant anchors and make insufficient adjustments from the anchor.

When the car dealer stated the high price, this acted as an anchor, even though you knew it was too high. You used a weak anchor.

That you ultimately purchased the car for too much suggests you insufficiently adjusted for that weak anchor.

35.5.2 A wealthy person

You see a person who drives a luxury car and wears designer clothes. You decide they must be wealthy, even though you have no other information about this person.

a) What heuristic could lead to this belief?

The representativeness heuristic could cause this belief.

Under the representativeness heuristic, people evaluate probabilities by the degree to which A is representative of (similar to) B.

In this case, a person driving a luxury car and wearing designer clothes is highly representative of a wealthy person. You place a high probability on them being wealthy.

b) Explain how using this heuristic would differ from using Bayes' rule in this situation.

Under Bayes' rule, the probability that someone is wealthy is a function of:

• the probability that any particular person in the population is wealthy (the base rate that forms your prior probability)

- the probability that a wealthy person will drive a luxury car and wear designer clothes
- the probability that a non-wealthy person will drive a luxury car and wear designer clothes.

They take that prior probability and update it based on the evidence they have observed.

Bayes' rule differs from that under the representativeness heuristic in that it considers the population's base rate. What proportion of people are wealthy? The representativeness heuristic does not. The representativeness heuristic is largely based on the probability that a wealthy person will drive a luxury car and wear designer clothes relative to a non-wealthy person - that is, how representative clothing and cars are of wealth.

Chapter 36

Biases in probability judgment

Summary

- The conjunction fallacy occurs when people judge the probability of two events occurring together as higher than the probability of one of those events occurring alone, violating basic probability rules.
- Base-rate neglect is the tendency to ignore general probabilities (base rates) when making judgments about specific cases, often leading to errors in conditional probability calculations.
- Probability matching is a suboptimal strategy where people predict outcomes in proportion to their observed frequencies, rather than always choosing the most likely outcome.
- The gambler's fallacy is the mistaken belief that if an event occurs more frequently than normal during a given period, it will occur less frequently in the future (or vice versa).
- The hot hand fallacy is the belief that a person who has experienced success with a random event has a greater chance of further success in additional attempts. Analysing this belief can be complicated by statistical biases in measurement and the potential existence of genuine "hot hands" in some contexts.

36.1 Introduction

In this part, I introduce several biases in probability judgment:

• The conjunction fallacy

- Base-rate neglect
- Probability matching
- The gambler's fallacy
- The hot hand fallacy

36.2 The conjunction fallacy

The conjunction fallacy occurs when someone judges the probability of the conjunction of two events to be greater than the probability of one or both events.

For example, if we have two outcomes, A and B, the probability of both A and B occurring - that is, the conjunction of A and B - should be less than or equal to each of the individual probabilities.



Figure 36.1: The conjunction of P(A) and P(B)

The most famous example of the conjunction fallacy comes from Tversky and Kahneman (1983). They asked students to read the following statement:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Tversky and Kahneman asked the students to rank the following statements

from most to least probable:

- 1. Linda is a teacher in elementary school.
- 2. Linda works in a bookstore and takes Yoga classes.
- 3. Linda is active in the feminist movement.
- 4. Linda is a psychiatric social worker.
- 5. Linda is a member of the League of Women Voters.
- 6. Linda is a bank teller.
- 7. Linda is an insurance salesperson.
- 8. Linda is a bank teller and is active in the feminist movement.

Note that "8. Linda is a bank teller and is active in the feminist movement" is a conjunction of "3. Linda is active in the feminist movement" and "6. Linda is a bank teller".

Tversky and Kahneman found in a sample of students that 88% ranked 3 before 8 before 6. "6. Linda is a bank teller" was rated less probable than "8. Linda is a bank teller and is active in the feminist movement".

To understand why this is an error, recall that the probability of the conjunction of two outcomes is as follows:

$$P(A \cap B) = P(A|B)P(B) = P(B|A)P(A)$$

If P(A|B) < 1 and P(B|A) < 1, $P(A \cap B)$ must be less than P(A) or P(B).

One explanation for why people make this error relates to the representativeness heuristic.

Tversky and Kahneman constructed the description of Linda to be representative of a feminist and unrepresentative of a bank teller. If people use the representativeness heuristic to order the statements, they will likely rank 8 above 6.

The Linda Problem is one of the most heavily debated experiments in the social sciences.

For example, Hertwig and Gigerenzer (1999) argue that people infer nonmathematical meaning to the word "probability", taking it to mean "plausible" or "credible".

While this is possibly a fair critique of the Linda problem, other illustrations of the conjunction fallacy appear more robust.

For example, Tversky and Kahneman (1983) created this example involving rolls of a die:

Consider a regular six-sided die with four green faces and two red faces. The die will be rolled 20 times and the sequence of greens (G) and reds (R) will be recorded. You are asked to select one sequence, from a set of three, and you will win \$25 if the sequence you chose

appears on successive rolls of the die. Please check the sequence of greens and reds on which you prefer to bet.

- 1. RGRRR
- 2. GRGRRR
- 3. GRRRRR

65% of experimental subjects chose sequence 2. It appears more "representative" of a die with four green faces and two red faces. But note that 1 is contained within 2 and is strictly more likely. The fact subjects are betting on the outcome should remove doubt about interpretation.

36.3 Base-rate neglect

The base rate is the probability of an outcome unconditional on any evidence.

For example, if 1% of the population has COVID-19 and the remainder doesn't, the base rate of COVID-19 is 1%. If you were to obtain evidence that someone has COVID-19, such as a positive COVID-19 test, you would use that base rate in determining the conditional probability that they have the disease.

Base rate neglect is the failure to consider an event's base rate when making a judgment.

36.3.1 The cab problem

One illustration of base-rate neglect comes from the cab problem by Tversky and Kahneman (1982). It involves the following story:

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. Participants are given the following data:

- 1. 85% of the cabs in the city are Green, 15% are Blue.
- 2. A witness identified the cab as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colours 80% of the time.

What is the probability that the cab involved in the accident was Blue rather than Green?

In the experiment, the median and modal answer was 80%.

The correct answer is 41%.

The experimental result indicates confusion between conditional probabilities. The experimental participants were confusing the probability of the witness identifying a blue cab given that the cab was blue, with the probability of the cab being blue given that the witness identified it as blue. However, we need to use Bayes' rule to calculate the probability of the cab being blue, given that the witness identified it as blue.

 $\underbrace{P(\text{claim blue}|\text{blue})}_{80\%} \neq \underbrace{P(\text{blue}|\text{claim blue})}_{\text{Requires Bayes' rule}}$

The experimental subjects effectively neglected the rarity of blue cabs. A witness seeing a blue cab is representative of what would occur if the cab were blue.

The correct answer is as follows:

$$\begin{split} P(\text{blue}|\text{claim blue}) &= \frac{P(\text{claim blue}|\text{blue})P(\text{blue})}{P(\text{claim blue})} \\ &= \frac{P(\text{claim blue}|\text{blue})P(\text{blue})}{\left(\begin{array}{c} P(\text{claim blue}|\text{blue})P(\text{blue}) \\ + P(\text{claim blue}|\neg\text{blue})P(\neg\text{blue}) \end{array}\right)} \end{split}$$

$$= \frac{0.8 \times 0.15}{0.8 \times 0.15 + 0.2 \times 0.85}$$
$$= 0.41$$

36.3.2 Medical diagnosis

We can also see base rate neglect in the context of diagnosing a rare disease.

Consider the following problem:

You test yourself for COVID-19. The following information is known:

- The probability that a person has COVID-19 is 1% (the prevalence).
- If a person has COVID-19, the probability that they test positive is 90% (the sensitivity).
- If a person does not have COVID-19, the probability that they nevertheless test positive is 9% (the false positive rate).

You test positive. What is the chance that you have COVID-19?

When problems of this nature are given to physicians, around 10 to 20% reason using Bayes' rule (for example, see Hoffrage et al. (2015)). The most common answers approximate the sensitivity, 90% for this example.

As for the cab problem, there is confusion between the conditional probabilities.

 $P(\text{COVID}|+\text{ve}) \neq P(+\text{ve}|\text{COVID})$

One hypothesis for this error is that a positive test is "representative" of someone with COVID-19. As a result, the test is given greater weight than the more general information about the base rate.

The correct answer is:

$$P(\text{COVID}|+\text{ve}) = \frac{P(+\text{ve}|\text{COVID})P(\text{COVID})}{P(+\text{ve})}$$
$$= \frac{P(+\text{ve}|\text{COVID})P(\text{COVID})}{\left(P(+\text{ve}|\text{COVID})P(\text{COVID}) + P(+\text{ve}|\neg\text{COVID})P(\neg\text{COVID})\right)}$$

$$= \frac{0.9 \times 0.01}{0.9 \times 0.01 + 0.09 \times 0.99}$$
$$= 0.092$$

36.3.3 Natural frequencies

Let us reconsider this medical problem with an alternative representation. This representation uses "natural frequencies".

You test yourself for COVID-19. The following information is known:

- Ten in every 1000 people have COVID-19 (the prevalence).
- Of these 10 people with COVID-19, nine will test positive (the sensitivity).
- Of the 990 people without COVID-19, about 89 nevertheless test positive (the false positive rate).

You test positive. What is the chance that you have COVID-19?

Seeing a representation in this manner makes the base rate (and the rate of false positives) much more salient, and leads to more accurate estimates of the conditional probabilities. We can see that the probability that we have

COVID-19, given we tested positive for COVID-19, equals the number of people who have COVID-19 who have tested positive, divided by the total number of positive tests:

$$P(\text{COVID}|+\text{ve}) = \frac{n(+\text{ve} \cap \text{COVID})}{n(+\text{ve})}$$
$$= \frac{9}{9+89}$$
$$= 0.092$$

Cosmides and Tooby (1996) first proposed using natural frequencies in this way. We derive natural frequencies by observing cases representatively sampled from a population.

Hoffrage and Gigerenzer (1998) reported that this change in representation increased the proportion of correct answers among physicians from 10% to 46%.

There is evidence that you can get further gains through a frequency tree representation (e.g. Spiegelhalter and Gage (2015)). Below is one such tree from Gigerenzer (2011), which they compare with a tree using conditional probabilities.



The numbers at the bottom of the conditional probability tree do not contain the base rate information. You can't simply compare them to calculate conditional probabilities. You need to refer to the middle layer. Conversely, the natural frequency tree contains all you need to calculate the conditional probability in the bottom row.

To illustrate this point, consider what happens if we convert the numbers at the bottom of the conditional probability tree into frequencies: 900 in 1000, 10 in 1000, 90 in 1000 and 910 in 1000. Gigerenzer calls these simple frequencies. While simple frequencies can make a problem more tractable, they do not allow us to calculate conditional probabilities. Simple frequencies are just a restatement of the probabilities. In contrast, natural frequencies are joint frequencies, such as the number of people who test positive and who have COVID-19.

36.3.3.1 Identifying a finch

The following example provides another illustration of the use of Bayes' rule and natural frequencies.

You are trying to spot a rare type of bird, the Darwin finch. It looks very similar to the Wallace finch, except for a slight difference in the shape of its beak. You know the following about the finches in your area:

- 99% of the finches are Wallace finches. The remaining 1% are Darwin finches.
- If you spot a Darwin finch, you will correctly identify it as a Darwin finch 95% of the time. The other 5% of the time, you identify it as a Wallace finch.
- If you spot a Wallace finch, you will correctly identify it as a Wallace finch 95% of the time. The other 5% of the time, you identify it as a Darwin finch.

You spot a finch and identify it as a Darwin finch.

What is the probability that the finch is a Darwin finch?

First, I use Bayes' rule to calculate the probability.

$$P(D|I) = \frac{P(I|D)P(D)}{P(I)}$$

= $\frac{P(I|D)P(D)}{P(I|D)P(D) + P(I|\neg D)P(\neg D)}$
= $\frac{0.95 \times 0.01}{0.95 \times 0.01 + 0.05 \times 0.99}$
= 0.16

The probability that it is a Darwin finch is 16%.

Next, I use natural frequencies to calculate that same conditional probability.

Suppose there are 10,000 finches.

That would mean there are 100 Darwin finches and 9,900 Wallace finches.

If I spotted these 100 Darwin finches, I would identify 95 as Darwin finches.

If I spotted a Wallace Finch, I would identify $0.05\times9900=495$ as Darwin Finches.

That means 95 of the 95+495=590 birds I identify as Darwin finches would be Darwin finches.

Therefore:

$$P(D|I) = \frac{95}{590}$$
$$= 0.16$$

Note that in this example, I have started with a number of finches, 10,000, which allows me to avoid fractions and small decimals. If I started with only 100 finches, I would later be talking about an unintuitive 4.95 finches. If you are using natural frequencies to solve a problem of conditional probability, you should choose a large enough number to avoid complicated fractions and decimals. Alternatively (or in conjunction), round any unintuitive numbers to the nearest whole number, giving you an approximate answer in your final calculation.

36.4 Probability matching

Probability matching is the tendency of people to mirror the probability distributions they observe in their predictions of events. For example, if asked to predict whether a die will show a six or not, they will predict six around one in six rolls.

The strategy of probability matching is not optimal for minimising prediction error.

Consider the following experimental setup:

- A red lamp that turns on with probability p = 0.70
- A green lamp that turns on with probability q = 0.30

Participants predict which light will turn on after observing a series of flashes.

What do participants do?

The predictions tend to reflect the actual probabilities of the two light bulbs being turned on. People tend to predict 70% of the time that the red lamp will come on and 30% of the time that the green lamp will come on.

With probability matching, the probability of a successful guess is:

$$p(\text{success}) = 0.7 \times 0.7 + 0.3 \times 0.3$$

= 0.58

A better strategy is always to select the event with the highest probability. In this example, participants should always predict that the red light will be turned on, giving them a 70% probability of a successful guess.

Similarly, for my earlier example of a die roll, the option with lowest error is always to predict that the die will not show a six.

36.5 The gambler's fallacy

The gambler's fallacy is the false belief that an outcome not recently realised in a sequence of independent draws is more likely to occur on the next draw.

For example, following three flips of a coin that all come up heads, a person experiencing the gambler's fallacy would believe that a tail is more likely on the next flip.

Using data from Rapoport and Budescu (1997), Rabin and Vayanos (2010) derived the probability of heads predicted by experimental subjects, given the last three flips being heads or tails. Following a sequence of three heads, they predict heads on the next flip with only 30% probability. But after three tails, they predict heads on the next flip with 70% probability.

3rd-to-last	2nd-to-last	Very last	Prob. next will be H (%)
Н	Н	Н	30.0
Т	Н	Н	38.0
Н	Т	Н	41.2
Н	Н	Т	48.7
Н	Т	Т	62.0
Т	Н	Т	58.8
Т	Т	Н	51.3
Т	Т	Т	70.0

One explanation for the gambler's fallacy is representativeness. For example, people do not see the sequence of coin flips HHHHHH as representative of flipping a fair coin six times. They see HHTTHH as more representative, even though both sequences have the same probability of occurring.

36.5.1 The law of small numbers

An alternative explanation is that people believe in the "law of small numbers" (Rabin, 2002). They overestimate the degree to which a small sample will resemble the population from which it is drawn. For example, if a fair coin is flipped six times, they will overestimate the likelihood the result will be three heads and three tails.

Imagine an urn filled with red balls and black balls. You draw balls from the

urn with replacement. The red balls are drawn with probability p and the black balls are drawn with probability 1 - p.

Assume Freddy knows the probabilities p and 1-p but (wrongly) assumes balls are drawn from the urn without replacement. If he believes there are N balls in the urn, he expects a sample of N balls to match p and 1-p exactly.

Under Freddy's beliefs, outcomes are correlated. Under the actual process, where balls are replaced, the outcomes are uncorrelated.

Imagine Freddy plays roulette. The roulette wheel contains 36 slots, 18 black and 18 red. Assume that Freddy believes there are 18 red and 18 black "balls in the urn".

Freddy observes four spins of the wheel before betting. He observes a sequence of four reds.

An unbiased belief would be that the sequence of reds tells him nothing about future draws because the outcomes are uncorrelated.

However, Freddy believes that, after four reds, black is more likely on the next spin. He is wrongly computing the probability based on a belief that only 14 reds remain along with 18 blacks.

$$\hat{P}(RRRRR|RRRR) = \frac{\text{reds}}{\text{reds} + \text{blacks}}$$
$$= \frac{18 - 4}{18 - 4 + 18}$$
$$= 0.438$$

In reality, P(RRRRR|RRR) = 0.5. Freddy is suffering from the gambler's fallacy.

36.6 The hot hand fallacy

A person subject to the hot hand fallacy believes a streak will persist despite each outcome being independent of the last.

For example, suppose a spectator observes a basketball player taking a series of shots during a game. The spectator then makes predictions based on the observed shots, with good shots predicted to be more likely following a streak of successful shots. After a series of good shots, they believe the player has a "hot hand".

Let's look at this example in more detail.

Suppose a person takes ten shots in a basketball game. In this image, a ball is a hit, an X is a miss.

To assess whether this person has a hot hand, we can look at their shots following a previous hit. For instance, in this sequence of shots, there are six occasions where we have a shot following a hit. Five successful shots, such as the highlighted seventh shot, are followed by another hit.



We can then compare the player's average shooting percentage with the proportion of shots they hit if the shot immediately before was a hit. If their hit rate after a hit is higher than their normal shot probability, we might say they get a hot hand.

Using this methodology, Gilovich et al. (1985) took shot data from various sources, including the Philadelphia 76ers and Boston Celtics, and examined the data for evidence of a hot hand. They also looked at whether there was a hit or miss after streaks of hits or misses.

From this data, they argued that the hot hand was an illusion. There was no evidence that a player was more likely to make a shot following a series of successful shots.

36.6.1 A bias in sequences

Gilovich et al. (1985) was the first of many examinations of whether there is a hot hand in sports (see Bar-Eli et al. (2006)). Through this research, there has been many methodological debates and arguments about whether there might be bias in the data, such as teams adjusting their defence in response to a player with a hot hand. However, the general trend in the literature was a finding of no evidence of a hot hand.

Miller and Sanjurjo (2018) provided a compelling critique of this position. They found a statistical bias in the analysis by Gilovich et al. (1985) and many others. The intuition behind the statistical bias is as follows.

Suppose you flip a coin three times. There are eight possible sequences of heads and tails. Each sequence has an equal probability of occurring.

Considering these sequences, if you were to flip a coin three times, and there is a head followed by another flip in that sequence, what is the expected probability that another head will follow that head?

Table 36.1 shows the proportion of heads following a previous flip of heads for each sequence. In the table's first row, HHH, the first flip is a head. Another

head follows that first flip. After the second flip, a head, we also have a head. There is no flip after the third head. 100% of the heads in that sequence followed by another flip are followed by a head.

In the second row of the table, HHT, a head follows 50% of the heads.

In the third row, there is one head followed by another flip, which is a tail. None of the heads in that sequence are followed by a head.

And so on until the last two rows, where there are no heads followed by another flip.

Now, back to our question. If you were to flip a coin three times, and there is a head followed by another flip in that sequence, what is the expected probability that another head will follow that head? It turns out the answer to this question is 42%. I get this number by calculating the expected probability of a head given any particular sequence. This is equal to the average of the probabilities in each sequence.

Table 36.1: Eight possible combinations of heads and tales across three flips

Flips	$p(H_{t+1} H_t)$
HHH	100%
HHT	50%
НТН	0%
HTT	0%
THH	100%
THT	0%
TTH	-
TTT	-
Expected probability	41.7%

That calculation contrasts with what we get we count across all of the sequences, where we see eight flips of head followed by another flip. Of the subsequent flips, four are heads and four are tails, which is the 50% you expect.

Why do we find that difference? By looking at these short sequences, we are introducing a bias. The cases of heads following heads tend to cluster together, such as in the first sequence, which has two cases of a head following a head. Yet the sequence THT, which has only one flip occurring after a head, is equally likely to occur as HHH. A tail appears more likely to follow a heads because of this bias, whereby the streaks tend to cluster together. The expected probability I get when taking a series of three flips is 42%, when in fact the actual probability of a head following a head is 50%. As the sequence of flips gets longer, the bias reduces in size, although it increases if we examine longer streaks, such as the probability of a head after three previous heads.

The net effect of this bias is that the measure of the proportion of heads following another head is biased downwards.

This bias is relevant to the analysis of the hot hand as it is present in the methodology of the papers that purportedly demonstrated that there was no hot hand in basketball, such as that by Gilovich et al. (1985). They effectively took short streaks of shots and calculated the proportion of hits followed by another hit. Their measure of the proportion of hits following a hit or sequence of hits is biased downwards. Like our calculation using coins, a calculation using that method results in a number lower than the actual probability of hitting a shot.

Conversely, the hot hand pushes the probability of hitting a shot after a previous hit up. If there is a hot hand, we should see more hits following a previous hit.

Now consider the net effect of these two forces. If there is a hot hand, the probability of hitting a shot after a previous hit should be higher than the average hit rate. The biased methodology pushes the measure in the other direction. Together, the downward bias and the hot hand counteract each other. In the case of Gilovich et al. (1985), these two countervailing forces led to the conclusion by researchers that each shot is independent of the last.

However, if you use a methodology not subject to this bias, you get a true measure of the hot hand. And in the case of the Gilovich et al. data, removing the bias reveals a hot hand. Miller and Sanjurjo (2018) found that in the Gilovich et al. data the probability of hitting a shot following a sequence of three previous hits is 13 percentage points higher than after a sequence of three misses.

36.6.2 Alternative intuition

Here is another way of showing that there is a bias in this sequence.

To do this, we will use Bayes' rule with more than two variables. This operates in a similar manner to our previous use of Bayes' rule.

To understand this, suppose we have three possible outcomes, A, B and C. For these outcomes we can write the following probabilities:

$$P(A \cap B \cap C) = P(A \cap B|C)P(C) = P(A|B \cap C)P(B \cap C)$$
$$= P(B|A \cap C)P(A \cap C) = P(C|A \cap B)P(A \cap B)$$

And so on. We can write the joint probability of these events as varying combinations of the conditional probabilities.

Typically we derive Bayes' rule by equating any two of these equations. For instance, as $P(A|B \cap C)P(B \cap C) = P(B|A \cap C)P(A \cap C)$ we can rearrange this to write:

$$P(A|B \cap C) = \frac{P(B|A \cap C)P(A \cap C)}{P(B \cap C)}$$

We will use this equation in our example.

Now, suppose we flip three coins and select at random one of the flips that follows a heads. This means that if we select a flip that follows a head we will select either flip 2 or flip 3.

If we select flip 2, we know that flip 1 was a head. The first two flips in the sequence are either HT or HH.

However, we can also say that if we select flip 2, HT is twice as likely as HH. Why? Because if the first two coins were HH we could also have chosen flip 3.

That is, if the first two flips are HT, we can only select flip 2. We select flip 2 with 100% probability. If the first two flips are HH, we select flip 2 with 50% probability and flip 3 with 50% probability.

We are twice as likely to observe HT as HH, given we selected flip 2.

Using H_i or T_i to represent a head or tail on the *i*-th flip and X_i to represent selection of flip *i*, we can show the probability of a tail given we have selected flip 2 using Bayes' rule. Using the equation we derived earlier involving three potential outcomes:

$$\begin{split} P(T_2|H_1 \cap X_2) &= \frac{P(X_2|H_1 \cap T_2)P(H_1 \cap T_2)}{P(H_1 \cap X_2)} \\ &= \underbrace{\frac{P(X_2|H_1 \cap T_2)P(H_1 \cap T_2)}{\left(\frac{P(X_2|H_1 \cap T_2)P(H_1 \cap T_2)}{+P(X_2|H_1 \cap H_2)P(H_1 \cap H_2)}\right)} \end{split}$$

Expand denominator using formula for total probability

$$= \frac{1 \times 0.25}{1 \times 0.25 + 0.5 \times 0.25}$$
$$= \frac{2}{3}$$

$$\begin{split} P(H_2|H_1 \cap X_2) &= \frac{P(X_2|H_1 \cap H_2)P(H_1 \cap H_2)}{P(H_1 \cap X_2)} \\ &= \underbrace{\frac{P(X_2|H_1 \cap H_2)P(H_1 \cap H_2)}{\left(\frac{P(X_2|H_1 \cap T_2)P(H_1 \cap T_2)}{+P(X_2|H_1 \cap H_2)P(H_1 \cap H_2)}\right)} \end{split}$$

Expand denominator using formula for total probability

$$= \frac{0.5 \times 0.25}{1 \times 0.25 + 0.5 \times 0.25}$$
$$= \frac{1}{3}$$

As you can only select flip 2 if flip 1 is a head, we can also say that $P(T_2|H_1 \cap X_2) = P(T_2|X_2) = \frac{2}{3}$ and $P(H_2|H_1 \cap X_2) = P(H_2|X_2) = \frac{1}{3}$. That is, the probability of a tail given we have selected flip 2 is 2/3. The probability of a head given we have selected flip 2 is 1/3. We are twice as likely to observe T_2 as H_2 , given we have selected flip 2.

We don't see the same bias if we select flip 3.

If we select flip 3, we know that flip 2 was a head. But the fact we select flip 3 does not tell us anything about what flip 3 is, as flip 3 itself does not influence the choice of flip. Whether flip 3 is a head or tail is independent of the choice of flip 3 or the outcome of flip 2.

Accordingly:

$$P(T_3|H_2 \cap X_3) = P(T_3) = 0.5$$

$$P(H_3|H_2 \cap X_3) = P(H_3) = 0.5$$

We now combine the results of our examination of the second and third flip.

We are equally likely to select flip 2 or flip 3 as flips 1 and 2 will each be heads with 50% probability. If both are heads, we select one randomly. Given we have selected a flip, what is the probability that the following flip is a head?

$$\begin{split} P(H) &= P(X_2) \times P(H_2|X_2) + P(X_3) \times P(H_3|X_3) \\ &= 0.5 \times 0.33 + 0.5 \times 0.5 \\ &= 0.417 \end{split}$$

What does this mean for measurement of the hot hand?

As for before, if I take a sequence of three flips and I look at a flip after a head, if there is at least one head, the probability that flip is a head is 0.42. This is despite the coin flips being independent. It appears I have a cold hand.

Use that same methodology in a scenario where there is a hot hand, the bias will counteract the hot hand and make it harder to detect, if it can be detected at all.

36.6.3 The hot hand fallacy for truly random sequences

Despite the evidence that there is a hot hand in some sports, there is strong evidence that there still exists a "hot hand fallacy". People see streaks in truly random processes, with each outcome independent of the last.

For example, Ayton and Fischer (2004) found that when people predict the results of a roulette wheel's spins, they increase their confidence in their predictions after a series of successes. Their confidence increases despite the outcome being random. Interestingly, they also exhibit the gambler's fallacy in what they predict.

Chapter 37

Heuristics and the bias-variance trade-off

Summary

- Heuristics, while often associated with errors, can be powerful and effective decision-making tools due to the bias-variance trade-off.
- The bias-variance trade-off suggests that as model complexity increases, bias tends to decrease, but variance (sensitivity to data fluctuations) increases. There's an optimal level of complexity that balances these two factors.
- Simple heuristics can sometimes outperform more complex decisionmaking strategies by finding a better bias-variance trade-off, leading to lower overall error despite using less information.
- The "Take the Best" heuristic, which uses cues in order of validity to make predictions, has been shown to perform as well as or better than multiple regression across various environments.
- The gaze heuristic, used for catching balls, demonstrates how a simple rule (maintaining a constant angle of gaze) can effectively solve a complex problem without requiring complex calculations, even if it results in a seemingly biased movement pattern.

37.1 Introduction

Much of the heuristics and biases literature of Kahneman, Tversky and those who followed in their footsteps focuses on the errors that can be caused by using heuristics. However, there are also powerful reasons why we use heuristics in

decision making.

One of the strongest arguments for the use of heuristics relates to what is called the bias-variance trade-off.

Suppose you are trying to make a prediction or develop an estimate based on historical data. There is a true underlying process that is generating the data. You plan to build a predictive model that should approximate the underlying process. You have a noisy data sample with which to develop it and you are trying to decide which predictors to include.

For example, you want to predict the level of dropout in a school. You have possible predictors such as attendance rates, family socio-economic status, the school's average SAT score, and the degree of parental involvement in the child's schooling. Which of those should you include in your model?

Bias is the degree to which there are erroneous assumptions in your model. The classic case of bias is when you have failed to include a relevant predictor. If you exclude relevant predictors, your predictive model will not include relevant relations between the predictors and the target output you are trying to predict. In the school example, to the extent any of these factors are linked to dropout rates, excluding them can bias your prediction.

However, the inclusion of too many predictors can lead to what is called variance, which is an error that arises because of the sensitivity of the model to fluctuations in the data you use to develop the model. It ultimately involves giving too much weight to irrelevant or marginally relevant information.

For example, if you included the school colours in your model, it may appear to give you a better model due to noise. But as soon as you used that model to make a new prediction, the inclusion of the irrelevant variable would likely backfire.

The following image gives one conception of bias and variance. An unbiased predictor will tend to centre on the target. A low variance predictor will tend to cluster. A low variance, low bias estimate is the best outcome.



However, as the term bias-variance trade-off suggests, you typically can't choose the minimum bias, minimum variance option. There is a trade-off between the two. As model complexity increases, bias tends to decrease, but variance tends to go up. There is an optimal level of complexity.



Model complexity

The result of this bias-variance trade-off means that heuristics can sometimes be better than more complex decision making strategies. This is not just because they are tractable for the human mind - unlike, say prospect theory calculations or Bayesian updating - but also because they find a better bias-variance tradeoff. Despite our focus on how heuristics can cause biases decisions, they can also lead to lower error.

37.2 Simple heuristics

In a chapter in Simple Heuristics That Make Us Smart, Czerlinski et al. (1999) describe a competition between some simple heuristics and multiple regression. Both were used to predict outcomes across 20 environments, such as school dropout rates and fish fertility.

One simple heuristic in their competitions was "Take the Best". This heuristic operates by working through variables in order of validity in predicting the outcome. For example, if you want to know which of two schools has the highest dropout rate, you ask which of the many possible predictive cues has the highest validity. If student attendance rate has the highest validity, and one school has lower attendance than the other, infer that that school with the lower attendance has the higher dropout rate. If the attendance rate is the same, look at the next cue.

Depending on the precise specifications, the result of the competition across the 20 environments was either a victory for Take the Best or at least equal performance with multiple regression. This is impressive for something that is less computationally expensive and ignores much of the data (or in other words, is biased).

The reason for this success was that the simpler models had lower variance. This enabled lower or similar total error to the more complex models that included all variables.

37.3 Example: The gaze heuristic

As another example of a heuristic in operation, consider the gaze heuristic.

The gaze heuristic is a tool that people - and dogs - use to catch balls. The heuristic is simply this: maintain the ball at a constant angle of gaze. If you move to keep this angle constant, you will end up where the ball lands. Obviously, this is easier than calculating where you should be from the velocity of the ball, the angle of flight, the effect of wind resistance and so on.



Figure 37.1: Source: Gigerenzer (2021)

But it results in a strange pattern of movement. Suppose you are close to the point where the ball is first hit into the air. As it rises you will tend to back away from the ball. As it then starts to fall, you will move back in. If it is hit up to the side of you, you will move to the ball in a curve. If you examine the path you took to catch the ball, you might call the curve a bias. However, it is actually the result of a very effective decision-making tool.

There are also some circumstances where the gaze heuristic works better than others. It tends to work best when the ball is already high in the air. If you catch sight of a ball hit straight up before it has risen far, using the heuristic for its entire flight could require first running away from the ball and then toward it.

Understanding this is a much richer understanding than saying that the catcher is biased because they did not run straight to where the ball was going to land. It also points to the power of heuristics. Try to train someone to run straight to where a ball will land and watch them fail. Don't see heuristics as poor cousins of "more rational" approaches.

Chapter 38

Overconfidence

Summary

- Overconfidence is a robust finding in judgment and decision-making research, but it actually encompasses three distinct phenomena: overprecision, overestimation, and overplacement.
- Overprecision is the tendency to believe our predictions or estimates are more accurate than they are, often demonstrated in confidence interval tasks.
- Overestimation is the belief that we can perform better than we realistically can, typically occurring for difficult tasks but not for easy ones.
- Overplacement is the erroneous judgment that we are better than others, which varies with task difficulty. People tend to overplace on easy tasks and underplace on difficult ones.
- Overconfidence has been linked to real-world phenomena such as excessive firm entry in markets and over-trading in financial markets, with experiments and studies demonstrating how it can lead to suboptimal outcomes in these domains.

38.1 Introduction

De Bondt and Thaler (1995) wrote "Perhaps the most robust finding in the psychology of judgment and choice is that people are overconfident."

Take the following examples:

• A person is asked to estimate the length of the Nile by providing a range that the respondent is 90% sure contains the correct answer. For example, they might answer that there is a 90% probability that the Nile is between
2500km and 5000km long. However, when people answer this question, the estimate typically contains the correct answer only 50% of the time.

- PGA golfers typically believe they sink around 75% of 6-foot putts some even believe they sink as many as 85% when the average is closer to 55%.
- 93% of American drivers rate themselves as better than average. 25% of high school seniors believe they are in the top 1% in their ability to get along with others.

There are many similar examples, all making the case that people are generally overconfident.

But despite each being labelled as overconfidence, note that these examples are three different phenomena.

Overprecision is the tendency to believe that our predictions or estimates are more accurate than they are. The typical study seeking to show overprecision asks for someone to give confidence ranges for their estimates, such as estimating the length of the Nile.

Overestimation is the belief that we can perform at a level beyond that which we realistically can. The evidence here is mixed. We typically overestimate when attempting a difficult task, such as a six-foot putt. But on easy tasks, the opposite is often the case – we tend to underestimate our performance. Whether over or underestimation occurs depends upon the domain.

Overplacement is the erroneous relative judgement that we are better than others. Obviously, we cannot all be better than average. But this relative judgement, like overestimation, tends to vary with task difficulty. For easy tasks, such as driving a car, we overplace and consider ourselves better than most. But, people will rate themselves below average for a skill such as drawing or identifying plants from the Amazon. People don't suffer from pervasive overplacement. Whether they overplace depends on what the situation is.

You might note that we tend to both underestimate and overplace our performance on easy tasks. We can also overestimate but underplace our performance on difficult tasks.

So, are we both underconfident and overconfident at the same time? The blanket term of overconfidence does little justice to what is occurring.

The conflation of these different effects under the umbrella of overconfidence often plays out in stories of how overconfidence (rarely assessed before the fact) led to someone's fall. For instance, evidence that people tend to believe they are better drivers than average (overplacement) is not evidence that overconfidence led someone to pursue a disastrous corporate merger (overestimation).

38.2 Firm entry

An example of overconfidence in action can be seen in firm entry.

Most new businesses fail within a few years. For example, one study of US manufacturers found over 60% of entrants had exited within five years and almost 80% within 10 years.

Camerer and Lovallo (1999) ran an experiment to test whether business failure may be due to optimism about their relative skill.

The lab experiment involved a set of markets. Those who chose to participate in a market were paid a set amount according to their rank within the market. Those ranked within the "market capacity" would share a payment of \$50. Those beyond the market capacity would be penalised \$10. Accordingly, if there are 5 entrants above market capacity, the expected payoff of all entrants is zero. More than that and it is negative.

The rank in the market was determined by either luck, through a random draw, or a test of skill involving logic puzzles or trivia questions about sports or current events.

In each round of the experiment, the market capacity was announced to the players, along with whether the payoffs in the market were based on luck or skill. The participants were then asked to forecast the expected number of entrants (for which they earn a payment if correct) and decide simultaneously and without communicating whether to enter into the market. Subjects were then told how many participants had entered.

After all of the rounds, students solved puzzles or took the trivia quiz to determine their skill rank.



The results of the experiment showed that more participants entered the market when the ranking was based on skill than if based on random draw. This indicates a belief that their skill level will rank them higher than a random draw: they are above average.

An interesting element to this experiment was that for some of the markets the participants were recruited by being asked if they would like to volunteer for an experiment in which performance would depend on their performance on sports or current event trivia questions. Hence the pool in those markets would be stronger than typical. In those markets with self-selected participants, market entry was even higher, and payoffs were negative in most rounds. This suggests the self-selected entrants were overconfident in their skill due to what Camerer and Lovallo call "reference group neglect". The participants seem to neglect that the others in the reference group also self-selected in to the experiment and think they are skilled too.

Moore et al. (2007) also ran an experiment on firm entry and found, like Camerer and Lovallo, that entrepreneurs overweight personal factors and underweight competitors when making entry decisions. However, when they varied the task difficulty, they found excess entry only when the industry appeared an easy one in which to compete. When it appeared difficult, too few entered. People overplaced in easy markets and underplaced in hard ones.

38.3 Trading

Another domain where overconfidence has been argued to play a role is related to trading.

A consistent finding in the analysis of trading behaviour is that more trading leads to poorer outcomes. The higher transaction costs are not compensated through higher returns.

To test whether over-trading may be linked to overconfidence, Barber and Odean (2001) examined investors by gender. Men tend to be more overconfident than women - a point they support with evidence of overprecision, overplacement and overestimation. If overconfidence leads to more trading, you would predict that men would trade more than women.

Barber and Odean (2001) examined trading account data from over 35,000 households for the period from February 1991 to January 1997. They found that men traded 45 percent more than women, reducing their net returns by 2.65 percentage points a year as opposed to 1.72 percentage points for women.

Chapter 39

Heuristics and biases exercises

39.1 Detecting a terrorist

Every month 100 million people fly on commercial airlines. Imagine 10 of them are terrorists.

Airport security are able to correctly identify that a person is a terrorist in 99% of cases and a non-terrorist in 99.9% of cases.

a) Intuitive responses to questions of this type tend to involve much higher probabilities. Discuss how intuitive responses could err due to confusion of conditional probabilities.

? Answer

If a person confuses $P(\text{terrorist} \mid \text{identified})$ with $P(\text{identified} \mid \text{terrorist})$ they will wrongly assume the probability that someone identified as a terrorist is a terrorist is 99%. This is a common explanation for mistakes of this nature: e.g. identification of cabs problem discussed in Section 36.3.1.

b) State and solve the question in part a) in terms of natural frequencies.

? Answer

Number of passengers: 100,000,000 Number of terrorists: 10 Number of terrorists identified as terrorists: $10 \times 0.999 \approx 10$ Number of non-terrorists identified as terrorists: $0.001 \times 100,000,000 =$ 100,000

Proportion of people identified as terrorists who are terrorists = $10/(10 + 100000) \approx 1$ in 10000

39.2 A fire alarm

You know the following statistics about fire:

- The probability of your house catching fire on any particular day is 1 in 10,000
- Your fire alarm correctly detects a house fire 95% of the time
- The probability that your fire alarm sounds on a day when there is no fire (a false alarm) is 1 in 100.

a) Your alarm goes off. What is the probability that your house is on fire?

Answer $P(\text{fire}|\text{alarm}) = \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm})}$ $= \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm}|\text{fire})P(\text{fire}) + P(\text{alarm}|\neg\text{fire})P(\neg\text{fire})}$ $= \frac{0.95 \times 0.0001}{0.95 \times 0.0001 + 0.01 \times 0.9999}$ = 0.0094The probability of a fire if the alarm goes off is 0.94%.

b) Many people given this problem estimate the probability of the house being on fire as close to 95%. Provide one possible explanation for this error.

? Answer

People often confuse P(A|B) with P(B|A). In this case, such confusion would lead them to conclude that P(fire|alarm) = P(alarm|fire) = 95%. You might also think of this as anchoring on the 95% and insufficiently adjusting from there.

Alternatively, people sometimes act as though they have assumed uniform

priors: e.g. 50:50 as to whether a fire or not. In that case:

$$\begin{split} P(\text{fire}|\text{alarm}) &= \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm})} \\ &= \frac{P(\text{alarm}|\text{fire})P(\text{fire})}{P(\text{alarm}|\text{fire})P(\text{fire}) + P(\text{alarm}|\neg\text{fire})P(\neg\text{fire})} \\ &= \frac{0.95 \times 0.5}{0.95 \times 0.5 + 0.01 \times 0.5} \\ &= 0.9896 \end{split}$$

c) Express and solve this problem using natural frequencies.

? Answer

- Your house will catch fire on 100 out of 1,000,000 days. (You could choose any base number of days I chose 1 million as gives round numbers for the following items.)
- Your fire alarm will correctly detect a house fire on 95 of those days.
- You will have a false alarm on 9,999 out of the 999,900 days without fire.

$$P(\text{fire}|\text{alarm}) = \frac{95}{95 + 9999}$$

= 0.0094

39.3 The law of small numbers

Lincoln observes performance by fund manager Neville. Neville may be a skilled, mediocre or unskilled manager:

- A skilled fund manager has a 75% chance of beating the market each quarter.
- A mediocre fund manager has a 50% chance of beating the market each quarter.
- An unskilled fund manager has a 25% chance of beating the market each quarter.

Lincoln knows these odds.

The performance of a fund manager is independent from quarter to quarter.

Consider the model we used to examine behaviour involving a belief in the law of small numbers whereby the decision maker acts as though the process has the character of balls being drawn out of an urn without replacement. Lincoln develops his beliefs using this model with N = 12.

a) Lincoln thinks Neville is mediocre. What does Lincoln believe is the probability that Neville beats the market in the first quarter?

? Answer

Lincoln thinks the realisation of Neville's performance is like drawing from an urn with N = 12 balls. Because he believes Neville is mediocre, he thinks half (6) of the balls are good-performance balls and half (6) are bad-performance balls. The likelihood of drawing a good ball (G) the first quarter is 6/12.

$$\hat{P}(G) = \frac{G}{N}$$
$$= \frac{6}{12}$$
$$= 0.5$$

b) Neville beats the market in the first quarter. What does Lincoln believe is the probability he does it again in the second quarter?

? Answer

In Lincoln's mind, the balls are not replaced once drawn. If Neville has a good first quarter, Lincoln believes that only five good balls remain. Therefore, Lincoln believes that the probability of Neville having a good second quarter is 5/11.

$$\hat{P}(GG|G) = \frac{G-1}{N-1}$$
$$= \frac{5}{11}$$

c) Neville beats the market again. What does Lincoln believe is the probability that he will do so in the third quarter?

? Answer

After two good quarters, Lincoln believes only four good balls remain. Therefore, the probability of Neville having another good quarter is 4/10.

$$\hat{P}(GGG|GG) = \frac{G-2}{N-2}$$
$$= \frac{4}{10}$$

d) Lincoln observes Jill, who he believes is a skilled fund manager. What does Lincoln believe is the probability of her having 10 consecutive periods of outperformance?

? Answer

Lincoln believes that in 12 quarters Jill will have nine quarters of outperformance. As a result, he does not believe it is possible for her to have ten consecutive periods of out-performance. After nine periods, only three balls are left in the urn. None of those balls are good.

e) What psychological bias does Lincoln's behaviour reflect? Explain.

? Answer

Each time Neville has a good quarter, Lincoln thinks it is less likely that Neville will have another. This is an example of gambler's fallacy. Lincoln thinks that Neville's sequence of performances should be the typical sequence of a mediocre fund manager, with the same number of good and bad quarters. This leads Lincoln to expect bad quarters to be more likely after a sequence of good quarters.

Similarly for Jill, Lincoln expects a string of success to correct itself and her record to revert to the average for a skilled manager.

39.4 Heuristics

For each of the following experiments from Tverksy and Kahneman (1974), explain what heuristic may be leading to the belief or decision.

a) Tverksy and Kahneman (1974) write:

Subjects were shown brief personality descriptions of several individuals, allegedly sampled at random from a group of 100 professionalsengineers and lawyers. The subjects were asked to assess, for each description, the probability that it belonged to an engineer rather than to a lawyer. In one experimental condition, subjects were told that the group from which the descriptions had been drawn consisted of 70 engineers and 30 lawyers. In another condition, subjects were told that the group consisted of 30 engineers and 70 lawyers. The odds that any particular description belongs to an engineer rather than to a lawyer should be higher in the first condition, where there is a majority of engineers, than in the second condition, where there is a majority of lawyers. Specifically, it can be shown by applying Bayes' rule that the ratio of these odds should be $(.7/.3^2)$, or 5.44, for each description. In a sharp violation of Bayes' rule, the subjects in the two conditions produced essentially the same probability judgments.

💡 Answer

This behaviour might be explained by representativeness. Tverksy and Kahneman (1974) write:

[S]ubjects evaluated the likelihood that a particular description belonged to an engineer rather than to a lawyer, by the degree to which this description was representative of the two stereotypes, with little or no regard for the prior probabilities of the categories.

The subjects used prior probabilities correctly when they had no other information. In the absence of a personality sketch, they judged the probability that an unknown individual is an engineer to be .7 and .3, respectively, in the two base-rate conditions. However, prior probabilities were effectively ignored when a description was introduced, even when this description was totally uninformative.

Suppose one samples a word (of three letters or more) at random from an English text. Is it more likely that the word starts with r or that r is the third letter? ... [M]ost people judge words that begin with a given consonant to be more numerous than words in which

b) Tverksy and Kahneman (1974) write:

the same consonant appears in the third position. They do so even for consonants, such as r or k, that are more frequent in the third position than in the first.

💡 Answer

This behaviour might be explained by availability. Tverksy and Kahneman (1974) write:

People approach this problem by recalling words that begin with r (road) and words that have r in the third position (car) and assess the relative frequency by the ease with which words of the two types come to mind. Because it is much easier to search for words by their first letter than by their third letter, most people judge words that begin with a given consonant to be more numerous than words in which the same consonant appears in the third position.

c) Tverksy and Kahneman (1974) write:

Two groups of high school students estimated, within 5 seconds, a numerical expression that was written on the blackboard. One group estimated the product

 $8\times7\times6\times5\times4\times3\times2\times1$

while another group estimated the product

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$$

...

The median estimate for the ascending sequence was 512, while the median estimate for the descending sequence was 2,250. The correct answer is 40,320.

? Answer

This behaviour might be explained by anchoring and adjustment. Tverksy and Kahneman (1974) write:

To rapidly answer such questions, people may perform a few steps of computation and estimate the product by extrapolation or adjustment. Because adjustments are typically insufficient, this procedure should lead to underestimation. Furthermore, because the result of the first few steps of multiplication (performed from left to right) is higher in the descending sequence than in the ascending sequence, the former expression should be judged larger than the latter.

d) Tverksy and Kahneman (1974) write:

In considering tosses of a coin for heads or tails ... people regard the sequence H-T-H-T-T-H to be more likely than the sequence H-H-H-T-T-T ... [or] ... the sequence H-H-H-T-H.

💡 Answer

This behaviour might be explained by representativeness. Tverksy and Kahneman (1974) write:

People expect that a sequence of events generated by a random process will represent the essential characteristics of that process even when the sequence is short. ...

[P]eople expect that the essential characteristics of the process will be represented, not only globally in the entire sequence, but also locally in each of its parts. A locally representative sequence, however, deviates systematically from chance expectation: it contains too many alternations and too few runs. Another consequence of the belief in local representativeness is the well-known gambler's fallacy. After observing a long run of red on the roulette wheel. for example, most people erroneously believe that black is now due, presumably because the occurrence of black will result in a more representative sequence than the occurrence of an additional red.

39.5 Overconfidence

Consider the following three statements. Suppose that each statement is an instance of overconfidence. For each statement name and define the form of overconfidence that provides the best explanation for the students' beliefs.

a) 90% of students believe they will score above the class average in the final exam.

? Answer

Overplacement.

Overplacement is the erroneous relative judgement that we are better than others.

b) 90% of students believe they will receive a high distinction.

💡 Answer

Overestimation.

Overestimation is the belief that we can perform at a level beyond that which we realistically can.

c) Arthur believes with 90% probability that he will score between 74% and 76% in the final exam.

Answer

Overprecision.

Overprecision is the tendency to believe that our predictions or estimates are more accurate than they actually are.

39.6 Lethal events

When people are asked the frequency of lethal events, they are often inaccurate. The following table lists those events most subject to under- or over-estimation of the frequency.

Most overestimated	Most underestimated
All accidents	Diabetes
Motor vehicle accidents	Stomach cancer
Tornadoes	Stroke
Flood	Tuberculosis
All cancer	Asthma
Fire and flames	Emphysema
Venomous bite or sting	
Homicide	

What heuristic could lead to this pattern of overestimation and underestimation? Why?

? Answer

The most overestimated events tend to be vivid events that are often the subject of news. The most underestimated are much less vivid and likely receive less coverage.

This pattern could be driven by the availability heuristic. When using the availability heuristic, people judge the frequency of events by the ease with which instances of those events come to mind. When asked to estimate the frequency of vivid events often in the news, instances of those events will easily come to mind. The availability heuristic will lead these events to be judged more probable.

Conversely, people will find it harder to call to mind those events which are less vivid and newsworthy, leading them to judge those events as being less frequent.

Part VII

Game theory

Summary

- Game theory studies strategic interaction between players, where outcomes depend on others' behaviour and beliefs about strategies.
- Components of a game include players, actions, strategies, information available, and payoffs.
- Players are typically assumed to be rational optimisers who understand the game and assume the same of other players.
- Games can be categorised as cooperative vs non-cooperative based on the ability to negotiate binding contracts.
- Games can also be categorised as simultaneous vs sequential, based on the timing of decisions and knowledge of others' actions.

Introduction

In many situations, your outcome depends on others' behaviour. Their outcome depends on your behaviour.

Similarly, your strategy will depend on your belief about others' strategy. Their strategy depends on their beliefs about your strategy.

Game theory studies this strategic interaction between players. We can solve strategic problems using the tools of game theory.

Components of a game

A game has the following components:

- First, the **players** of the game. Most of the games we examine in these notes involve two players.
- Second, the **actions** that each player can take; for example to contribute to a common pool or to defect.
- Third, the **strategies** that comprise a complete contingent plan of action. That is, for any given scenario or action by another player, a strategy specifies the action to be taken by the player.
- Fourth, the **information** available to players. In these notes, we generally assume perfect information.
- And finally, the **payoffs**. This comprises a complete summary of the value to each player of each set of actions.

The players

In game theoretical analysis, we typically assume that the players are rational optimisers who understand the game that they are playing. By rational, we mean that the player is aware of their alternatives, forms expectations about any unknowns, has preferences that conform to the axioms of completeness and transitivity and they choose the best option using some optimisation algorithm.

We also assume that the players assume other players are also rational optimisers who understand the game.

We weaken this assumption when we analyse behavioural game theory.

Types of games

There are many different types of games analysed in game theory. Some of the delineations between these games are as follows.

Cooperative versus non-cooperative games

First, games are often divided into cooperative and non-cooperative games.

In non-cooperative games, players are not allowed to negotiate binding contracts. In cooperative games, players can negotiate binding contracts that allow them to implement joint strategies.

In these notes, I will focus on non-cooperative games.

Simultaneous or sequential games

Second, games can involve simultaneous or sequential moves.

In a simultaneous move game, you make decisions without knowing the action of your rival.

In sequential games, players make sequential decisions knowing the other player's action. We will examine both of these types of game in these notes.

Chapter 40

Simultaneous-move one-shot games

Summary

- Simultaneous-move one-shot games involve players making decisions without knowing their rival's actions. These games are typically represented in a "normal" or "strategic" form using a matrix.
- The prisoner's dilemma is a classic example of such a game, where two prisoners must decide whether to confess or remain silent.
- In game theory, a dominant strategy is one that yields the highest payoff regardless of the opponent's actions. In the prisoner's dilemma, confessing is the dominant strategy for both players.
- A Nash equilibrium occurs when each player's strategy is the best response to the other's strategy. In the prisoner's dilemma, the Nash equilibrium is (Confess, Confess), despite this not being the optimal outcome for either player.
- The examples showcase different types of simultaneous-move one-shot games, demonstrating concepts such as multiple equilibria, coordination problems, and the tension between individual and collective outcomes.
 - The driving game demonstrates a coordination game with two Nash equilibria: (Left, Left) and (Right, Right).
 - Matching pennies is a zero-sum game with no pure-strategy Nash equilibria.
 - The stag hunt game shows a coordination game with two Nash equilibria: (Stag, Stag) and (Hare, Hare), highlighting the tension between cooperation and individual safety.
 - The public goods game illustrates the conflict between individual and collective interests, with the Nash equilibrium being no contribution,

despite full contribution being Pareto optimal.

40.1 Introduction

In a simultaneous-move one-shot game, you make decisions without knowing the action of your rival. This can be interpreted as either players making decisions at the same time or players making decisions before knowing the decisions of their rivals.

40.2 The normal form

We usually write simultaneous move one-shot games in the "strategic" or "normal" form. In this form, all of the monetary or non-monetary outcomes are represented in a matrix.

I will now illustrate the normal form of the game with a game called the prisoner's dilemma.

40.2.1 The prisoner's dilemma

The prisoner's dilemma is a classic simultaneous-move one-shot game. A pair of criminals have been captured following a crime. The police have sufficient evidence to convict them of a minor crime (e.g. trespass), but insufficient evidence to convict them of the major crime that has occurred (e.g. theft of the crown jewels).

The police place each prisoner in a separate cell where they cannot communicate with each other. The police then offer both prisoners a deal: confess and they will let them go free despite the minor crime, but they will then have the evidence required to give their criminal partner a massive sentence for the serious crime.

If neither confesses, the police will have insufficient evidence to get a conviction for the major crime, so they will both receive a short sentence for the minor crime. If both confess, they will both get a longer sentence, but with some reduction in sentence relative to if they didn't confess.

The normal form of the game is as follows:

Prisoners A and B have two actions available: to confess and to stay silent. The numbers in the matrix represent the payoffs from each combination of actions, in this case, the number of years they will serve in prison. A higher number is therefore a worse outcome. The left number in each cell of the matrix represents the payoff to the row player, Prisoner A. The number on the right of each matrix cell is the payoff to the column player, Prisoner B.

		Prisoner B	
		Confess	Silent
Prisoner A	Confess	5, 5	0, 20
	Silent	20, 0	1, 1

Figure 40.1: The prisoner's dilemma

For example, if both Prisoner A and Prisoner B choose to confess, they each receive a prison sentence of five years. If Prisoner A confesses and Prisoner B remains silent, Prisoner A gets off without a prison sentence, whereas Prisoner B gets twenty years.

Equipped with the normal form of the game, we can determine what each player wants to do in response to each action of the other player.

For example, we can see that if Prisoner B confesses, Prisoner A can either confess and receive five years in prison, or remain silent and receive 20 years in prison. They would choose to confess.

We indicate the preferred action in response to another player's action by circling the relevant payoff. For example:

		Prisoner B	
		Confess	Silent
Prisoner A	Confess	5,5	0, 20
	Silent	20, 0	1, 1

Figure 40.2: The prisoner's dilemma: representing preferred actions

If Prisoner B remains silent, Prisoner A could either confess and escape without a sentence, or remain silent and receive a sentence of one year in prison. They would prefer to confess.

We can then work through the same process for Prisoner B's actions.

If Prisoner A confesses, Prisoner B can either confess and receive five years in prison, or remain silent and receive 20 years in prison. They would choose to

confess.

If Prisoner A remains silent, Prisoner B could either confess, and escape without a sentence, or remain silent, and receive a sentence of one year in prison. They would prefer to confess.

Indicating this set of preferred actions in response to that of the other player gives us this completed matrix.



Figure 40.3: The prisoner's dilemma solved

40.3 Dominant strategies

Before examining this matrix further, I will now introduce the concept of the dominant strategy.

A strategy is dominant if it gives a higher payoff than every other strategy, for every strategy that your rivals play.

A strategy is strictly dominant if it gives a strictly higher payoff than every other strategy, for every strategy that your rivals play.

If you have a strictly dominant strategy, you should play it for sure.

In a dominant strategy equilibrium, all players choose a dominant strategy.

In the prisoner's dilemma, both players have a dominant strategy to confess. No matter what the other player does, confessing is better than remaining silent.

40.4 Nash equilibrium

Another important concept is the Nash equilibrium.

A set of strategies is a Nash equilibrium if every player is playing a best response to their rivals' strategies. No one has an incentive to change strategy.

A Nash equilibrium is self-enforcing and stable. If the players agree to play a certain way, they'll both do it. Unilateral deviations are not worthwhile.

The prisoner's dilemma has a single Nash equilibrium: (Confess, Confess). Visually, where the preferred response of both players to the other player's action falls within the same cell, this indicates a Nash equilibrium.

40.5 Simultaneous-move one-shot game examples

In this part, I will show some examples of simultaneous-move one-shot games.

40.5.1 The driving game

Consider the following game between two players deciding what side of the road to drive on. They can drive on the left or the right. If they both drive on the left or right when they approach each other, they will successfully pass. If one drives on the left and the other on the right, they will crash.

		Driver 2	
		Left	Right
Driver 1 -	Left	10, 10	-10, -10
	Right	-10, -10	10, 10

Figure 40.4: The driving game

What are the Nash equilibria?

If Driver 2 drives on the left, Driver 1 can either successfully drive on the left, or drive on the right and crash. They would choose to drive on the left. If Driver 2 drives on the right, Driver 1 can either successfully drive on the right, or drive on the left and crash. They would choose to drive on the right.

Similarly, if Driver 1 drives on the left, Driver 2 can either successfully drive on the left, or drive on the right and crash. They would choose to drive on the left. If Driver 1 drives on the right, Driver 2 can either successfully drive on the right, or drive on the left and crash. They would choose to drive on the right.

We can see from the matrix that there are two Nash equilibria. The Nash equilibria are (Left, Left) and (Right, Right). If both drivers are driving on the left, neither has an incentive to change their strategy. If both drivers are driving on the right, again, neither has an incentive to change.



Figure 40.5: The driving game solved

40.5.2 Matching pennies

The next game, called matching pennies, involves two players, Even and Odd, who each have a penny. Each player must select one side of the penny and simultaneously show the penny to the other player. If the pennies match, Even wins. If they don't match, Odd wins.

		Odd	
		Heads	Tails
Even	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

Figure 40.6: Matching pennies

What are the Nash equilibria?

To determine this, we work through the matrix as we did in the previous example.

If Odd shows heads, Even can either show heads and win, or show tails and lose. They would choose to show heads. If Odd shows tails, Even can either show tails and win, or show heads and lose. They would choose to show tails.

Similarly, if Even shows heads, Odd can either show tails and win, or show heads and lose. They would choose to show tails. If Even shows tails, Odd can either show heads and win, or show tails and lose. They would choose to show tails.



Figure 40.7: Matching pennies solved

There are no pure-strategy Nash equilibria for this game. For any combination of heads and tails, one of the players would want to change their choice.

There are what are called "mixed-strategy Nash equilibria" in this game, but mixed-strategy equilibria are beyond the scope of this subject.

40.5.3 The stag hunt

Consider the "stag hunt game" between two players deciding what animal they will hunt. Both hunters need to cooperate to catch the stag. They can catch a hare by themselves, but it provides less meat.

		Hunter 2	
		Stag	Hare
Hunter 1 -	Stag	3, 3	0, 1
	Hare	1, 0	1, 1

Figure 40.8: The stag hunt game

What are the Nash equilibria?

If Hunter 2 hunts the stag, Hunter 1 can either hunt the stag and catch it, or hunt the hare and catch it. They would choose to hunt the stag as it gives a payoff of 3 compared to 1. If Hunter 2 hunts the hare, Hunter 1 can either hunt the stag and not catch it, or hunt the hare and catch it. They would choose to hunt the hare as it gives a payoff of 1 compared to 0.

Similarly, if Hunter 1 hunts the stag, Hunter 2 can either hunt the stag and

catch it, or hunt the hare and catch it. They would choose to hunt the stag as it gives a payoff of 3 compared to 1. If Hunter 1 hunts the hare, Hunter 2 can either hunt the stag and not catch it, or hunt the hare and catch it. They would choose to hunt the hare as it gives a payoff of 1 compared to 0.



Figure 40.9: The stag hunt game solved

The Nash equilibria are (Stag, Stag) and (Hare, Hare). On either pair of strategies, neither player has incentive to change. It is an open question, however, as to which Nash equilibrium might emerge if they were to play the game.

40.5.4 The public goods game

The final game I will consider in this part is the public goods game.

In this game, each participant is given an initial endowment.

Each participant secretly and simultaneously chooses how much of their endowment they wish to contribute to a public pot.

The money in the public pot is multiplied by some amount and split evenly between the players. Typically, the multiple applied to the pot is greater than 1, but less than the number of players.

For example, five players might each be given \$10, with the pot doubled. Suppose they each contribute \$5 of their \$10 endowment to the pot. The \$25 contributed to the pot is multiplied by 2 to a total of \$50. Each player then receives \$10 from the pot, giving them \$15 in total.

In Nash equilibrium in the public goods game, nobody transfers anything to the pot. Any contributions are split between all players, so if there are more players than the multiple, which is normally the case by design, contributions result in a loss to that individual player.

Consider the previous game, but this time Player E contributes nothing. There is then \$20 in the pot, which is doubled to \$40. The pot is then split equally between the players, each receiving \$8 from the pot. The result is that Player E is better off having not contributed, ending with \$18, compared to the \$15 they would have received had they contributed \$5 like the other players.



Figure 40.10: The public goods game



Figure 40.11: The public goods game with defection $% \left({{{\mathbf{F}}_{{\mathbf{F}}}} \right)$

The Pareto optimal result, however, is for all players to contribute their full endowment and each receives back their multiplied contribution. However, the Pareto optimal result is not stable, as each player has an incentive to defect and contribute nothing.



Figure 40.12: The public goods game with cooperation

Chapter 41

Sequential games

Summary

- Sequential games involve players making decisions in order, with knowledge of previous actions. These games are typically represented in "extensive form" using a game tree.
- The centipede game is an example of a sequential game where players take turns deciding to "take" (end the game) or "pass" (continue), with increasing payoffs at each step.
- A subgame perfect Nash equilibrium is a solution concept where players' strategies constitute a Nash equilibrium in every subgame of the original game.
- The centipede game is solved using backward induction, resulting in a subgame perfect Nash equilibrium where the first player "takes" at the very first opportunity, despite this leading to a lower payoff than if players cooperated by passing.
- The other examples illustrate sequential games and their subgame perfect Nash equilibria, demonstrating concepts like backward induction and the importance of credible threats.
- The ultimatum game shows that rational play leads to minimal offers being accepted, with subgame perfect Nash equilibria of offering \$1 (or \$0) and accepting.
- The dictator game, with no strategic interaction, predicts the dictator will keep the entire endowment.
- The trust game demonstrates how lack of trust can lead to suboptimal outcomes, with the subgame perfect Nash equilibrium being no investment and no return.

41.1 Introduction

In sequential games, players make sequential decisions knowing the action of the other player.

41.2 The extensive form

Sequential games can be shown in what is called the "extensive form" representation. The extensive form representation explicitly shows the timing of play.

Payoffs are represented in a game tree.

I will now illustrate the extensive form with a game called the centipede game.

41.2.1 The centipede game

This centipede game has six decision nodes. At each node, a player can "take", and end the game, or they can "pass", increasing the total payoff. The other player then has a move.

The numbers 1 and 2 along the top of the centipede represent the decision nodes for two players. At the first node, player 1 has the choice to take or pass. If player 1 passes, player 2 has the choice to take or pass, and so on. At the final node, the game ends regardless of what player 2 chooses.



The payoff when a player takes and ends the game is represented by the numbers in the brackets. The first number is the payoff for player 1 and the second number is the payoff for player 2. For example, if player 1 takes at the first node, they receive a payoff of 1 and player 2 receives a payoff of 0. At the final node, if player 2 passes they receive a payoff of 5 and Player 1 receives a payoff of 6. If player 2 takes at that final node, they receive a payoff of 6 and Player 1 receives a payoff of 4.

41.3 Subgame perfect Nash equilibrium

Before examining this game, I will introduce the concept of a subgame perfect Nash equilibrium.

A subgame is a part of a game that can be played as a game itself. It begins at a single node and contains every successor node.

41.3.1 Solving the centipede game

For example, this final stage of the centipede game is a subgame.



As is this subset of the game.



A Nash Equilibrium is subgame perfect if every player plays the Nash Equilibrium in every subgame

We can solve for the subgame perfect Nash equilibrium of sequential games by backward induction. To do that we solve for the decision nodes at the end of the game first and then work our way back to the beginning of the game.

In our centipede game, using backward induction, player 2 at the final node will "take" for a payoff of 6 instead of passing for a payoff of 5. When marking choices in a sequential game, it is often useful to mark the option taken by the player, or that not taken, in addition to indicating the payoff they would receive.



At the node immediately before, player 1 will "take" for a payoff of 5 instead of passing, given player 2 will then take, giving player 1 a payoff of 4.



Therefore, at the node before, player 2 will take for a payoff of 4 instead of passing for a payoff of 3.

Therefore, at the node before, player 1 will take for a payoff of 3 instead of passing for a payoff of 2.

Therefore, player 2 at the node before will take for a payoff of 2 instead of passing for a payoff of 1.

And therefore, player 1 at the first node will take for a payoff of 1 instead of passing for a payoff of 0.



There is a unique subgame perfect equilibrium for the centipede game: $S_1 = (take, take, take)$ and $S_2 = (take, take, take)$, where S_1 and S_2 are the set of strategies for player 1 and player 2 respectively.

In the subgame perfect Nash equilibrium of the centipede game, player 1 takes at the first node.

41.4 Sequential game examples

In this part, I will discuss some sequential games and their subgame perfect Nash equilibria.

41.4.1 The ultimatum game

The first example is the ultimatum game.

The ultimatum game involves two players: the proposer and the responder.

The proposer is given a fixed amount of money m. They then offer a portion x of the sum m to the responder.

The responder can either accept or reject the offer. They make this decision knowing the fixed amount m held by the proposer and the offer x.

If the responder accepts, the responder receives the offer x and the proposer gets the remainder m-x. If the responder rejects, both players receive nothing.



Figure 41.1: The ultimatum game

Below is the extensive form of the ultimatum game with m = \$10 and an assumption that the offer must be a whole dollar amount. At the first node is the proposer. They can choose to offer any dollar sum between \$0 and \$10. Whatever the choice, the responder is at the next node. They can choose to accept or reject the offer. The payoffs of each set of actions is indicated in the brackets at the bottom of the game tree, with the first number being the proposer's payoff and the second number being the responder's payoff.



Figure 41.2: Ultimatum game tree

If we work through this game by backward induction, we can see that for any non-zero amount, the responder will accept the offer. The only time they might not accept is where the offer is 0, but they still might.

Given this, the proposer will offer \$0 or \$1 only.



Figure 41.3: The ultimatum game Nash equilibrium

We can say that there are two subgame perfect Nash equilibria. The first is for the proposer to offer \$1 and the responder to accept if offered \$1 and reject if offered \$0. The other (weak) subgame perfect Nash equilibrium is an offer of \$0 and acceptance.

More generally, game theory makes a clear prediction on the outcome of the ultimatum game. If the players have monotonic preferences - that is, more is better - the responder accepts any x > 0 (and possibly even if x = 0) and the the proposer offers the smallest amount the proposer can offer.

Where the strategy space is continuous (that is the offer could always be made smaller) the only subgame perfect Nash equilibrium is for the proposer to offer \$0 and the receiver to accept.

41.4.2 The dictator game

The next example is the dictator game.

In the dictator game, the dictator is given a fixed amount of money m. They then offer a portion x of the sum m to the receiver. The game then ends.

Exchange is unilateral. Receivers have an empty strategy set.



Figure 41.4: The dictator game

The standard game theory prediction is no interaction whatsoever. The dictator

maximises their payoff by keeping all of the endowment themselves, receiving payoff m (which is bolded).



Figure 41.5: The dictator game solved

41.4.3 The trust game

The final example is the trust game.

The trust game involves two players: a sender and a receiver

Both the sender and receiver are given an initial sum m.

The sender sends a share x of their m to the receiver. This amount x is often called the investment.

Before the investment is received by the receiver, it is multiplied by some factor k.

Therefore, the receiver receives kx.

The receiver then returns to the sender some share y of their total allocation m + kx.

The final outcome is the sender has m - x + y and the receiver has m + kx - y. We can represent these payoffs as:

$$(m-x+y,m+kx-y)$$

The extensive form of the game is as follows.



Figure 41.6: The trust game

Here is a numerical example.

Suppose the sender and receiver are given an initial sum of \$10.

The sender decides to send \$5 of their \$10 to the receiver.

This is multiplied by a factor of 3. Therefore, the receiver receives \$15 and now has \$25.

The receiver then returns to the sender 7.50 of their 25.

The final outcome is (10 - 5 + 7.50, 10 + 15 - 7.50) = (12.50, 17.50).



Figure 41.7: The trust game example

If both receivers have utility function u(x) = x the only subgame-perfect equilibrium is that the receiver will keep all their money, so the sender sends nothing.

We can see this by backward induction. The receiver can either return y for a payoff of 10 + 3x - y or return 0 for a payoff of 10 + 3x. The receiver will return 0.

One way to think about this problem is that the receiver is effectively playing a dictator game.

Working backwards, the sender therefore has a choice between sending x for a payoff of 10 - x or sending 0 for a payoff of 10. The sender will send 0.



Figure 41.8: The trust game solved

Relative to the Pareto optimal outcome whereby the sender's full endowment is tripled and they receive a positive return on their investment, both players are worse off under the equilibrium outcome.

Chapter 42

Asymmetric information

Summary

- Information asymmetry in markets can lead to market failures, even with rational agents. This is illustrated by the "market for lemons" example, where buyers' inability to observe product quality before purchase can result in only low-quality goods being sold.
- The winner's curse is a phenomenon in common-value auctions where the winning bidder tends to overpay. This occurs because winning implies that other bidders valued the item lower, suggesting the winner's estimate may be too high.
- The Nash equilibrium strategy in such auctions is to bid more conservatively than your private estimate of the value.

42.1 Introduction

To date in this section on game theory, I have assumed perfect information. That is, all players know the rules of the game, the available actions and the payoffs from each set of actions.

I will now explore two examples where we relax this assumption and allow the parties to have different information. However, we will retain the assumption of rational behaviour.

42.2 The market for lemons

This example draws on the work of Akerlof (1970).
An agent decides to buy a used car. Price p is fixed and quality is unobservable.

Suppose there are two types of cars: good cars and lemons. A car is good with probability q and a lemon with probability 1 - q.

The seller knows the type. To the seller, good cars are worth \$10,000 and lemons \$5,000.

To potential buyers, good cars are worth \$15,000 and lemons \$7,500.

Before the purchase, the buyer knows the types of cars in the market and the frequency of each. They only discover the type of car, however, after the purchase.

Given both car types are worth more to buyers than sellers, there should exist advantageous trades for both parties for both types of car. Selling is an efficient solution.

But what happens?

Let μ be the probability that a car that is sold is good. If sellers are willing to sell their good cars, $\mu = q$. If not, $\mu = 0$.

Therefore, the expected value of a car to a buyer is:

$$E=\mu 15000+(1-\mu)7500=7500+7500\mu$$

Hence, the buyer will be willing to pay up to $p = 7500 + 7500\mu$.

Given the value of each type of cars to sellers, they will sell a lemon if $p \ge 5000$ and a good car if $p \le 10000$.

Combining the conditions for the buyer and seller, a lemon will be sold if the price lies between the minimum required by the seller for the lemon and the maximum the buyer is willing to pay for the lemon. That is:

$$5000 \le p \le 7500 + 7500\mu$$

This relationship holds regardless of the value of μ , so the seller will always be willing and able to sell the lemon.

They will be able to sell the good car if:

$$10000 \le p \le 7500 + \mu7500$$

This relation can only hold if $\mu \geq 1/3$.

Assuming risk neutral buyers, we are left with two possible equilibria.

If $q \ge 1/3$, sellers sell both types of cars:

$$\mu = q \le 1/3 \to 10000 \le p^* \le 7500 + \mu7500$$

If q < 1/3, sellers sell only lemons:

$$\mu = 0 \rightarrow 5000 \le p^* \le 7500$$

Generalising what is happening here:

- 1. When buyers cannot observe product quality, sellers have an incentive to pass off lemons as good cars.
- 2. Rational buyers expect this seller behaviour and they lower their willingness to pay.
- 3. Sellers cannot sell good cars at high prices even though buyers would be willing to pay high prices for good cars.
- 4. At the lower prices, sellers only offer to sell lemons.

Information asymmetry is sufficient to result in a market failure even if the agents are rational.

42.3 The winner's curse

The second example involves a phenomenon called the winner's curse.

The winner's curse occurs in the context of common-value auctions.

A common-value auction is an auction in which the item for sale has the same value to all the bidders.

Examples include stocks, which all have one value, and oil, where the amount of oil in a tract is the same for all oil companies.

Common-value auctions contrast with private-value auctions in which bidders have different valuations for the item for sale. This typically occurs where the item's valuation reflects bidder tastes, such as art.

The winner's curse is a phenomenon in common value auctions whereby the winner tends to experience a loss.

Petroleum engineers invented the term in discussing why oil companies in the Gulf of Mexico had poor results in the 1950s through 1970s. Oil companies in the Gulf acquired drilling rights through auctions. Their rights tended to lead to losses or less in profits than expected. In hindsight, the winning bids were unreasonably high.

The winner's curse is widely documented in experimental settings and has been observed in corporate environments.

42.3.1 Winner's curse example

I will now walk through a numerical example of the winner's curse.

Company 1 and company 2 hire a geologist to estimate the value of an oil field. The honest geologist of each company privately reports their estimated valuation to the company. Company 1 learns v_1 and company 2 learns v_2 .

 v_1 and v_2 are uniformly distributed between \$0 and \$100 and independent.

Assume the true value of the oil field is the mean of v_1 and v_2 :

$$V = \frac{v_1 + v_2}{2}$$

The two companies simultaneously bid for the field in a first-price auction. The highest bid wins and pays their bid.

What should a company bid in this auction?

Suppose both companies bid the private valuation they receive. Company 1 receives $v_1 = 50$, bids \$50 and wins.

If they win, $v_1 = $50 > v_2$.

On average, in this state of the world company 2's signal is \$25 (due to the uniform distribution). The average value of the tract is therefore:

$$\bar{V} = (50 + 25)/2 = \$37.50$$

The result is that company 1 has, on average, profit of 37.50 - 50 = -12.50. That is, a loss of 12.50.

Company 1 now decides to change strategy and bid less than the valuation they receive. What if $v_1 = 50$ and company 1 bids \$37.50 instead. We will assume that company 2 continues to bid v_2 and company 1 wins.

If company 1 wins, $37.50 > v_2$.

On average, in this state of the world, company 2's signal is \$18.75 (due to the uniform distribution). The average value of the tract is therefore:

$$(50 + 18.75)/2 = $34.37$$

Company 1's profit is, on average, 34.37 - 37.50 = -3.13.

Company 1 now decides to bid only half the valuation. What if $v_1 =$ \$50 and company 1 bids \$25. We again assume company 2 bids v_2 and company 1 wins.

If company 1 wins, $25 > v_2$.

On average, in this state of the world company 2's signal is 12.50 (due to the uniform distribution), so the average value of the tract is (50+12.50)/2 = 31.25.

Company 1's profit is 31.25 - 25 = 6.25 on average. However, they will win only 25% of the time.

This analysis also has a complication in that it does not account for the fact that company 2 is also a strategic player. We assumed company 2 bids v_2 , but as for company 1, this strategy would lead to an expected loss for company 2.

So what does each firm do at equilibrium?

As each firm will have the same strategy at equilibrium, we can solve for company 1 assuming company 2 does the same strategy in response. At equilibrium, we can also assume that each company will have an expected profit of zero as each company would otherwise have an incentive to change their bid to gain a share of the positive profit.

Company 1 will win if $\delta v_1 > \delta v_2$; in other words, if $v_1 > v_2$. On average, in this state of the world company 2's signal is $0.5v_1$ (due to the uniform distribution), so that the average value of the tract is $(v_1 + 0.5v_1)/2 = 0.75v_1$.

Company 1's profit is:

$$\pi_1 = 0.75 v_1 - \delta v_1 = (0.75 - \delta) v_1$$

Profit is zero when $\delta = 0.75$. The Nash equilibrium is that both parties bid 75% of their private valuation.

In summary, bidding based purely on your own valuation fails to take into account that you only win if the other player's signal is low.

Alternatively, we may say that winning the auction is bad news regarding the value of the field. This is the winner's curse.

Because of the winner's curse, the Nash equilibrium is to bid more conservatively.

The mistake that oil companies make is ignoring or underestimating the winner's curse. If an oil company wins an auction, it's likely because its geologists have the highest estimates of the field's value. But if all other geologists have lower estimates of the value, the company's geologists have probably overestimated it.

Chapter 43

Strategic moves and commitment

Summary

- Strategic moves are actions that change the game being played, typically from a single-stage to a two-stage game. They come in two forms: unconditional (commitments) and conditional (threats and promises).
- For commitments to be effective, they must be observable and irreversible.
- Conditional strategic moves must be credible. If carrying out a threat or promise is too costly for the player making it, the threat or promise won't influence the other player's actions.
- The game of "chicken" illustrates how strategic moves can change outcomes. Originally, this game has two pure-strategy Nash equilibria: (Straight, Swerve) and (Swerve, Straight). However, if one player commits to going straight (e.g., by removing their steering wheel), it changes the Nash equilibrium to favour that player.

43.1 Chicken

Consider the following game of chicken. Two players are driving toward each other. Whoever swerves first loses. If neither swerves, they crash and die.

		Player B	
		Straight	Swerve
Player A	Straight	-100, -100	10, 0
	Swerve	0, 10	1, 1

There are two pure-strategy Nash equilibria: (Straight, Swerve) and (Swerve, Straight). If the other player swerves, they want to go straight. If the other player goes straight, they want to swerve.



Now consider a new scenario.

As they are driving toward each other, Player A rips the steering wheel out of their car and throws it out the window. They will now drive straight no matter what Player B does.



This is effectively a new game. What is the Nash equilibrium?

The Nash equilibrium is (Straight, Swerve). Player A wins the game of chicken.



43.2 Strategic moves

The option to commit to a course of action, as in this game of chicken, is an example of a strategic move.

A strategic move changes the game you are playing from a single-stage game to a two-stage game. In the first stage, you make your strategic move. In the second you play the original game.

Strategic moves come in two forms:

- Unconditional strategic moves, which we call commitments
- And conditional strategic moves, which we call threats and promises

43.2.1 Unconditional strategic moves

An unconditional strategic move is a commitment. For example, removing your steering wheel in chicken is a commitment.

The commitment needs to be observable and irreversible.

If your opponent cannot observe your commitment, you can claim to have made the commitment when you have not.

If your commitment is reversible, the game remains as if you had never made it.

43.2.2 Conditional strategic moves

A conditional strategic move involves specifying to your opponent how you will respond to each move.

A threat involves specifying negative consequences to the other player if they do not play as you wish. "If you don't clean your room, you won't get dessert."

A promise involves specifying positive consequences to the other player if they play as you wish. "If you clean your room, you can have dessert."

Threats and promises only achieve their objective if they are credible. That is, they only work if the other player believes they will be carried out as stated.

Sticking to a commitment and carrying out a threat or a promise typically reduces the possible actions of the player. If the proposer loses too much from carrying out a threat or promise, they will not carry it out.

43.2.3 Example: Complaining is costly

Here is an example.

You threaten to complain about poor service by a company. Complaining is costly.



We work through this problem by backward induction. At the final node for the customer, they can complain for a payoff of -1 or not complain for a payoff of 1. They will not complain.

The company, therefore, has a choice between providing good service for a payoff of 1 or bad service for a payoff of 2. They will provide bad service. The company has the same payoff for bad service regardless of the presence of the threat to complain as the threat is not credible.

For the customer's initial choice of whether to threaten to complain, it does not matter either way. Regardless of their threat, they receive bad service.

The result is two sub-game perfect Nash equilibria: (Threatens to complain, Bad service, Does not complain) and (No threat, Bad service).



Chapter 44

Game theory exercises

The cold war 44.1

The year is 1964 and the Soviet Union and the United States are in the midst of the cold war.

Suppose each player is considering whether they should act aggressively (hawk) or peacefully (dove). If one plays hawk while the other plays dove, they win the cold war. If both play hawk, there is a nuclear armageddon.

The payoffs (x, y) of each option for the Soviet Union and United States is as follows:

		United States	
		Hawk	Dove
Soviet Union	Hawk	-1000, -1000	10, 0
Soviet Union	Dove	0, 10	5, 5

I Inited States

a) What is the Nash equilibrium of this game? What other game does this resemble?



b) In the movie Dr Strangelove, the Soviet Union created a doomsday machine that would detonate automatically if there was a nuclear strike. The fallout would render the earth uninhabitable. The doomsday machine could not be deactivated and would explode if any attempt was made.

Explain how the doomsday device could act as a commitment device?



c) In the movie, the Soviet Union failed to inform the United States of the existence of the device. How could this failure to inform undermine its effectiveness as a commitment device?

Answer

A commitment will only be effective if it is both observable and irreversible. While the doomsday machine is irreversible, by not being observable it will not change the response of the United States. The United States will think they are playing the game analysed in part a), not that in part b).

44.2 Hiring

Robyn is hunting for a new employee. Robyn's company uses highly-technical equipment and needs to invest heavily in training the new employee. If the new employee leaves straight after training, Robyn's company will suffer a net loss from the employee. If the employee stays long-term, they will have a large gain.

Robyn approaches Sean and asks if he is interested in a long-term role with the company.

Sean is interested in the training as he could use it to boost his career, but sees less benefit in staying long-term. He considers whether he should say he is interested or not.

The extensive form of the game is laid out below, with the payoffs (x, y) being for Robyn and Sean respectively.



a) Will Robyn offer the position to Sean?

? Answer

We work through the problem by backward induction.

Sean can get 2 by leaving after training or 1 by staying. He leaves after training.

When considering whether he will state that he is interested, he could get 2 for stating he is interested (as he will later leave) versus nothing for saying he is not interested. He states he is interested.

Robyn compares the -1 she gets for hiring Sean (as he will leave) with the zero for no offer. She does not make an offer.



b) What sort of strategic move could help John? What could make the move credible?

Answer

One option is to sign a binding contract with penalties if he leaves early. Any penalty greater than -1 would make staying more attractive. A contract is both observable and irreversible (at least without mutual agreement).

44.3 Investment

Linda is looking for investment opportunities. She identifies a promising cryptobased start-up created by Marco. Marco is looking for seed funding.

Linda can invest \$10.

If Linda invests, her investment will triple in value. Marco can then decide to either shut down the start-up and keep the \$30 or maintain the start-up in the market and pay a \$15 dividend to each of Linda and himself.

If Linda does not invest, Linda keeps the \$10. The start-up gets \$0.



a) Draw the extensive form representation of the above sequential game.

b) What is the equilibrium of this game if Linda and Marco are purely self-interested?



44.4 War

Two city states, Atlantis and El Dorado, are divided by a body of water. In the middle is an island that both states claim sovereignty over.

To establish their claims, both states have built a bridge to the island. Atlantis then sent troops to the island.

El Dorado is deciding whether to attack Atlantis's troops to reclaim the island or to concede.

If El Dorado attacks, Atlantis need to decide whether to defend against the attack or to retreat back across the bridge.

If El Dorado attacks and Atlantis defends, both countries will suffer large losses.

These decisions and the payoffs (x, y) from each decision for El Dorado and Atlantis respectively are as follows.



a) What is the subgame-perfect equilibrium of this game?



b) An adviser to the Atlantis army suggests that they burn the bridge behind them to remove the option of retreat.

Draw the new extensive form game that would emerge if Atlantis had the option of burning the bridge. What is the subgame-perfect equilibrium?



If we work through this game by backward induction, starting with the upper branch:

- Atlantis would prefer to retreat (payoff of 2) compared to fighting (payoff of 1).
- El Dorado would prefer to attack (payoff of 3) compared to conceding (payoff of 2).

For the lower branch:

• El Dorado would prefer to concede (payoff of 2) compated to attack (payoff of 1).

For Atlantis's final decision, they would prefer to burn (payoff of 3) compared to not burning (payoff of 2).



The subgame-perfect equilibrium is (Burn, Retreat; Attack, Concede).

44.5 Buying a car

Hayley wants to buy a car. The used-car salesman can sell her a good car (for which he earns a small profit) or a lemon (for which he earns a large profit).

The payoffs (x, y) for each decision are indicated in the game tree below, with x being the Hayley's satisfaction and y being the salesman's profit.



a) Assume the salesman only cares about his profit. What would the salesman do if Hayley chooses to purchase? Why?

Answer

The salesman will compare payoffs of 8 for selling a lemon and 5 for selling a good car. He will choose to sell a lemon.

b) Given the anticipated choice of the salesman, would Hayley purchase the car? Why?

? Answer

Hayley will compare a payoff of 0 for no purchase and a payoff of -1 for buying a car that will be a lemon. She chooses not to purchase.

c) Suppose Hayley can take legal action if she is sold a lemon. If Hayley is successful in court, she will be refunded the purchase price but would suffer a cost of -5 due to the effort involved. Would this change the outcome? Why?



Working by backward induction, Hayley has a choice between taking legal action for a payoff of -5 or not complaining for a payoff of -1. She will not complain.

The rest of the game plays out as per questions a) and b). There is no change to the outcome as she cannot commit to complain in advance. The threat to take legal action is not credible.

d) Suppose Hayley has a reputation of being quick to anger and always carrying out her threats. Suppose Hayley would experience satisfaction of +6 from taking legal action (in addition to the effort cost of -5). The salesman knows this and believes it to be a credible commitment. Would this change the outcome? Why?



The salesman now has a choice between a payoff of 0 for selling the lemon (as Hayley takes legal action and the sale is refunded) and 5 for selling a good car. He sells the good car

Hayley now has a choice of a payoff of 0 for not purchasing a car, and 5 for purchasing. She makes the purchase and gets a good car.

Part VIII

Behavioural game theory

In our analysis of game theory, I assumed rational agents in that they use all available information and can successfully determine their best action given their opponent's (also rational) action.

But what if agents have limited rationality or vary in their rationality?

In this part, I will examine several departures from rationality.

The first is level-k thinking, in which the agents are assumed to have a certain level of reasoning. For example, a level-0 agent would choose an action randomly. A level-1 agent would assume that the opponent is level-0 and choose the best response to that. A level-2 agent would assume that the opponent is level-1 and choose the best response to that. And so on. The players try to be one step ahead of their opponents.

The second departure involves the degree to which the players account for asymmetric information. We consider what happens if players act as though everyone has the same information or if they fail to appreciate the informational advantage they have relative to less-informed agents.

The third departure involves emotions. We consider the role of emotions in enabling players to commit to courses of action that they otherwise could not credibly stick with.

Chapter 45

Level-k thinking

Summary

- Level-k thinking is a model of strategic reasoning where players try to be "one step ahead" of others. A level-k player assumes others are level-(k-1) and responds optimally to that assumption.
- Level-0 players are non-strategic, often modelled as randomizing across all strategies. Level-1 players assume others are level-0, level-2 assumes others are level-1, and so on.
- The p-beauty contest game illustrates level-k thinking. In this game, players choose numbers, aiming to be closest to (p) times the average of all choices. Different levels of thinking lead to different number choices, with experimental results often showing a mix of level-1, level-2, and level-3 players.

45.1 Introduction

The idea behind level-k thinking is that a player forms an expectation of what others will do and tries to be "one step ahead".

That is, a level-k player plays the best response to level-(k-1) players.

Level-0 players do not engage in strategic thinking. This is usually modelled as randomisation across all strategies.

Level-1 players assume other players are level-0 and act optimally conditional on this assumption.

Level-2 players assume other players are level-1 and act optimally conditional on this assumption.

And so on.

45.2 Examples

45.2.1 The beauty contest

To understand level-k thinking, consider the following thought experiment from Keynes (1936).

[P]rofessional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.

This thought experiment has since been developed into a game, the p-beauty contest (Moulin (1986)).

In the p-beauty contest, each of n players pick a number $y \in [0, 100]$.

The winner is the player whose chosen number is closest to the mean of all the chosen numbers (\bar{y}) multiplied by a parameter p. That is, the winner is the player with their chosen number closest to $p\bar{y}$.

p is typically chosen such $0 \le p \le 1$, with p = 1/2 and p = 2/3 common.

How might level-k players play this game?

Suppose p = 2/3.

A level-0 player does not think strategically. We will have the level-0 player randomly select a number between 0 and 100.

The level-1 player will play the best response to level-0 players. If level-0 players select across the interval [0, ..., 100], the best response is:

$$y_1 = \frac{2}{3}\bar{y} = \frac{2}{3} \times 50 = 33.3$$

The level-2 player will play the best response to level-1 players. If all other players are level-1 and select 33.3, the best response is:

$$y_2 = \frac{2}{3}\bar{y} = 2/3 \times 33.3 = 22.2$$

The level-3 player will play the best response to level-2 players. If all other players are level-2 and select 22.2, the best response is:

$$y_3 = \frac{2}{3}\bar{y} = 2/3 \times 22.2 = 14.8$$

And so on.

The following charts come from Nagel (1995), with p = 1/2 and p = 2/3. The charts show the distribution of chosen numbers in the p-beauty contest.

The chart with p = 2/3 spikes at 22.2 and 33.3, suggesting players are playing at level-2 and level-1, respectively. This matches with other experimental evidence on the p-beauty game, with few level-0 players. Most are level-1, level-2 and level-3.



Figure 45.1: Choices by players of the p-beauty game

The lab evidence doesn't necessarily imply that level-k is the "right model". Data and theory appear to match, but it is hard to know whether this is how subjects are thinking.

Finally, it is worth noting that in the Nash equilibrium, each player picks 0. This is because the best response to all other players picking 0 is to pick 0. For any higher number, everyone has an incentive to lower their choice. However, if playing against level-k players, selecting 0 is not the best approach.

45.2.2 The assignment game with level-k thinking

Let's consider another example of level-k thinking involving a game called the assignment game.

Each player needs to decide if they will work or shirk. If they both work, they receive a good payoff. They receive an ever better payoff, however, if they shirk while the other works.

		Play	Player B	
		Work	Shirk	
Player A -	Work	7,4	1,9	
	Shirk	9,-1	0,0	

Working through the payoffs for each player, if player B works, player A is better off shirking, receiving payoff of 9. If player B shirks, player 1 is better off working, receiving payoff of 1. If player A works, player 2 is better off shirking, receiving payoff of 9. If player A shirks, player B is still better off shirking, receiving payoff of 0.

There is a unique Nash equilibrium (work, shirk), with shirk the dominant strategy for Player B.

		Player B	
		Work	Shirk
Discon	Work	7,4	1,9
Flayer A	Shirk	9,1	0,0

Consider, however, if instead of fully rational agents, we have level-k thinkers playing the game.

In this case, the outcome of the game will depend on the level of thinking of each player.

If both players are level-0, they will each play randomly.

At level-1, each player will play the best response to level-0 players. Each player determines this by calculating their best response to the random strategy of the other player.

For player A, their expected payoffs are calculated using the 50% probability with which player B could play each action.

The expected payoff from playing work is:

$$\frac{1}{2} \times 7 + \frac{1}{2} \times 1 = 4$$

The expected payoff from playing shirk is:

$$\frac{1}{2} \times 9 + \frac{1}{2} \times 0 = 4.5$$

A level-1 player A chooses to shirk.

For player B, their expected payoff from playing work is:

$$\frac{1}{2}\times 4+\frac{1}{2}\times -1=1.5$$

Their expected payoff from playing shirk is:

$$\frac{1}{2}\times9+\frac{1}{2}\times0=4.5$$

A level-1 player B also chooses to shirk.

If a player has a dominant strategy, they discover it at k = 1. Any level-k thinker will always uses the dominant strategy for $k \ge 1$. In that case, we know that any player B with $k \ge 1$ will shirk.

What if each player is level 2?

Player A calculates their best response to a level-1 player B. A level-1 player B always plays shirk. Player A's best response to shirk is to work. The level-2 player A works.

Although we know a level-2 player B will shirk as as shirk is their dominant strategy, we can show this by considering their best response to a level-1 player A. A level-1 player A always plays shirk. Player B's best response to shirk is to shirk. The level-2 player B shirks.

At a certain level of thinking, the players will discover the Nash equilibrium. Here, they have discovered it at level-2 thinking. For any higher level of thinking, they will remain at the Nash equilibrium. That is, if players endowed with level $k = \bar{k}$ rationality play Nash, all players with $k > \bar{k}$ play Nash.

Level-k	Player A	Player B
k = 0	Random	Random
k = 1	Shirk	Shirk
k = 2	Work	Shirk
k = 3	Work	Shirk
k = 4	Work	Shirk

45.2.3 Centipede game

Another example of level-k thinking is the centipede game.

This centipede game has six stages. At each stage, a player can "take" and end the game or they can "pass", increasing the total payoff. The other player then has a move.

The numbers 1 and 2 along the top of the centipede represent the decision nodes for two players. At the first node, player 1 has the choice to take or pass. If player 1 passes, player 2 has the choice to take or pass, and so on. At the final node, the game ends regardless of what player 2 chooses.



The payoff when a player takes and ends the game is represented by the numbers in the brackets. The first number is the payoff for player A and the second number is the payoff for player B.

There is a unique subgame perfect equilibrium for the centipede game: $S_1 = (take, take, take)$ and $S_2 = (take, take, take)$, where S_1 and S_2 are the set of strategies for player A and player B respectively. We solve for this in Section 41.3.1.

What do people do when playing the centipede game in the lab?

People tend to pass until a few stages before the end (depending on the length of the centipede) and then take. They do not play the Nash equilibrium strategy.

Can level-k thinking provide insight into this behaviour?

Suppose a level-0 player passes until the end. They are possibly lucky if they are player A playing against another level-0 player.

A level-1 player B would take at (4,6) as the level-0 player A would pass until then. A Level-1 player A would be planning to take (6,5) at the end as they believe the level-0 player B will keep passing.

A level-2 player B would plan to take at the final stage (4,6) as they believe the level-1 player A passes. A level-2 player A would take the payoff at (5,3) as they believe a level-1 player B would take at (4,6). A level-3 player B would plan to take at (2,4) as they believe the level-1 player A will take at (5,3). A level-3 player A would plan to take at (5,3) as they believe a level-2 player B would take at (4,6).

And so on.

Chapter 46

Asymmetric information and the curse of knowledge

Summary

- Information asymmetry can cause market failures even with rational agents, but there's evidence that people often don't behave optimally in these situations.
- The "curse of knowledge" occurs when better-informed agents fail to ignore their private information when predicting less-informed agents' actions, even when it's in their interest to do so.
- Conversely, less-informed agents often fail to account for others' informational advantages. This is demonstrated in the "market for lemons" example, where a "cursed" buyer might overpay for a car, not realizing only lemons would be sold at their offered price.
- Similarly, in the "winner's curse" scenario, a "cursed" bidder in an oil field auction might overbid by underestimating the significance of winning, failing to appreciate that winning implies the other bidder's information was less favourable.

46.1 Introduction

We saw earlier in our examination of the market for lemons and the winner's curse that asymmetric information can cause market failures even if agents are fully rational. However, the rational agents account for the information and behaviour of others and as a result, behave optimally despite that market imperfection.

There is substantial empirical evidence that people do not behave in this way.

For example, people tend to underestimate the extent to which informational differences drive others' behaviour. They often act as if others have the same information set that they do. Where an agent has information that another doesn't, this phenomenon is known as the curse of knowledge.

Further, better-informed agents often fail to take advantage of their informational advantage against less-informed agents because they don't understand the link between information and behaviour.

46.2 The curse of knowledge

The idea behind the curse of knowledge is that better-informed agents should ignore the additional information they hold when predicting the actions of lessinformed agents. Experimental evidence shows that people are unable to ignore their private information even when it is in their interests to do so.

For example, Newton (1990) had students participate in an experiment in one of two roles: "Tapper" and "Listener".

Tappers received a list of 25 well-known songs and were asked to "tap out" the rhythm of one of the songs.

Listeners tried to identify the song based solely on the taps.

Tappers predicted that listeners would identify 50% of the songs.

Listeners only identified 3 of 120 songs correctly (a rate of about 2.5%).

46.3 The market for lemons

While that experiment involved agents who had more information than the other players - they knew the song - we also see failures where the other player has additional information but the agent does not account for that fact.

We can explore this idea in the market for lemons.

Recall our earlier example in Section 42.2 involving the purchase of a used car.

There are two types of cars, good cars and lemons, and only the seller knows the type. The buyer knows that the seller has this information.

A car is good with probability q and a lemon with probability 1 - q. To the seller, good cars are worth \$10,000 and lemons \$5,000. To potential buyers, good cars are worth \$15,000 and lemons \$7,500.

A "cursed" buyer doesn't think that the seller's decision whether to trade depends on the seller's knowledge of the car.

Suppose that q = 0.2. That is, only 20% of the cars are good.

Suppose the cursed buyer believes that cars are sold with equal probability regardless of type.

In that case, the expected value of a car to a buyer is:

$$\hat{\mathbf{E}} = 0.2 \times 15000 + 0.8 \times 7500$$

= 9000

A buyer would be willing to pay up to \$9000 for a car.

At that price, however, the seller of a good car would not be willing to sell. The market will comprise only lemons, which sellers are more than happy to sell. The buyer will pay \$9000 for a car worth only \$7500 to them.

46.4 The winner's curse

We can also explore this phenomenon in the winner's curse. Recall our example of the winner's curse in Section 42.3 on bidding for an oil field.

Company 1 and company 2 hire a geologist to estimate the value of an oil field. The honest geologist of each company privately reports their estimated valuation to the company. Company 1 learns v_1 and company 2 learns v_2 .

 v_1 and v_2 are uniformly distributed between 0 and 100 and independent.

Assume the true value of the oil field is the mean of v_1 and v_2 :

$$V = \frac{v_1 + v_2}{2}$$

The two companies simultaneously bid for the field in a first-price auction. The highest bid wins and pays their bid.

Assume company 1 is cursed and therefore assumes that company 2's bid is independent of v_2 . Company 1 assumes v_2 is on average \$50 and that company 2 always bids.

Company 1's expected profit, if they bid v_1 , is:

$$\begin{split} \hat{\mathbf{E}}[\pi_1|\text{bid } v_1] &= \frac{1}{2}\pi_1(\text{lose}) + \frac{1}{2}\pi_1(\text{win}) \\ &= \frac{1}{2} \times 0 + \frac{1}{2} \bigg(\frac{v_1 + 50}{2} - v_1 \bigg) \\ &= \frac{1}{4} (50 - v_1) \end{split}$$

We can see that:

$$\hat{\mathbf{E}}[\pi_1|\mathrm{bid} v_1]>0 \Leftrightarrow v_1>50$$

That is, company 1 expects to make a profit if they receive a private valuation of more than \$50.

However, as shown in Section 42.3, this bidding approach leads to, on average, a loss. Company 1 under-appreciates that company 2 is more likely not to bid when company 2's information is bad. Therefore, company 1 under-appreciates the extent to which winning the auction is bad news.

Chapter 47

Emotions

Summary

- Emotions are mental states that signal positive or negative outcomes and can function as commitment devices in strategic interactions.
- Emotions like guilt and anger can serve as credible commitments, enabling trust and cooperation or deterring cheating, even when the emotional response appears "irrational" in the short term.
- In everyday scenarios, like complaining about bad service or playing "chicken," emotions can change the game's payoff structure and equilibrium, making threats credible and influencing others' behaviour to one's advantage.

47.1 Introduction

Emotions are mental states that signal positive or negative outcomes.

One function of emotions may be to act as a commitment device:

- The emotion of guilt can constrain a desire to "cheat" where cheating delivers a higher pay-off. This in turn may allow people to trust you.
- The emotion of anger may lead you to punish someone even where delivering the punishment also harms you. This in turn may lead people to be less likely to cheat you.

While this behaviour may appear "irrational", it allows people to make credible commitments that in turn allow them to enter beneficial trades and cooperative arrangements, while being less likely to being cheated. Consider the following quote from Richard Nixon:

I call it the Madman Theory, Bob. I want the North Vietnamese to believe I've reached the point where I might do anything to stop the war. We'll just slip the word to them that, "for God's sake, you know Nixon is obsessed about communism. We can't restrain him when he's angry—and he has his hand on the nuclear button" and Ho Chi Minh himself will be in Paris in two days begging for peace.

Pushing the nuclear button is not in Nixon's interest, and from a purely rational perspective may not be a credible threat. But if a madman has his finger on the button, the calculation changes.

47.2 Examples

47.2.1 Complaining for bad service

Recall our earlier example in Section 43.2.3, illustrated below, of a customer threatening to complain if they receive bad service. Complaining is costly.

We determined this by backward induction. At the final node for the customer, they can complain for a payoff of -1 or not complain for a payoff of 1. They will not complain.

The company, therefore, has a choice between providing good service for a payoff of 1 or bad service for a payoff of 2. They will provide bad service. The company has the same payoff for bad service regardless of the presence of the threat to complain as the threat is not credible.

For the customer's initial choice of whether to threaten to complain, it does not matter either way. Regardless of their threat, they receive bad service.



But what if the customer gets a strong sense of satisfaction from complaining worth +3? Then their payoffs become as follows:



The threat to complain is now credible. If they receive bad service, they complain for a payoff of 2 rather than not complain for a payoff of 1.

The company now provides good service following a threat to complain. Absent that threat, they would provide bad service.



47.2.2 Chicken

As another example, recall the game of chicken. Two players are driving toward each other. Whoever swerves first loses. If neither swerves, they crash and die.

		Player B	
		Straight	Swerve
Diawar A	Straight	-100, -100	10, 0
Player A	Swerve	0, 10	1, 1

There are two pure-strategy Nash equilibria: (Straight, Swerve) and (Swerve, Straight). If the other player swerves, they want to go straight. If the other player goes straight, they want to swerve.

		Player B	
		Straight	Swerve
Player A	Straight	-100,-100	100
Player A	Swerve	0,10	1, 1

Now suppose player A is crazy. They are afraid of nothing and will never swerve. Player B knows this.

Player A's craziness acts as a commitment device similar to that of removing the Steering Wheel. If player A will not swerve, player B will.

		Player B	
		Straight Swerve	
Player A	Straight	-100, -100	10, 0

The Nash equilibrium is (Straight, Swerve). The crazy player A wins the game of chicken.

		Player B	
		Straight	Swerve
Player A	Straight	-100,-100	100
Chapter 48

Behavioural game theory exercises

48.1 Penalty kick

A soccer player (the striker) has a penalty kick. The striker is deciding whether to kick to the left or right. If the goalkeeper dives in the correct direction, the goalkeeper will stop the ball and the two sides will tie. Otherwise, the striker will score a goal and win.

Lately, the striker has been having trouble kicking to the right, sometimes missing the goals even when the goalkeeper doesn't dive in that direction.

The expected payoffs for each combination of actions are as follows, with the payoff (x, y) being the payoffs for the striker and goalkeeper respectively:

		Goalkeeper		
		Right	Left	
Striker	Right	0, 5	8, 2	
	Left	10, 0	0, 5	

Are there any pure-strategy Nash equilibria? If so, what are they?

There are no pure-strategy Nash equilibria. Whatever the striker does, the goalkeeper wants to match. If the goalkeeper matches, the striker wants to change.

a) Suppose the striker and goalkeeper are level-k thinkers.

If they were level-0, both would choose right or left with equal probability.

What would each player do if they were a level-1 thinker? Explain.

? Answer

A level-1 striker will assume they are playing a level-0 goalkeeper. They will estimate the the payoff from each action responding to the random play of a level-0 goalkeeper.

$$\begin{split} E[U_S(R)] &= 0.5 \times 0 + 0.5 \times 8 \\ &= 4 \end{split}$$

$$E[U_S(L)] = 0.5 \times 10 + 0.5 * \times 0$$

= 5

The level-1 striker has a higher expected payoff for kicking left, so kick left.

A level-1 goalkeeper will assume they are playing a level-0 striker. They will estimate the the payoff from each action responding to the random play of a level-0 striker.

$$\begin{split} E[U_G(R)] &= 0.5 \times 5 + 0.5 \times 0 \\ &= 2.5 \end{split}$$

$$\begin{split} E[U_G(L)] &= 0.5 \times 2 + 0.5 * \times 5 \\ &= 3.5 \end{split}$$

The level-1 goal keeper has a higher expected payoff for going left, so go left.

b) What would each player do if they were a level-2 thinker? Explain.

? Answer

A level-2 striker will assume they are playing a level-1 goalkeeper. They believe the level-1 goalkeeper will go left, so they will go right (payoff 8 compared to payoff 0).

A level-2 goalkeeper will assume they are playing a level-1 striker.

They believe the level-1 striker will go left, so they will go left (payoff 5

48.2 Hide and seek

In the hide-and-seek game, the Hider selects one of the four boxes marked A, B, A and A. The Seeker guesses the box selected by the hider.

The Seeker wins if they find the Hider. Otherwise, the Hider wins.



The payoffs are as follows. I have labelled the end boxes A1 and A2 to distinguish the "A"s from each other.

		Seeker			
		A1	В	Α	A2
Hider	A1	0,1	1,0	1,0	1,0
	В	1,0	0,1	1,0	1,0
	Α	1,0	1,0	0,1	1,0
	A2	1,0	1,0	1,0	0,1

Assume a level-0 seeker or hider selects a box by hiding in or looking in the "most salient" hiding spots. They choose A1 or A2 on the ends with p = 0.3 each, or B (because it is different) with p = 0.35. They hide or look in less salient middle A with probability $1 - 2 \times 0.3 - 0.35 = 0.05$.

(Note that this assumption for the level-0 agents is different to what we have assumed to date. We have typically assumed a level-0 agent randomly chooses an action.)

a) What box do the level-1 hider and seeker choose?

Answer

The level-1 hider calculates the expected payoff from hiding in each of the boxes if playing against a level-0 seeker.

 $E[U(A1)] = 0 \times 0.3 + 1 \times 0.35 + 1 \times 0.05 + 1 \times 0.3 = 0.7$ $E[U(B)] = 1 \times 0.3 + 0 \times 0.35 + 1 \times 0.05 + 1 \times 0.3 = 0.65$ $E[U(A)] = 1 \times 0.3 + 1 \times 0.35 + 0 \times 0.05 + 1 \times 0.3 = 0.95$ $E[U(A2)] = 1 \times 0.3 + 1 \times 0.35 + 1 \times 0.05 + 0 \times 0.3 = 0.7$

The level-1 hider hides in the least salient box A. The level-1 seeker calculates the expected payoff from looking in each of the boxes if playing against a level-0 hider.

 $E[U(A1)] = 1 \times 0.3 + 0 \times 0.35 + 0 \times 0.05 + 0 \times 0.3 = 0.3$ $E[U(B)] = 0 \times 0.3 + 1 \times 0.35 + 0 \times 0.05 + 0 \times 0.3 = 0.35$ $E[U(A)] = 0 \times 0.3 + 0 \times 0.35 + 1 \times 0.05 + 0 \times 0.3 = 0.05$

$$E[U(A2)] = 0 \times 0.3 + 0 \times 0.35 + 0 \times 0.05 + 1 \times 0.3 = 0.3$$

The level-1 seeker looks in Box B. If both the hider and seeker are level-1, the hider wins.

b) What box do the level-2 hider and seeker choose?

? Answer

The level-2 hider knows that the level-1 seeker chooses B. They select any box apart from B with equal probability, all of which they believe will give a pay-off of 1.

The level-2 seeker knows that the level-1 hider will select A. They select A.

The level-2 seeker wins with probability 1/3.

c) What box do the level-3 hider and seeker choose?

? Answer

The level-3 hider knows that the level-2 seeker chooses A. They select any box apart from A with equal probability, all of which they believe will give a pay-off of 1.

The level-3 seeker knows that the level-2 hider will select any box except B with equal probability. They select one of A1, A2 or A with equal probability.

The level-3 seeker wins with probability $1/3 \times 1/3 + 1/3 \times 1/3 = 2/9$.

d) What box do the level-4 hider and seeker choose?

? Answer

The level-4 hider knows that the level-3 seeker selected A1, A2 and A with equal probability. They select B.

The level-4 seeker knows that the level-3 hider selected any box apart from A with equal probability. They also select those boxes with equal probability.

The level-4 seeker wins with probability 1/3.

48.3 Matching pennies (with a twist)

Consider the following two-player game:

		В		
		Х	Y	
A -	х	6, 6.1	0, 0	
	Y	0, 0	6.1, 6	

a) What are the two pure-strategy Nash equilibria of this game?

? Answer

The two pure-strategy Nash equilibria of this game are (X,X) and (Y,Y). That is, if the players are jointly playing either of those combinations of strategies, neither has an incentive to deviate. Their response is a best response to the other players' actions.

b) Suppose players in this game think according to the level-k model. Assume a level-0 agent randomises between options with equal probability.

What would player A and player B do if they were level-1 players?

? Answer

Remember the idea behind level-k thinking: given their own cognitive level, a player forms an expectation of what others will do and tries to be "one step ahead of them".

We work out the utility of each option.

First, for player A:

$$EU_A^1(X) = 0.5 \times 6 + 0.5 \times 0 = 3$$

$$EU^1_A(Y) = 0.5 \times 0 + 0.5 \times 6.1 = 3.05$$

Player A chooses Y if they are a level-1 player. Next, for player B:

$$EU_B^1(X) = 0.5 \times 6.1 + 0.5 \times 0 = 3.05$$

$$EU_B^1(Y) = 0.5 \times 0 + 0.5 \times 6 = 3$$

Player B chooses X if they are a level-1 player. If both players are level-1, they will fail to coordinate.

c) What would player A and player B do if they were level-2 players?

? Answer

We again work out the utility of each option: First, for player A. They know that a level-1 player B will select X. Accordingly:

$$EU_{4}^{2}(X) = 1 \times 6 + 0 \times 0 = 6$$

$$EU_{A}^{2}(Y) = 1 \times 0 + 0 \times 6.1 = 0$$

Player A chooses X if they are a level-2 player. Next, for player B. They know that a level-1 player A will select Y. Accordingly:

$$EU_B^2(X) = 0 \times 6.1 + 1 \times 0 = 0$$

$$EU_B^2(Y) = 0 \times 0 + 1 \times 6 = 6$$

Player B chooses Y if they are a level-2 player. If both players are level-2, they will fail to coordinate.

d) What would player A and player B do if they were level-3 players?

? Answer

We again work out the utility of each option: First, for player A. They know that a level-2 player B will select Y. Accordingly:

 $EU^3_A(X)=0\times 6+1\times 0=0$

$$EU_A^3(Y) = 0 \times 0 + 1 \times 6.1 = 6.1$$

Player A chooses Y if they are a level-3 player. Next, for player B. They know that a level-2 player A will select X. Accordingly:

$$EU_B^3(X) = 1 \times 6.1 + 0 \times 0 = 6.1$$

 $EU_B^3(Y) = 1 \times 0 + 0 \times 6 = 0$

Player B chooses X if they are a level-3 player. If both players are level-3, they will fail to coordinate.

Of interest, each player tries to coordinate on the strategy that the other player would be better off coordinating on. That is, Player A receives 6 from successful coordination choosing X, which is less than the 6.1 Player A would get from coordinating on Y.

Given your answers to b) through d), what mix of level-k players might explain the mis-coordination described above?

e) When this game is played in the laboratory, the players mis-coordinate. About 3/4 of the row players (player A) choose X while about 3/4 of the column players (player B) choose Y.

? Answer

That 3/4 of Player "A"s choose X and 3/4 of Player "B"s choose Y suggests there are many level-2 players (or possibly level-4). They each assume that the other player is level-1 and has picked the option with the highest payoff for themselves. They are effectively trying to coordinate with the other player by assuming that the other will seek their highest paying option. However, if both do this, both receive nothing.

48.4 Buying a car

Suppose that you are considering purchasing a car.

You believe that the seller values it between \$1000 and \$5000, with an equal probability that it has a value at any point in this range. That is, you believe it is uniformly valued to the seller between \$1000 and \$5000.

The seller knows the car and its true value.

Assume that whatever the car is worth to the seller, it is worth 1.33 times that to you (so a car worth \$2400 to the owner is actually worth \$3200 to you).

a) What offer should you make to ensure that you will not lose money?

? Answer

Suppose you offer x. If the seller accepts, the value must be between \$1000 and x.

As the value evenly distributed across that interval, its average value would be:

$$1000 + \frac{x - 1000}{2} = 500 + \frac{x}{2}$$

The expected value of the car to you will be:

$$\frac{4}{3}\left(500 + \frac{x}{2}\right)$$

To ensure you don't lose you want:

$$\frac{4}{3}\left(500 + \frac{x}{2}\right) > x$$

Solving this out, you expect to make a profit where x <\$2000.

b) Suppose you are cursed player and you believe sellers will take the average optimal action of selling whenever they are offered more than \$3000. As a result,

you decide to offer \$3000. What is your accepted profit if the seller accepts your offer?

? Answer

If seller accepts, the value must be between \$1000 and \$3000. If value evenly distributed across that interval, its average value would be \$2000.

Given it is worth 1.33 times more to you, it would be worth \$2,667 on average.

You would lose, on average, \$333.

48.5 Advice

Agent A is going to their financial adviser to buy some life insurance. The adviser can sell them insurance that does not cover heart attacks but for which the adviser receives a huge sales commission (bad insurance). Or the adviser can sell Agent A comprehensive insurance for which their sales commission is lower (good insurance).

The payoffs (x, y) for each decision are indicated in the game tree below, with x being the satisfaction of Agent A and y being the satisfaction of the adviser.



a) Assume the adviser only cares about the payoffs indicated. What would the adviser do if Agent A chooses to purchase?

Answer

The adviser will compare payoffs of 4 for selling bad insurance and 2 for selling good insurance. They will choose to sell bad insurance.

b) What would Agent A do, anticipating the choice of the adviser?

💡 Answer

Agent A will compare a payoff of 0 for no purchase and a payoff of -2 for purchase (knowing that they will be sold bad insurance). They will choose not to buy insurance.

c) Suppose now that Agent A can complain to the regulator if they are sold bad insurance. If Agent A is successful, they can cancel the insurance but would suffer a cost of -3 due to the effort involved. Would this change the outcome of the game?



Working by backward induction: Agent A has a choice between complaining for a payoff of -3 or not complaining for a payoff of -2. They do not complain.

The rest of the game plays out as per questions a) and b). There is no change to the outcome as they cannot commit to complain in advance (at least in this version of the game). The threat to complain is not credible.

d) Suppose Agent A has a reputation for seeking revenge and would experience satisfaction of +4 from complaining to the regulator (in addition to the effort cost of -3) as reciprocation for the action of the adviser. How would this change the outcome of the game?



(as Agent A complains and the insurance is cancelled) and 2 for selling good insurance. They sell the good insurance.

Agent A now has a choice of a payoff of 0 for not purchasing insurance and 2 for purchasing. They make the purchase.

Part IX

Social preferences

Summary

- Social preferences, or other-regarding preferences, encompass three main types: distribution, reputation, and reciprocity.
- Distribution preferences relate to how people care about resource division, driven by altruism or inequality aversion.
- Reputation preferences involve concern for what others think and fear of social stigma from selfish behaviour.
- Reciprocity preferences reflect how people care about others' intentions and respond in kind to their actions.
- The ultimatum and dictator games demonstrates social preferences, as dictators and proposers often offer more than the minimum and responders sometimes reject non-zero offers.

Introduction

People do not care solely about their own outcomes. They care about the outcomes and actions of others. These preferences are known as social preferences, or sometimes "other-regarding preferences".

In this part, I will examine the three types of social preferences: distribution, reputation, and reciprocity.

Distribution refers to how people care about the division of resources. This can be driven by either altruism, which is the desire to help others, or inequality aversion, which is concerned with the fairness of the distribution and the relative gaps between individuals.

Reputation relates to how people care about what other people think. People fear the social stigma that can result from "selfish" behaviour.

Reciprocity relates to how people care about the intentions of others and how they often respond in kind to their actions.

Two examples

The results of the following games are evidence of social preferences.

The ultimatum game

Recall our earlier discussion of the ultimatum game.

The ultimatum game involves two players: the proposer and the responder.

The proposer is given a fixed amount of money m. They then offer a portion x of the sum m to the responder.

The responder can either accept or reject the offer. They make this decision knowing the fixed amount m held by the proposer and the offer x.

If the responder accepts, the responder receives the offer x and the proposer gets the remainder m-x. If the responder rejects, both players receive nothing.



Generally, if the players have monotonic preferences and the offer strategy set is discrete:

- The responder accepts any x > 0.
- The proposer offers the smallest non-zero amount the proposer can offer.

The other (weak) subgame perfect Nash equilibrium is an offer of \$0 and acceptance.

What do people do in the ultimatum game?

Unlike the game theoretic predictions, proposers rarely offer the minimum amount, and responders often reject non-zero offers.

For example, Henrich et al. (2001) recruited subjects from 15 small-scale societies to play the ultimatum game. The mean offer in all societies was substantially above zero. The rejection rate was low but non-zero.

Group	Country	Mean offer ^a	Modes ^b	Rejection	Low- offer rejection
Group	Country	oner	Modes	Tate	Tate
Machiguenga	Peru	0.26	0.15/0.25	0.048	0.10
			(72)	(1/21)	(1/10)
Hadza	Tanzania	0.40	0.50	0.19	0.80
(big camp)			(28)	(5/26)	(4/5)
Hadza	Tanzania	0.27	0.20	0.28	0.31
(small camp)		(38)	(8/29)	(5/16)	
Tsimané	Bolivia	0.37	0.5/0.3/0.25	0.00	0.00
			(65)	(0/70)	(0/5)
Quichua	Ecuador	0.27	0.25	0.15	0.50
			(47)	(2/13)	(1/2)
Torguud	Mongolia	0.35	0.25	0.05	0.00
			(30)	(1/20)	(0/1)
Khazax	Mongolia	0.36	0.25		
Mapuche	Chile	0.34	0.50/0.33	0.067	0.2
S			(46)	(2/30)	(2/10)
Au	PNG	0.43	0.3	0.27	1.00
			(33)	(8/30)	(1/1)
Gnau	PNG	0.38	0.4	0.4	0.50
			(32)	(10/25)	(3/6)
Sangu	Tanzania	0.41	0.50	0.25	1.00
farmers			(35)	(5/20)	(1/1)
Sangu	Tanzania	0.42	0.50	0.05	1.00
herders			(40)	(1/20)	(1/1)
Unresettled	Zimbabwe	0.41	0.50	0.1	0.33
villagers			(56)	(3/31)	(2/5)
Resettled	Zimbabwe	0.45	0.50	0.07	0.57
villagers			(70)	(12/86)	(4/7)
Achuar	Ecuador	0.42	0.50	0.00	0.00
			(36)	(0/16)	(0/1)
Orma	Kenya	0.44	0.50	0.04	0.00
			(54)	(2/56)	(0/0)
Aché	Paraguay	0.51	0.50/0.40	0.00	0.00
			(75)	(0/51)	(0/8)
Lamelarae	Indonesia	0.58	0.50	0.00	0.00
			(63)	(3/8)	(4/20)

TABLE 1-THE ULTIMATUM GAME IN SMALL-SCALE SOCIETIES

_

Note: PNG = Papua New Guinea. ^a This column shows the mean offer (as a proportion) in the ultimatum game for each society. ^b This column shows the modal offer(s), with the percentage of subjects who make modal offers (in parentheses). ^c The rejection rate (as a proportion), with the actual numbers given in merenthese

^d The rejection rate for offers of 20 percent or less, with the actual numbers given in parentheses. ^e Includes experimenter-generated low offers.

These results cannot be explained by examining only the outcomes to the individual. We need to consider their social preferences.

The dictator game

Recall our earlier discussion of the dictator game.

In the dictator game, the dictator is given a fixed amount of money m. They then offer a portion x of the sum m to the receiver. The game then ends.

Exchange is unilateral. Receivers have an empty strategy set.



The standard game theory prediction is that the dictator offers nothing. The dictator maximises their payoff by keeping all of the endowment themselves, receiving payoff m.

However, in experiments, dictators tend to give a positive sum of money. The following shows distributions reported by Engel (2011). Most players offer more than zero, suggesting preferences beyond simply maximising their own payoff.



Figure 48.1: Distribution of amount given in the dictator game (Engel, 2011).

Chapter 49

People care about resource distribution

Summary

- Distributional preferences relate to how people value the relative distribution of resources among people.
- Distributional preferences can be incorporated into economic analysis by including others' outcomes in utility functions.
- Altruism is a type of distributional preference where an individual places a positive weight on others' outcomes. It can be pure (genuine concern for others) or impure (deriving satisfaction from doing good).
- Inequality aversion is a type of distributional preference where people dislike having less than others and may also dislike having more than others.
- We can also model other distributional preferences, including statusseeking, Rawlsian and utilitarian preferences.

49.1 Introduction

In 2008, when the Sacramento Bee created a database of California state worker salaries, University of California (UC) employees could learn about their colleagues' pay for the first time. Researchers at three UC campuses studied the impact by randomly emailing some employees about the database (Card et al., 2012).

Employees who learned they earned below the median for their department and role reported lower job satisfaction and began searching for new positions. The

effect was strongest among workers in the bottom quartile of their unit's pay scale. A follow-up two to three years later suggested more lower-paid employees had left their positions.

In contrast, employees above the median showed no changes in satisfaction or behaviour after learning about peer salaries.

This research shows that people care about their position relative to others, not just their absolute circumstances. The dissatisfaction and subsequent jobseeking of low-paid UC employees revealed 'distributional preferences', preferences about the allocation of resources across individuals.

To understand how social comparisons influence economic decisions, we can model these distributional preferences mathematically. Economists model distributional preferences by extending utility functions to include others' outcomes, just as we model preferences over personal consumption.

This approach allows us to analyse how social considerations influence economic decision-making in contexts ranging from charitable giving to workplace compensation. Understanding distributional preferences can help us explain behaviour that deviates from pure self-interest in experimental games. We can use them to examine how inequality and social comparisons affect economic decisions. This can inform the design of workplace arrangements and economic policies that account for social preferences

In what follows, I examine two forms of distributional preferences in detail altruism and inequality aversion - and consider other social preferences. Each helps explain different aspects of observed behaviour, from charitable giving motivated by altruism to rejection of unfair offers driven by inequality aversion.

49.2 Altruism

Altruism is concern for the outcomes of others, with positive weights on others' wellbeing.

To incorporate altruism, we give a positive weight to the utility of others in the utility function. Here's a simple example where an individual's utility depends on both their own outcome and another person's:

$$U_i(x_i, x_j) = x_i + \alpha x_j$$

The utility of agent i, U_i , is a function of the the outcome for agent i, x_i , and agent j, x_j . α is some number greater than zero, with utility increasing linearly with an increased payoff for either agent.

Altruism comes in two forms. Pure altruism reflects genuine concern for others' wellbeing. Impure altruism occurs when people get a 'warm glow' from doing good, regardless of the actual impact on others.

49.2.1 Example: the public goods game

Consider a public goods game where two players each have \$10 and must simultaneously decide how much to contribute to a public project. The total contributions are multiplied by 1.5 and split equally between both players, regardless of their individual contributions.



Figure 49.1: Public goods game process

This game captures real-world situations: work teams deciding effort on shared projects, neighbors contributing to community improvements, or countries investing in climate change mitigation.

For this scenario, if both players contribute their full \$10, the pot becomes \$30 and each receives \$15.



Figure 49.2: Public goods game process with cooperation

If neither contributes, each keeps their \$10.



Figure 49.3: Public goods game process with no cooperation

If one contributes \$10 and the other \$0, the \$15 pot is split equally so the contributor finishes with \$7.50 and the non-contributor end up with \$17.50.

A self-interested player would contribute nothing. For any contribution by the other player, keeping your money always gives a higher payoff. However, both players would be better off if they both contributed everything.



Figure 49.4: Public goods game process with defection

Now suppose both players have altruistic preferences, placing a positive weight on the other's outcome (albeit lower than the weight on their own outcome):

$$U_i(x_i, x_j) = x_i + 0.75 x_j$$

With these preferences, we can calculate whether contributing is optimal if the other person contributes. If the agent contributes, they receive utility from their \$15 and the \$15 received by the other agent.

$$U_i(\text{contribute}) = 15 + 0.75 \times 15$$
$$= 26.25$$

If they do not contribute, they receive utility from their higher payoff of \$17.50, but the lower payoff of \$7.50 for the other agent acts as a drag.

$$\begin{split} U_i(\text{free ride}) &= 17.50 + 0.75 \times 7.50 \\ &= 23.125 \end{split}$$

With these altruistic preferences, contributing is optimal when the other player contributes. The weight placed on the other player's payoff ($\alpha = 0.75$) makes cooperation the best response to cooperation.

Contributing is actually the best response to any contribution by the other player. For example, here are the payoffs in response to a contribution of \$0 by the other player. If the agent contributes, they receive utility from their \$7.50 and the other player's \$17.50.

$$U_i(\text{contribute}) = 7.5 + 0.75 \times 17.5$$
$$= 20.625$$

If they do not contribute, they receive utility from their higher payoff of \$10, but the lower payoff of \$10 for the other player reduces their utility more than they gain for the increased personal payoff.

$$\begin{split} U_i(\text{free ride}) &= 10 + 0.75 \times 10 \\ &= 17.50 \end{split}$$

An altruistic player would contribute even if the other player is not an altruist.

This example shows how moderate altruism can overcome the free-rider problem in public goods provision. The positive weight two altruistic players place on the other's outcome makes mutual contribution a Nash equilibrium.

49.2.2 Limitations of the altruism model

While altruism helps explain behaviours like charitable giving, it has limitations.

Recall the ultimatum game. The first player (the proposer) receives a sum of money (say \$10) and must offer some portion of it to the second player (the responder). The responder then has two options:

- Accept the offer: both players get the proposed split
- Reject the offer: both players get nothing



Altruism can explain why proposers make generous offers, but it fails to explain why responders often reject low offers, knowingly reducing both players' payoffs. Such rejection would require negative weights on outcomes, contradicting the altruism model.

49.3 Inequality aversion

Inequality aversion offers an explanation for such rejections. We dislike being worse off than others, and to a lesser extent, we may feel uncomfortable being better off than others. This captures the intuition that people care about relative standing. Think of how employees react to learning their co-workers earn more, or the guilt some feel about their wealth relative to those in poverty.

49.3.1 A model of inequality aversion

To understand how inequality aversion affects decision-making, we need a formal way to represent these preferences. Fehr and Schmidt (1999) developed an influential model that captures both the dislike of having less than others and the potential discomfort of having more.

In their model, the utility of agent i, u_i , is a function of the payoff of agent i and the payoff of agent j.

$$u_i(x_i, x_j) = x_i - \alpha \max\{x_i - x_i, 0\} - \beta \max\{x_i - x_j, 0\} \alpha > 0 \quad \text{and} \quad \beta > 0$$

Each term in this equation represents a distinct psychological effect:

- x_i is their own outcome, such as their payoff or salary
- $\alpha \max\{x_j x_i, 0\}$ captures the sting of disadvantageous inequality, like discovering your colleague earns more than you
- $\beta \max\{x_i x_j, 0\}$ represents the discomfort of advantageous inequality, like feeling guilty about your high salary compared to others or the guilt of making a low offer in the ultimatum game.

The max function ensures that only the relevant term is included. If agent i has more than the other player, with x_i higher than x_j , the disadvantageous inequality term is zero. If agent i has less, the advantageous inequality term is zero.

To present the utility function in a form you may find more intuitive without the max function, we can rewrite it with utility of agent i being their own payoff, minus a penalty for either advantageous or disadvantageous inequality.

$$u_i(x_i, x_j) = x_i - \begin{cases} \beta(x_i - x_j) & \text{ if } x_i \geq x_j \\ \alpha(x_j - x_i) & \text{ if } x_i < x_j \end{cases}$$

Typically $\alpha > \beta$ as people dislike having less than others more than they dislike having more than others. We could also set $\beta < 0$ for an agent that likes to be better off than others.

To visualise how inequality aversion affects utility, the following graph shows a plot of utility as a function of x_i .

This utility function has a kink at x_j where agent *i* moves from having less to more than agent *j*. Below x_j , utility increases due to both the increasing material outcome for agent *i*, plus the reduction in inequality. If β is between 0 and 1 as in this diagram, the utility of agent *i*, $U(x_i)$, continues to increase in x_i above x_j , but at a decreasing rate, as inequality degrades the benefits of having more.



Figure 49.5: An inequality aversion utility function

49.3.2 Example: the ultimatum game

This theoretical model of inequality aversion helps explain real-world behaviors that cannot explained by pure self-interest. This includes behaviour in the ultimatum game, a classic example where people often act against their material self-interest.

A purely self-interested responder should accept any non-zero offer, as something is better than nothing. Yet people often reject unfair offers, such as job candidates rejecting a low salary even though this leaves both parties worse off. The pain of inequality outweighs the cash.

Suppose the responder in the ultimatum game has inequality averse preferences with $\beta = 0.25$ and $\alpha = 0.5$.

$$u_i(x_i, x_j) = x_i - 0.5 \max\{x_j - x_i, 0\} - 0.25 \max\{x_i - x_j, 0\}$$

What offers x would the responder reject where the proposer has \$10 to split between them?

If the responder rejects, the payoff to the proposer and responder is zero. That is:

$$x_P = x_R = 0$$

If the responder accepts, the responder receives x, and the proposer keeps the remainder. That is:

$$\begin{aligned} x_P &= 10 - x \\ x_R &= x \end{aligned}$$

The responder will accept if the utility of accepting is greater than the utility of rejecting.

$$U_R({\rm accept}) > U_R({\rm reject})$$

I will examine offers above and then below \$5. First, I will look at offers above.

We substitute the utility function into the inequality. As the α term equals zero, I do not include it here. We then enter the ultimatum game outcomes for acceptance and simplify.

$$U_R(\text{accept}) > U_R(\text{reject})$$

$$x_R - \beta(x_R - x_P) > 0$$

$$x - 0.25(x - (10 - x)) > 0$$

$$x - 0.25(2x - 10) > 0$$

This inequality will hold for any $\beta < 1$ as $x \ge 2x - 10$ for any $x \le 10$. Therefore, the condition will hold for the agent with $\beta = 0.25$. Recall that if $\beta < 1$ the responder has higher utility from a higher payoff but at a decreasing rate when they have more than the proposer.

When the offer is below \$5, the responder gets less than the proposer, so only α matters. Again, I substitute the utility function into the inequality, enter the ultimatum game outcomes and simplify.

$$U_R(\text{accept}) > U_R(\text{reject})$$

$$x_R - \alpha(x_P - x_R) > 0$$

$$x - 0.5(10 - x - x) > 0$$

$$x > 2.50$$

The responder with $\alpha = 1/2$ will reject offers below \$2.50.

We can plot the utility of the responder as a function of the offer x. As the offer is not independent of the proposer's payoff, I will derive the shape of the utility curve as a function of $x_R = x$. I do this by substituting the payoffs and values of alpha into the utility function.

$$\begin{split} U_R(x_P, x_R) &= x_R - \alpha \max\{x_P - x_R, 0\} - \beta \max\{x_R - x_P, 0\} \\ &= x - 0.5 \max\{10 - 2x, 0\} - 0.25 \max\{2x - 10, 0\} \end{split}$$

We can also write this equation to explicitly show the slope of the curve. I first manipulate the equation into a form showing the slope times (x) plus intercept. I then simplify.

$$\begin{split} U_R(x) &= \begin{cases} (1+2\alpha)x - 10\alpha & \text{ if } x < 5 \\ (1-2\beta)x + 10\beta & \text{ if } x \geq 5 \end{cases} \\ &= \begin{cases} 2x-5 & \text{ if } x < 5 \\ 0.5x+2.5 & \text{ if } x \geq 5 \end{cases} \end{split}$$

This diagram shows the responder's utility curve as a function of the offer x. The slope of the curve for x < 5 is $1 + 2\alpha = 2$ and for $x \ge 5$ is $1 - 2\beta = 0.5$. These slopes differ from our earlier inequality aversion utility plot because, here, any gain for the responder means an equal loss for the proposer.

This relationship creates a double effect on inequality. Consider an offer increase from \$4 to \$5: the responder's payoff rises by \$1, but the gap between players shrinks by \$2. This explains why the slope is steeper when behind.

The opposite happens when the responder is ahead. If the offer increases from \$6 to \$7, the responder gains \$1 but inequality grows by \$2. This larger penalty for inequality acts as a drag on utility, making the slope flatter.

The curve has a kink at x = 5 where the responder moves from having less to more than the proposer. The point where the curve crosses the x-axis represents the point below which the responder would reject the offer.

49.3.3 Example: the dictator game

While the ultimatum game shows how inequality aversion affects responses to unfair offers, the dictator game helps us understand how inequality aversion might influence giving behaviour.

In the dictator game, the dictator makes a unilateral offer to the receiver. The game then ends. The receiver has an empty strategy set.



Figure 49.6: Utility curve for the responder in the ultimatum game with inequality aversion $% \left({{{\mathbf{F}}_{\mathrm{s}}}^{\mathrm{T}}} \right)$



In this version of the game, the dictator has a constrained choice set and must decide between the allocations (0, 1) and (1, 5). That is, the dictator must choose between 0 for themselves and 1 for the receiver, or 1 for themselves and 5 for the receiver. The dictator's $\alpha = 1/2$. As the dictator has less than the other player under each distribution, α is the relevant parameter.



We can calculate the dictator's utility of each allocation. The (0,1) allocation gives a utility of zero for the dictator's payoff minus alpha times the inequality.

$$\begin{split} U_D(0,1) &= 0 - \alpha \times (1-0) \\ &= -1/2 \times 1 \\ &= -1/2 \end{split}$$

The (1,5) allocation gives a utility of one for the dictator's payoff minus alpha times the larger inequality.

$$U_D(1,5) = 1 - \alpha \times (5-1)$$

= 1 - 1/2 × 4
= -1

The dictator prefers to allocate (0,1) as the cost of the inequality is greater than the benefit from the increased payoff.

 $\alpha>0$ can also account for the rejection of low offers in the ultimatum game.

49.4 Other distributional preferences

While altruism and inequality aversion explain many observed behaviours, they are not the only types of social preference. Rearranging the inequality aversion model reveals the direct weights people place on others' outcomes. Where $x_i > x_j$, agent *i* places weight of β on the other's outcome and $1 - \beta$ on their own. Where $x_i < x_j$, agent *i* places weight of $-\alpha$ on the other's outcome and $1 + \alpha$ on their own.

$$u_i(x_i,x_j) = \begin{cases} \beta x_j + (1-\beta) x_i & \text{ if } x_i \geq x_j \\ -\alpha x_j + (1+\alpha) x_i & \text{ if } x_i < x_j \end{cases}$$

Arranging in this way makes the weight given to the outcomes for each agent more transparent. It provides a more intuitive way to consider some distributional preferences.

Charness and Rabin (2002) used this form with parameters σ and ρ (where $\sigma = -\alpha$ and $\rho = \beta$). We'll follow their approach since having parameters with the same sign makes the analysis more intuitive.

$$u_i(x_i, x_j) = \begin{cases} \rho x_j + (1 - \rho) x_i & \text{ if } x_i \geq x_j \\ \sigma x_j + (1 - \sigma) x_i & \text{ if } x_i < x_j \end{cases}$$

49.4.1 Forms of distributional preferences

We can adjust the values of ρ and σ to capture many forms of distributional preferences. Some are as follows.

An altruistic agent has $\sigma > 0$ and $\rho > 0$. They weight others' payoffs positively in all situations. When someone else gains a dollar, it directly increases the agent's utility, regardless of whether they are ahead or behind. Compared to our earlier altruism model, this model allows for different degrees of altruism depending on the relative payoffs of the two agents.

An inequality-averse agent has $1 \ge \rho > 0 > \sigma$. They like gains for others who are behind but dislike falling further behind themselves. This inequality is equivalent to $\alpha > 0$ and $\beta > 0$ in the Fehr-Schmidt model.

A status-seeking agent has $0 > \rho \ge \sigma$. Not only do they dislike others gaining, they actively prefer having more than others. Their utility goes up when either they get more or the other player gets less.

A classically self-interested agent has $\rho = \sigma = 0$. Their utility depends only on their own payoff.

A Rawlsian agent has $\rho = 1$ and $\sigma = 0$, making their utility $u_i(x_i, x_j) = \min\{x_i, x_j\}$. They maximise the minimum payoff across all players, showing pure concern for the worst-off individual.

A utilitarian agent has $\rho = \sigma = 1/2$, making their utility $u_i(x_i, x_j) = x_i + x_j$. They weight everyone's payoff equally and simply sum the total, caring about efficiency rather than distribution.

49.4.2 Example: the trust game

Let's use this form of the utility function to analyse the outcomes in a trust game, where an investor must decide whether to trust an entrepreneur with their money. The trust game provides a lens for understanding how social preferences affect economic transactions where success requires sequential cooperation and trust. The trust game captures two key real-world dynamics. First, one party must choose whether to make themselves vulnerable to exploitation by trusting the other. Second, the trusted party must then decide whether to honour or betray that trust.



This structure mirrors some common economic situations. Investors must decide whether to trust entrepreneurs with capital before knowing if they will act in good faith. Business partners often need to commit resources before knowing if their counterparts will uphold their end of the deal. Even simple transactions, like paying in advance for services, involve one party trusting another to deliver as promised.

In the exercises in Section 44.3, I considered a scenario between Linda, a potential investor, and Marco, an entrepreneur. This mirrors common situations where investors must decide whether to trust entrepreneurs with capital, and entrepreneurs must choose between short-term gain and maintaining their reputation. Their choices will depend not just on monetary payoffs, but on how each values the other's outcomes.

Linda is looking for investment opportunities. She identifies a promising crypto-based start-up created by Marco. Marco is looking for seed funding.

Linda can invest \$10.

If Linda invests, her investment will triple in value. Marco can then decide to either shut down the start-up and keep the \$30 or maintain the start-up in the market and pay a \$15 dividend to each of Linda and himself. If Linda does not invest, Linda keeps the \$10. The start-up gets \$0.



Marco is effectively playing a dictator game. If he were purely self-interested, he would shut down and keep the \$30. As a result, a purely self-interested Linda would not invest.



Let's examine both players' actions if their utility functions place weight on the outcome of the other. U_L and U_M are Linda and Marco's utility. x_L and x_M are the outcomes for Linda and Marco.

Linda's utility function shows moderate concern for Marco's outcomes:

$$U_L(x_M, x_L) = \begin{cases} \frac{2}{3} x_M + \frac{1}{3} x_L & \text{ if } x_L \ge x_M \\ \frac{1}{3} x_M + \frac{2}{3} x_L & \text{ if } x_L < x_M \end{cases}$$

When Linda is ahead, she places more weight on Marco's payoff ($\rho = 2/3$) than her own $(1 - \rho = 1/3)$, suggesting significant altruism. When behind, she still values Marco's payoff positively ($\sigma = 1/3$), but prioritises her own outcome $(1 - \sigma = 2/3)$.

Marco's preferences are more asymmetric:

$$U_M(x_L, x_M) = \begin{cases} \frac{3}{4}x_L + \frac{1}{4}x_M & \text{ if } x_M \geq x_L \\ \\ x_M & \text{ if } x_M < x_L \end{cases}$$

He shows concern for Linda's outcomes only when he's ahead, otherwise focusing solely on his own payoff. This asymmetry reflects a mix of self-interest and conditional altruism.

If Marco and Linda know each other's utility functions, what is the equilibrium with these distributional preferences?

If Linda chooses trust, Marco has a choice between 15 each and 30 for himself. Marco calculates the utility of each option.

$$U_M(15, 15) = \frac{3}{4}(15) + \frac{1}{4}(15)$$
$$= 15$$
$$U_M(0, 30) = \frac{3}{4}(0) + \frac{1}{4}(30)$$
$$= 7.5$$

Marco receives higher utility by paying the dividend to Linda.

We can also calculate Linda's utility for each of Marco's options if she chooses to trust.

$$U_L(15, 15) = \frac{2}{3}(15) + \frac{1}{3}(15)$$
$$= 15$$
$$U_L(0, 30) = \frac{1}{3}(30) + \frac{2}{3}(30)$$
$$= 10$$

Linda's utility is higher if Marco pays a dividend.

For the other node, if Linda does not invest, she will keep \$10. Marco will have nothing. Each has the following utility from that failure to invest.

$$\begin{split} U_M(10,0) &= 0 \\ U_L(0,10) &= \frac{2}{3}(0) + \frac{1}{3}(10) \\ &= 3.33 \end{split}$$

Putting those payoffs into the extensive form of the game, we get the following, with payoffs of (15,15) if Linda invests and Marco returns the dividend, (10,7.5) if Linda invests and Marco does not return the dividend, and (3.33,0) if Linda does not invest.



To solve for the game's equilibrium, we use backward induction, starting with Marco's decision if Linda invests.

Marco can return a dividend for utility 15 or shut down for utility 7.5. He chooses to return the dividend. As a result, Linda will invest for utility 15, rather than not invest for utility 3.33. Linda invests.



Chapter 50

Reputation

Summary

- People care about their social image and reputation, often fearing the stigma associated with selfish behaviour. This concern can influence decision-making beyond purely strategic considerations.
- Andreoni and Bernheim's (2009) experiment demonstrates this effect: dictators in a modified dictator game changed their offers based on whether the receiver would infer that the offer came from them.

50.1 Introduction

People care about what other people think. They fear the social stigma that can result from "selfish" behaviour.

Partly, this is for strategic reasons. For example, to attract reciprocal behaviour, people may need to be aware of your intentions.

However, there is also evidence that people care about what other people think.

50.2 Example

One example of this comes from Andreoni and Bernheim (2009), who ran a nonanonymous dictator game.

Each dictator was endowed with \$20.

A computer then chose a distribution between the dictator and the receiver, selecting either (\$0, \$20) or (\$20, \$0) with equal probability. The dictator observes the computer's allocation, but the receiver does not.

The computer's allocation is then implemented with a probability p. Otherwise, the dictator's allocation is made. This probability is known to both the dictator and the receiver.

If the dictator's choice is to be implemented, the dictator makes a split of the \$20, offering x to the receiver. The receiver learns only the allocation. They do not learn the dictator's choice.



Distributional preferences predict that p should not affect the dictator's choice. The

dictator should only think about the situation in which their choice matters.

However, the experimental results did not conform with this prediction. Dictators condition their decision on the common knowledge of p.

This chart shows how offers change with (p) when the computer's offer of 0 if selected. The x-axis shows p equal to 0, 0.25, 0.5 and 0.75. Each line represents a different bucket of offers. The red line is the proportion of dictators offering 0. The blue line represents the proportion of participants offering \$10, a 50:50 split.



FIGURE 3.—Distribution of amounts allocated to partners, condition 0.

For p of 0.5 or 0.75 and a computer allocation of 0 to the receiver, most dictators will offer 0. If the receiver receives a low allocation, the receiver will likely infer it is due to the computer's decision. They will not blame the dictator.

For p of 0 and a computer allocation of 0 to the receiver, more than half of dictators will offer \$10 to the receiver. In this case, if the receiver receives a low allocation, the receiver will infer it is due to the dictator's decision, not the computer's.

For a p of 0.25, a slim majority of participants offer zero, with some plausible deniability due to the 25% probability of the offer coming from the computer.

These results suggest the dictator cares about their reputation in the eyes of the receiver.

Chapter 51

Reciprocity

Reciprocity involves like-for-like behaviour. Kindness is responded to with kindness. Unkindness is responded to with unkindness.

Reciprocity might be considered to have two forms.

The first is instrumental reciprocity. Agents reciprocate behaviour due to the longterm benefits of sustained cooperation. The behaviour is motivated by the positive trade-off between long-term and short-term gains.

The second is intrinsic reciprocity. Agents reciprocate behaviour despite the absence of long-term gains.

51.1 Intentions

We can see evidence for reciprocity in how people respond to the intentions of others in the ultimatum game.

The ultimatum game involves two players: the proposer and the responder.

The proposer is given a fixed amount of money m. They then offer a portion x of the sum m to the responder.

The responder can either accept or reject the offer. They make this decision knowing the fixed amount m held by the proposer and the offer x.

If the responder accepts, the responder receives the offer x and the proposer gets the remainder m - x. If the responder rejects, both players receive nothing.


Consider this variation of the ultimatum game. In each of two scenarios, the proposer has a constrained choice of offers.

In scenario 1, the proposer has a choice between offering a split of \$8 for the proposer and \$2 for the responder or \$5 for the proposer and \$5 for the responder. Responders tend to reject offers of \$2.



In scenario 2, the proposer has a choice between offering a split of \$8 for the proposer and \$2 for the responder or keeping the full \$10 for themselves. Responders tend to accept offers of \$2.



In the first scenario, responders reject offers of \$2. In the second, they accept them. How can (\$8, \$2) be better than (\$0, \$0) in one scenario but not in the other?

Distributional concerns cannot explain rejection in this case. An offer of (\$8, \$2) leads to the same distribution in both scenarios.

Instead, responders do not base their decision on the outcome alone. They use their knowledge of the proposer's options, consider the proposer's intentions and reciprocate them.

A proposer who offers \$2 instead of \$0 is seen as having good intentions. A proposer who offers \$2 instead of \$5 is seen as having bad intentions.

51.2 The trust game

The trust game, which I introduced in Section 41.4.3, provides another potential example of reciprocity.

The trust game involves two players: a sender and a receiver

Both the sender and receiver are given an initial sum m.

The sender sends a share x of their m to the receiver. This amount x is often called the investment.

Before the receiver receives the investment, it is multiplied by some factor k. Therefore, the receiver receives kx.

The receiver then returns to the sender some share y of their total allocation m + kx.

The final outcome is the sender has m - x + y and the receiver has m + kx - y.



The game theoretic equilibrium is for the receiver to return nothing, so the sender sends nothing.

Contrast this with what happens in experimental settings.

Senders tend to send a positive amount, typically around half of their endowment.

Receivers tend to send back a bit less than is sent.

These two figures from N. D. Johnson and Mislin (2011) illustrate the distribution of investments and returns in 162 replications of the trust game.



Figure 51.1: Distribution of proportion sent and proportion returned (N. D. Johnson and Mislin, 2011).

One possible explanation for this behaviour is that the receiver feels they should reciprocate the sender's investment. They are responding to the sender's intentions. The sender trusts that some of their investment will be repaid due to reciprocation.

The receiver's behaviour is also consistent with altruism and inequality aversion.

51.3 The public goods game

As a final example of reciprocity, consider the public goods game.

Each of N participants is given an initial endowment.

Each participant secretly and simultaneously chooses how much of their endowment they wish to contribute to a public pot.

The money in the public pot is multiplied by some amount and split evenly between the players. Typically, the multiple applied to the pot is greater than one but less than the number of players.



In Nash equilibrium in the public goods game, nobody transfers anything to the pot. Any contributions are split between all players, so if there are more players than the multiple, which is normally the case by design, contributions result in a loss to that individual player.

This is not what we see when people play the public goods game in the lab.

In a meta-analysis, Zelmer (2003) found an average contribution of 38% of the endowment. The amount contributed increased with the marginal per capita return; that is the higher k/N.

One possible explanation is that players trust that the other players will contribute, so they desire to reciprocate the expected contributions from others.

Another explanation hinges on social norms. Where a norm of behaviour exists, people tend to follow it.

Chapter 52

Social preferences exercises

52.1 Fehr-Schmidt preferences

Alby has the following distributional preferences:

$$u_A(x_A,x_j) = \underbrace{x_A}_{(1)} \underbrace{-\alpha \max\{x_j - x_A, 0\}}_{(2)} \underbrace{-\beta \max\{x_A - x_j, 0\}}_{(3)}$$

where:

 x_A is the outcome for Alby

 x_j is the outcome for any agent j with whom Alby interacts.

a) For $\alpha > 0$ and $\beta > 0$, what are preferences of this form are normally called?

? Answer

Inequality aversion.

b) For $\alpha > 0$ and $\beta > 0$, describe the role of each of the three terms labelled (1), (2) and (3) in the utility function.

? Answer

The first term captures the utility of Alby's own outcome. The second term captures Alby's dislike of having less than others. The third term captures Alby's dislike of having more than others.

c) Explain the intuition for why we normally set $\alpha > \beta$.

Answer

Most people dislike having less than others more than they dislike having more than others. In some instances, $\beta < 0$ in which case people like having more than others - they are fine with inequality as long as it is to their advantage.

52.2 Charness-Rabin preferences

Bob has the following distributional preferences:

$$u_B(x_B,x_j) = \begin{cases} \rho x_j + (1-\rho) x_B & \text{ if } x_B \geq x_j \\ \sigma x_j + (1-\sigma) x_B & \text{ if } x_B < x_j \end{cases}$$

where:

 x_B is the outcome of the game for Bob

 x_i is the outcome of the game for any agent j with whom Bob interacts.

a) For $1 \ge \rho \ge 0 \ge \sigma$, what are preferences of this form are normally called?

? Answer

Inequality aversion.

b) For $1 \ge \rho \ge 0 \ge \sigma$, describe the role of the terms in each of the two equations where $x_B \ge x_j$ and $x_B < x_j$.

? Answer

 σ and ρ are the weight that is Bob gives to the outcome for agent j. σ is applied where Bob's outcome is better than or equal to that of agent j, and ρ where it is worse.

The residual $1 - \sigma$ and $1 - \rho$ is the weight that Bob gives to his own outcome.

c) Explain the intuition why we normally set $\rho > \sigma$ for the utility function.

? Answer

People tend to be more willing to see others have better outcomes when those others are worse off than them. Therefore, ρ should be greater than σ so that the agent cares more about the other agent when they are the one receiving more.

d) What values of ρ and σ would result in a utility function where Bob is purely self-interested?

💡 Answer

If Bob were purely self interested, ρ and σ would have a value of zero. In that case, agent j's outcomes would not enter into the utility function. The utility function would become $u_B(x_B) = x_B$.

e) What value must σ have to explain Bob's rejection of low offers in the ultimatum game?

? Answer

Sigma must be negative such that, if agent i accepts, the decrease in utility from agent j's payoff would be larger than the utility gain agent i would receive from its own payoff.

f) Consider the following two scenarios involving the Ultimatum game.

Scenario 1: A proposer has a choice between offering a split of (\$8, \$2) or (\$5, \$5). In experiments with this choice, responders tend to reject offers of (\$8, \$2).

Scenario 2: A proposer has a choice between offering a split of (\$8, \$2) or (\$10, \$0). In experiments with this choice, responders tend to accept offers of (\$8, \$2).

A utility function of the type that Bob has cannot result in this behaviour. Explain why.

? Answer

In both cases, the outcome is (\$8, \$2). This would result in the same level of utility regardless of the other option that the proposer had. If compared with the outcome (\$0, \$0) from rejecting the offer, the action should therefore be the same.

An explanation for the difference between scenarios is that people care not just about outcomes, but also the intentions of those with whom they interact. In that circumstance, the good (or otherwise) intentions of the proposer in offering either less than they could or as much as they could would shape the responder's action. However, intentions do not enter into Bob's utility function.

52.3 A government grant

Jason, an academic, is seeking an industry research grant from the government and Bezzle Bank. Under the grant terms, a \$50,000 contribution by Bezzle Bank will be matched by government. The total research funds of \$100,000 are then under Jason's control.

Jason could use these for research proposal A that benefits both himself and the bank, giving the bank \$100,000 in business benefit and Jason \$100,000 in career benefit. Alternatively, Jason could use the funds to research a new idea B, giving him \$150,000 in career benefit, but little of use for the bank.

Bezzle Bank knows Jason's options and is considering whether to make a contribution.

a) Show that the Nash equilibrium outcome is for the bank to make no contribution.

? Answer

I have represented this problem in the following extensive form game. The possible payoffs for Bezzel Bank are \$50,000 for proposal A (\$100,000 minus it's \$50,000 contribution) and -\$50,000 for proposal B (loss of its contribution). If Bezzle Bank does not contribute, both Jason and the Bank have a payoff of zero.

We work through the problem by backward induction.

First, Jason has a choice between Proposal A for a payoff of \$100,000 and Proposal B for a payoff of \$150,000. He chooses proposal B.

Bezzle Bank therefore has a choice between contributing 50,000 for a payoff of -550,000 or not contributing to remain at the status quo. Bezzle Bank does not contribute.



b) The bank decides to make the contribution and Jason decides to focus on research proposal A. Give possible reasons for this deviation from the Nash equilibrium.

? Answer

One explanation involves reciprocation. Jason is responding to the Bank's intentions and wants to reciprocate them. Relatedly, Bezzle Bank trusts Jason based on a belief that Jason will reciprocate their intentions.

Another explanation may relate to reputation. Jason might care about his reputation and be reluctant to take an action that harms it. This may be due to concerns about his self-image, or more instrumental reasons such as a poor reputation harming future opportunities.

Altruism on the part of the bank may also be a factor. They may gain utility from the benefit to Jason (or the broader societal benefit from the research project).

Part X

Practice problems

This section comprises practice problems drawn from past assessments delivered in the subject.

Chapter 53

Expected utility theory problems

53.1 A 60:40 gamble

Penny is an expected utility maximiser with utility function $u(x) = \ln(x)$ and wealth of \$300.

Penny is offered the following bet A:

- a 60% probability to win \$150
- a 40% probability to lose \$100.

a) Does Penny accept bet A?

? Answer

Penny compares the utility of taking versus not taking the bet:

$$U(\mathbf{A}) = p_1 u(x_1) + p_2 u(x_2)$$

= 0.6 ln(W + 150) + 0.4 ln(W - 100)
= 0.6 ln(450) + 0.4 ln(200)
= 5.785
$$U(W) = ln(W)$$

$$= \ln(300)$$

= 5.704

U(A) > U(W), so Penny accepts the bet.

b) Following some bad economic news, Penny wealth declines to \$150.

Penny is offered bet A again. Does Penny accept the bet?

? Answer

Penny compares the utility of taking versus not taking the bet:

$$\begin{split} U(\mathbf{A}) &= p_1 u(x_1) + p_2 u(x_2) \\ &= 0.6 \ln(W + 150) + 0.4 \ln(W - 100) \\ &= 0.6 \ln(300) + 0.4 \ln(50) \\ &= 4.987 \end{split}$$

$$\begin{split} U(W) &= \ln(W) \\ &= \ln(150) \\ &= 5.011 \end{split}$$

U(A) < U(W), so Penny rejects the bet.

53.2 Another 60:40 gamble

Gamble A is as follows:

(\$100, 0.6; -\$100, 0.4)

This is a gamble with a 60% chance of winning \$100 and a 40% chance of losing \$100.

a) Would a risk neutral decision-maker (who maximises expected value) be willing to pay \$10 to play gamble A? What is the most they would be willing to pay to play?

Answer

A risk neutral decision maker will accept any offer with positive expected value. The expected value of the bet is:

$$\begin{split} E[A] &= p_1 x_1 + p_2 x_2 \\ &= 0.6 * 100 + 0.4 * (-100) \\ &= \$ 20 \end{split}$$

The risk neutral decision maker would pay \$10 as this is less than the expected value of the bet. They would be willing to pay up to the expected value of the bet: \$20. At that point they would be indifferent between paying for the bet and refusing the bet.

b) Would an expected utility maximiser with wealth \$200 and utility function U(x) = ln(x) be willing to pay \$10 to play gamble A? What is the most they would be willing

to pay to play?

? Answer

The expected utility maximiser will play if their utility from playing and paying is greater than their utility of refusing.

$$U(W) = ln(W)$$

= $ln(\$200)$
= 5.2983174

$$\begin{split} E[U(A-c)] &= p_1 U(x_1) + p_2 U(x_2) \\ &= 0.6 U(W+100-10) + 0.4 U(W-100-10) \\ &= 0.6 ln(200+100-10) + 0.4 ln(200-100-10) \\ &= 5.2018524 \end{split}$$

U(W) > U(A - c) so the decision maker will not be willing to pay \$10. To determine the most they would be willing to pay, we will first check whether they will pay any positive sum. We will do that by examining the expected utility of the gamble with no payment.

$$\begin{split} E[U(A)] &= p_1 U(x_1) + p_2 U(x_2) \\ &= 0.6U(W+100) + 0.4U(W-100) \\ &= 0.6ln(200+100) + 0.4ln(200-100) \\ &= 5.2643376 \end{split}$$

U(W)>U() so the decision maker will not be willing to pay any amount. In fact, they would pay to avoid the bet.

To calculate how much, we determine what the certainty equivalent of the bet is:

$$U(CE) = E[U(A)]$$

$$ln(CE) = 5.2643376$$

$$CE = e^{5.2643376}$$

$$= \$$$

Having wealth of \$200 and the bet is the equivalent of having wealth of \$193.32. They would be willing to pay up to \$6.68 to avoid the bet.

c) Would the expected utility maximiser with utility function change their decision if they had \$1000 in wealth? Explain.

Answer

The expected utility maximiser will play if their utility from playing and paying is greater than their utility of refusing.

$$\begin{split} U(W) &= ln(W) \\ &= ln(\$1000) \\ &= 6.9077553 \end{split}$$

$$\begin{split} E[U(A-c)] &= p_1U(x_1) + p_2U(x_2) \\ &= 0.6U(W+100-10) + 0.4U(W-100-10) \\ &= 0.6ln(1000+100-10) + 0.4ln(1000-100-10) \\ &= 6.9128484 \end{split}$$

U(W) < U(A - c) so the decision maker is now willing to pay \$10. As the agent's wealth increases their utility function becomes increasingly linear (second derivative approaches zero) and they become closer to risk neutral. As a result, the positive expected value bet becomes increasingly attractive.

d) At what wealth is the expected utility maximiser with utility function U(x) = ln(x) indifferent between accepting gamble A or not?

Answer

The expected utility maximiser will be indifferent when:

$$\begin{split} U(W) &= E[U(A)]\\ U(W) &= 0.6U(W+100) + 0.4U(W-100)\\ ln(W) &= 0.6ln(W+100) + 0.4ln(W-100) \end{split}$$

There isn't a simple closed form solution to this equation, but we know from questions b) and c) that W is somewhere between \$200 and \$1000.

If we wanted to calculate exact solution, you could use tool such as Mathematica or Matlab to solve, write a short code to solve in R or even iterate toward a solution using Excel.

The expected utility maximiser is indifferent when W = \$256.81.

53.3 Shannon and Simon

Shannon and Simon are offered a choice between options A and B:

A: Lottery X = (0.5, -\$100; 0.5, -\$20). This is a gamble with a 50% chance of losing \$100 and a 50% chance of losing \$20.

B: A loss of \$30 for certain.

(a) Shannon is risk-neutral. Will Shannon choose A or B?

A risk-neutral decision-maker maximises expected value. The expected value of option A is:

$$\begin{split} \mathbf{E}[A] &= p_1 x_1 + p_2 x_2 \\ &= 0.5 \times -\$100 + 0.5 \times -\$20 \\ &= -\$60 \end{split}$$

The expected value of option B is -\$30. Shannon will choose the certain loss because -\$30 is a better outcome than -\$60.

(b) Simon is risk averse with wealth \$200 and utility function $U(x) = x^{0.5}$. Will Simon choose A or B?

? Answer

? Answer

Simon will select the option that gives the highest expected utility. The expected utility of option A is:

$$\begin{split} \mathrm{E}[U(A)] &= p_1 u(x_1) + p_2 u(x_2) \\ &= 0.5 \times (W - 100)^{0.5} + 0.5 \times (W - 20)^{0.5} \\ &= 0.5 \times (100)^{0.5} + 0.5 \times (180)^{0.5} \\ &= 11.71 \end{split}$$

The expected utility of option B is:

$$\begin{split} \mathrm{E}[U(B)] &= u(W+x) \\ &= (200-30)^{0.5} \\ &= 13.04 \end{split}$$

Simon will choose option B because E[U(A)] = 11.71 < 13.04 = E[U(B)].

(c) Draw a graph showing the choices faced by Simon, his utility curve and the expected utility of each option. Indicate the certainty equivalent of option A. Explain how the graph shows which option Simon will choose.

Answer

This chart shows the choices faced by Simon, his utility curve and the expected utility of each option. The horizontal axis is the outcome and the vertical axis is utility of each outcome. The utility curve is the function $U(x) = x^{0.5}$.

The two possible outcomes of gamble A are W-100=100 and W-20=180, which deliver U(W - 100) and U(W - 20) respectively. Each are labelled on the chart. The expected utility of gamble A is the weighted average of these two utilities and lies on the straight line between U(W - 100) and U(W - 20). As each outcome has a 50% chance of occurring, the expected utility of gamble A is the midpoint of this line (as is the expected value of the gamble). The vertical line from W + E[A] = 140 identifies that point, with the expected utility E[U(A)] marked on the vertical axis.

The certain outcome from option B, the loss of \$30 resulting in wealth of \$170 is also marked on the x-axis, leading to utility of U(W - B).

It can be seen that the expected utility of gamble A, E[U(A)], is less than the utility of the certain outcome U(W - B). Simon will therefore choose B.

The certainty equivalent of option A is identified as the point where U(CE) = E[U(A)]. This is identified by drawing a horizontal line from the expected utility of gamble A to the utility curve. The point where this line intersects the utility curve is the certainty equivalent of gamble A, shown by projecting a vertical line downward.

This diagram is not drawn to scale.



53.4 Starting a business

(a) Stacey, an expected utility maximiser, was considering quitting her job to start a new business. She calculated that her business proposal had a positive expected value but that there is a material risk of a loss. She decided not to start the business.

Use concepts from expected utility theory to explain why Stacey may have decided not to start the business.

? Answer

Stacey may have decided not to pursue the business opportunity as she is risk averse. Someone who is risk averse will value the expected value of a gamble less than the equivalent amount paid with certainty. This means that a positive value bet might be turned down.

(b) Use a graph to demonstrate your answer to part (a).

? Answer

This graph shows Stavey's utility curve. As she is risk averse, the curve is concave (at least over the domain of the decisions we are considering).

Suppose there are two possible outcomes from starting the business. a moderate loss and a large gain, each having equal probability. Each of the outcomes of the business are labelled. Stacey finishes with his wealth minus the business loss (W - L) or her wealth plus the business gains (W + G). If she does not start the business, her wealth remains at W. The utility of each possible outcome (U(W - L), U(W), U(W + G)) is also indicated on the vertical axis.

The expected value of starting the business is labelled (E[X]). As the business has positive expected value, E[X] is greater than W.

The expected utility of the business lies on the straight line between the utility of the two possible business outcomes. The place on the line is determined by the probability of winning and is in line with the expected value of the business. We can identify the expected utility of the business by projecting a line up from E[X] to the straight line.

Finally, the certainty equivalent of the business is also marked. As U(CE) = E[U(X)], we can identify the certainty equivalent by projecting a line up from E[U(X)] to the utility curve.

Due to the convex curve, we can see that E[U(X)] is less than U(W). Stacey prefers the certain outcome of W to starting the business. Alternatively, we can see that the certainty equivalent of the business is lower than current wealth.



(c) Later, Stacey received a large inheritance from her Aunt. She then decided to quit work and go ahead with her business plan.

Use concepts from expected utility theory to explain why Stacey may have changed her mind.

? Answer

That Stacey later decided to pursue the business opportunity is likely because she is more risk neutral (or possibly even risk seeking) at higher wealth. The result is that the positive expected value gamble is more likely to be accepted at the higher wealth level.

One mechanism by which this could occur is the utility function becomes increasingly linear with higher wealth (as is the case with $u(x) = \ln(x)$ or $u(x) = x^{0.5}$).

53.5 Risk neutrality versus risk aversion

Anika and Anthony are offered a choice between options A and B:

A: Lottery A = (0.5, \$100; 0.5, \$20). This is a gamble with a 50% chance of winning \$100 and a 50% chance of winning \$20.

B: \$40 for certain.

(a) Anika is risk-neutral. Will Anika choose A or B?

? Answer

A risk-neutral decision-maker maximises expected value. The expected value of option A is:

$$\begin{split} \mathbf{E}[A] &= p_1 x_1 + p_2 x_2 \\ &= 0.5 \times \$100 + 0.5 \times \$20 \\ &= \$60 \end{split}$$

The expected value of option B is \$40. Anika will choose option A because \$60 is greater than \$40.

(b) Anthony is risk averse with wealth \$100 and utility function $U(x) = \ln(x)$. Will Anthony choose A or B?

? Answer

Anthony will select the option that gives the highest expected utility. The expected utility of option A is:

$$\begin{split} \mathbf{E}[U(A)] &= p_1 u(x_1) + p_2 u(x_2) \\ &= 0.5 \times \ln(W + 100) + 0.5 \times \ln(W + 20) \\ &= 0.5 \times \ln(200) + 0.5 \times \ln(120) \\ &= 5.04 \end{split}$$

The expected utility of option B is:

$$\mathbf{E}[U(B)] = u(W+x)$$

 $= \ln(140)$

= 4.94

Anthony will choose option A because E[U(A)] = 5.04 > 4.94 = E[U(B)].

(c) Draw a graph showing the choices faced by Anthony, his utility curve and the expected utility of each option. Indicate the certainty equivalent of option A. Explain how the graph shows which option Anthony will choose.

Answer

Figure 53.2 shows the choices faced by Anthony, his utility curve and the expected utility of each option. The horizontal axis is the outcome and the vertical axis is utility of each outcome. The utility curve is the function $U(x) = \ln(x)$. The two possible outcomes of gamble A are W+20=120 and W+100=200, which deliver U(W + 20) and U(W + 100) respectively. Each are labelled on the chart. The expected utility of gamble A is the weighted average of these two utilities and lies on the straight line between U(W + 20) and U(W + 100). As each outcome has a 50% chance of occurring, the expected utility of gamble A is the midpoint of this line (as is the expected value of the gamble). The vertical line from W + E[A] = 160 identifies that point, with the expected utility E[U(A)] marked on the vertical axis.

The certain outcome from option B, the receipt of \$40 resulting in wealth of \$140 is also marked on the x-axis, leading to utility of U(W + B).

It can be seen that the expected utility of gamble A E[U(A)] is greater than the utility of the certain outcome U(W+B). Anthony will therefore choose gamble A.

The certainty equivalent of option A is identified as the point where U(CE) = E[U(A)]. This is identified by drawing a horizontal line from the expected utility of gamble A to the utility curve. The point where this line intersects the utility curve is the certainty equivalent of gamble A, shown by projecting a vertical line downward.

This diagram is not drawn to scale.



53.6 Lottery ticket

Buying a lottery ticket has a negative expected value. Andrew is an expected utility maximiser. He purchases a lottery ticket.

(a) What risk preferences (attitude to risk) does Andrew have?

? Answer

If an expected utility maximiser purchases a lottery ticket with negative expected value, he is risk seeking. He values the gamble over and above the expected value of the gamble.

(b) Use a graph to demonstrate your answer to part (a).

? Answer

Figure 53.3 shows Andrew's utility curve. As he is risk seeking it is convex (at least over the domain of the lottery).

Each of the outcomes of the lottery are labelled. And rew finishes with his wealth minus the cost of the lottery ticket (W - T) or his wealth minus the cost of the lottery ticket plus his lottery winnings (W - T + L). If he does not purchase the ticket, his wealth remains at W. The utility of each possible outcome (U(W - T), U(W), U(W - T + L)) is also indicated on the vertical axis.

The expected value of the lottery after buying the ticket is labelled (E[X]). As the lottery has a negative expected value, E[X] is less than W.

The expected utility of the lottery lies on the straight line between the utility of the two possible lottery outcomes. The place on the line is determined by the probability of winning and is in line with the expected value of the lottery. We can identify the expected utility of the lottery by projecting a line up from E[X] to the straight line.

Finally, the certainty equivalent of the lottery is also marked. As U(CE) = E[U(X)], we can identify the certainty equivalent by projecting a line up from E[U(X)] to the utility curve.

Due to the convex curve, we can see that E[U(X)] is greater than U(W). And rew prefers the lottery to the certain outcome of W. Alternatively, we can see that the certainty equivalent of the lottery is higher than current wealth. And rew would require a payment of at least CE - W to forgo his opportunity to partake in the lottery.



Chapter 54

Prospect theory problems

54.1 Question 1

Suppose Lisa has the following reference-dependent value function:

$$v(x) = \begin{cases} x^{1/2} & \text{where} \quad x \ge 0\\ -2(-x)^{1/2} & \text{where} \quad x < 0 \end{cases}$$

x is the change in Lisa's position relative to his reference point.

(a) Explain what element of Lisa's value function leads her to exhibit loss aversion.

? Answer

The multiplier of 2 applied to outcomes in the loss domain leads to greater weight for losses relative to gains.

(b) Lisa considers a choice between:

- a loss of \$100 for certain or
- gamble A: (0.6, -\$250; 0.4, 0). That is, gamble A gives a 60% chance of losing \$250 and a 40% chance of losing nothing.

Will Lisa prefer the certain loss or gamble A?

Answer

The weighted value of gamble A is:

$$v(A) = 0.6v(-250) + 0.4v(0)$$
$$= 0.6(-2(250)^{0.5}) + 0.4(0)$$
$$= -18.97$$

The value of the certain \$100 loss is:

$$v(\$100) = v(-100)$$

= -2(100)^{0.5}
= -20

As V(A) > V(-\$100), Lisa will prefer the gamble to the certain \$100 loss. Another way to think about this problem is to calculate the certainty equivalent of gamble A and compare that to the certain loss. The certainty equivalent is the amount of money that Lisa would be indifferent between receiving for certain and the gamble. A benefit of calculating the certainty equivalent is that you can then easily determine whether Lisa would prefer the gamble to a certain loss of any particular amount.

The certainty equivalent (CE) of gamble A is calculated as:

$$v(CE) = V(A) = -18.97$$

 $-2 \times (-CE)^{0.5} = -18.97$
 $(-CE)^{0.5} = 9.49$
 $CE = -90$

The certainty equivalent of gamble A is a loss of \$90. Lisa would be indifferent between losing \$90 for certain and the gamble. For any loss greater than \$90, Lisa would prefer the gamble. Given this information, we can say that Lisa will prefer the gamble to the certain loss of \$100.

(c) Explain what features of the value function leads Lisa to accept or reject the gamble.

? Answer

The bet has a lower expected value (a loss of \$150) relative to the certain outcome of a loss of \$100. Lisa would need to be risk seeking to find the bet attractive.

Lisa's diminishing sensitivity to losses leads her to be risk seeking in the loss domain, which is where the gamble is located. This is why she prefers the gamble to a certain loss.

Loss aversion does not affect her decision, as all possible outcomes (at least

under our assumed reference point) are in the loss domain.

(d) Draw a diagram illustrating the choice in part (b). Show the possible payoffs, Lisa's value function and the value of each option. Explain how this diagram illustrates Lisa's decision.

? Answer

The diagram below shows the possible payoffs, Lisa's value function and the value of each option.

The possible outcomes from the gamble are zero and a loss of \$250. The certain outcome on offer is a loss of \$100. The expected value of the gamble is $0.6 \times -$ \$250 + 0.4 × 0 = -\$150.

As Lisa is risk seeking, the value of the certain \$100 loss is lower than the weighted value of the gamble. This can be seen through v(A) being greater than v(-\$100). Lisa will, therefore, choose the gamble.



(e) Suppose Lisa were to win \$100 in a raffle. The win does not change her reference point. Could this win cause her to change her decision concerning the loss of \$100 and gamble A?

? Answer

The weighted value of gamble A is:

$$\begin{aligned} v(A) &= 0.6v(100 - 250) + 0.4v(100 - 0) \\ v(A) &= 0.6v(-150) + 0.4v(100) \\ &= 0.6(-2(150)^{0.5}) + 0.4(100)^{0.5} \\ &= -10.7 \end{aligned}$$

The value of the certain 100 loss is:

```
v(\$100) = v(100 - 100)v(\$100) = v(0)= -2(0)^{0.5}= 0
```

As V(A) < V(-\$100), Lisa will prefer the certain loss.

(f) Explain what features of the value function lead Lisa to accept or reject the gamble in part (e). How does this differ from your answer in part (c)?

? Answer

The bet has a lower expected value (a loss of \$150) relative to the certain outcome of a loss of \$100. Or to put it another way, the expected outcome of the gamble relative to her reference point (a loss of \$50) is less than the certain outcome of \$0 relative to her reference point.

Therefore, Lisa would need to be risk seeking across most of the relevant domain to find the bet attractive. She is risk seeking across some of the domain (loss), but risk averse across some (gain). The net result of that curvature is an insufficient level of risk seeking to make the bet attractive.

Similarly, loss aversion makes the bet less attractive, but it would still be rejected even in the absence of loss aversion. (We can see this by calculating the weighted value of the gamble without the loss aversion multiplier.)

(g) Draw a diagram illustrating the choices in parts (e). Show the possible payoffs, Lisa's value function and the value of each option. Explain how this diagram illustrates Lisa's decision after the raffle win.

Answer

This diagram illustrates Lisa's decision after the win.

The possible outcomes from the gamble relative to the reference point are a gain of 100 (a loss of zero plus the 100 win) and a loss of 150 (a loss of 250 plus the 100 win). The certain outcome on offer is zero (a certain loss of

\$100 plus the \$100 win). The expected value of the outcome after the gamble is $0.6\times-\$150+0.4\times100=-\$50.$

As the bet has negative expected value, plus Lisa is loss averse, V(A) is less than the alternative outcome of zero.



Figure 54.2: Lisa's consideration of gamble A and the $100\$ loss after the win

54.2 Question 2

Early in the COVID-19 pandemic, the hope was that vaccines would prevent transmission and death. Later, it became apparent that vaccination dramatically reduced death rates but did not wholly prevent it.

Use the concept of probability weighting to explain why citizens would see a massive reduction as much less valuable than elimination.

? Answer

There is strong evidence that we overweight certain events, when the probability of the event is one, relative to near certain events, such as when the probability is, say, 99%. This overweighting of certainty is effectively the same as overweighting very low-probability events.

The elimination of COVID-19 would be seen as a certain event, and therefore overweighted relative to a near-certain event, such as a 99% reduction in

deaths. Hence elimination would be seen as much more valuable even though near-elimination has nearly the same number of lives saved

Put another way, the remaining small probability of death would be overweighted, meaning the near elimination of deaths would be seen as having much less value than the complete elimination.

Chapter 55

Intertemporal choice problems

55.1 Work or party?

Ruby is a sophisticated present-biased agent with $\beta = 0.5$ and $\delta = 1$. She has utility function u(x) = x.

Ruby is deciding today (t = 0) whether she will either:

- work tomorrow (t = 1) for \$10 income to be received the day after she works (t = 2) or
- party tomorrow (t = 1) for immediate utility of 8 but no income.

That is, she is deciding between (2, 10) and (1,8).

a) What does Ruby decide?

? Answer

Ruby is sophisticated, so uses backward induction to decide her preferred course of action.

At t=2 there is no decision to make. Ruby bears the consequences of her earlier decisions.

At t = 1 she compares the discounted utility of the two options:

$$\begin{split} U_1(\mathrm{work}) &= \beta \delta u(\mathrm{work}) \\ &= 0.5 \times 1 \times 10 \\ &= 5 \\ U_1(\mathrm{party}) &= u(\mathrm{party}) \\ &= 8 \\ \end{split}$$
 At t ~=~ 1 the discounted utility partying is higher than that of working

 $(U_1(\text{party}) > U_1(\text{work}))$, so Penny prefers to party. At t = 0 Penny knows that she will party at t = 1 no matter what she decides at t = 0, so she accepts that she will party.

b) Ruby is able to commit to working at t = 1 by posting a letter at t = 0 declining the party invitation (at no cost). Once she sends the letter, she cannot change her mind. Will she decline the invitation?

? Answer

The presence of the commitment device allows Ruby to include the option of working when she makes a decision at t = 0. She will now compare the discounted utility of using the commitment device with the discounted utility of working.

$$\begin{split} U_0(\text{commit+work}) &= \beta \delta^2 u(\text{commit+work}) \\ &= 0.5 \times 1^2 \times 10 \\ &= 5 \\ U_0(\text{party}) &= \beta \delta u(\text{party}) \\ &= 0.5 \times 1 \times 8 \end{split}$$

 $U_0({\rm commit+work})>U_0({\rm party}),$ so Ruby commits to working by declining the invitation.

= 4

c) Suppose declining the party invitation comes at a cost to Ruby's reputation at t = 2. What is the largest utility cost that Ruby would be willing to incur such that she would still use the commitment device of declining the invitation?

? Answer

The largest cost she would be willing to incur is the cost where the discounted utility of each option is equal. Setting the cost as c:

$$\begin{split} U_0(\text{commit}) &= \beta \delta^2 u(\text{work}-c) \\ &= 0.5 \times 1^2 \times (10-c) \\ &= 5-0.5c \\ \\ U_0(\text{party}) &= \beta \delta u(\text{party}) \\ &= 0.5 \times 1 \times 8 \\ &= 4 \end{split}$$

Ruby will be indifferent where 5 - 0.5c = 4 or where c = 2. The largest utility cost she would be willing to incur is 2.

d) Victoria is a naive present-biased agent with $\beta = 0.5$ and $\delta = 1$. She has utility function u(x) = x.

Victoria faces the same choice as Ruby and also has the chance to decline at t = 0 the party invitation at a cost to Victoria's reputation at t = 2. Will Victoria decline the invitation?

? Answer

As Victoria has the same discount function as Ruby, we know from question b) that Victoria will decide to work at t = 0. We also know from question a) that she will then change her mind and party at t = 1.

However, as Victoria is naive she does not see her future self-control problem and does not have the foresight to realise at t = 0 that she will not work as planned. As a result, she would see no need in the commitment device and would not be willing to incur any cost to use it.

55.2 Surfboard vs NFT

James is a naive present-biased agent with $\beta = 0.5$, $\delta = 0.95$ and u(x) = x.

James is planning to purchase a surfboard. James also likes to purchase Bored Ape non-fungible tokens (NFTs).

The surfboard that James wants to purchase will be available for sale in 2 weeks (t = 2). Purchasing the surfboard will give him utility of 10 on the day of purchase.

James also knows a new Bored Ape NFT will be available for sale next week (t = 1). Purchasing the Bored Ape NFT will give him utility of 8 on the day of purchase.

James has enough money to buy only the surfboard or the Bored Ape NFT. He cannot afford both.

(a) What does James choose at t = 0? Explain.

? Answer

To determine what James chooses at t = 0, we need to compare the discounted utility of the two options.

The discounted utility of the NFT received at t = 1 is:

$$U_0(1,8) = 0.5 \times 0.95^1 \times 8$$

= 3.8

The discounted utility of the surfboard received at t = 2 is:

$$\begin{array}{l} U_0(2,10) = 0.5 \times 0.95^2 \times 10 \\ = 4.51 \end{array}$$

James chooses the option with the highest discounted utility, which is the surfboard at t = 2.

(b) At t = 1 the NFT becomes available for sale. James is sitting at his laptop contemplating what he should do. What does James choose?

? Answer

To determine what James chooses at t = 1, we need to compare the discounted utility of the two options.

The discounted utility of the NFT received at t = 1 is:

$$\begin{array}{l} U_1(1,8) = 0.95^0 \times 8 \\ = 8 \end{array}$$

The discounted utility of the surfboard received at t = 2 is:

$$U_1(2,10) = 0.5 \times 0.95^1 \times 10$$

= 4.75

James chooses the option with the highest discounted utility, which is the NFT at t = 1. He has changed his mind. This is because the NFT at t = 1 is no longer subject to the short-term discount factor β .

(c) Draw a graph illustrating James's choices and his decisions.

Answer

The following chart shows each of the two options presented to James, the NFT at t = 1 and the surfboard at t = 2. The line extended from each back to t = 0 represents the discounted utility of each option at time t.



(d) Jane is a sophisticated present-biased agent with $\beta = 0.5$, $\delta = 0.95$ and u(x) = x. Jane also likes surfing and Bored Ape NFTs. She faces the same set of choices as James in part a). What does Jane choose at t = 0? Explain.

? Answer

To determine what Jane chooses at t = 0, Janes considers the discounted utility of the two options by backward induction.

At t=2, Jane will choose the surfboard if she has not already purchased the NFT.

At t = 1, the discounted utility of the NFT is:

$$\begin{array}{l} U_1(1,8) = 0.95^0 \times 8 \\ = 8 \end{array}$$

At t = 1, the discounted utility of the surfboard received at t = 2 is:

$$\begin{split} U_1(2,10) &= 0.5 \times 0.95^1 \times 10 \\ &= 4.75 \end{split}$$

Jane would choose the NFT.

Jane now considers what she would choose at t = 0. She knows that no matter what she chooses now, she will choose the NFT at t = 1. Hence, she chooses the NFT at t = 0, the only feasible option.

(e) Jane discovers that she can place an early deposit of \$100 on the surfboard at t = 0. If she does so, she cannot afford the Bored Ape NFT at t = 1. Unfortunately, she will have to walk to the surf shop to place the deposit at a cost of 0.1 utility. What does Jane choose at t = 0? Explain.

🔮 Answer

Jane's logic runs as per part d) until she considers what she would choose at t = 0.

At t = 0, she has two choices available to her: place the deposit on the surfboard or purchase the NFT.

The discounted utility of the NFT at t = 1 is:

$$\begin{split} U_0(1,8) &= 0.5 \times 0.95^1 \times 8 \\ &= 3.8 \end{split}$$

The discounted utility of the surfboard at t = 2 when making the deposit is:

$$\begin{array}{l} U_0(2,10) = -0.1 + 0.5 \times 0.95^2 \times 10 \\ = 4.41 \end{array}$$

As the deposit binds Jane in the future, she can stick to her plan. She chooses buying the surfboard with the deposit rather than buying without the deposit as she knows that the latter choice is not viable: she will change her plans at t = 1.

55.3 Clearing credit card debt

Yvonne and Wendy each have a credit card with a credit limit of \$3,000. They had used the available credit on the card. They faced a 20% interest rate on their debts.

Yvonne and Wendy also had \$3,000 in their bank accounts. They were not earning any interest on this sum.

Yvonne used her savings to clear her credit card debt, saving her interest. Unfortunately, over the next three months, she used the credit card to make purchases and ended up with a credit card debt of \$3000 again. She is now paying 20% interest on her debt and has no savings.

Wendy put her savings into a term deposit, earning 5% interest. She cannot access her savings for the next year. She kept the credit card debt and made no further purchases but is still paying the 20% interest.

Use concepts from the subject material on intertemporal choice to explore why Yvonne

and Wendy might have each made the decisions that they did. You should not make any mathematical calculations in answering this question. Explain the intuition only.

Answer

Wendy is likely a sophisticated present-biased agent. She can see that if she pays off her credit card with her savings, she will likely spend in the future.

The term deposit and maxed-out credit card provide Wendy with a commitment device. She cannot access her savings for the next year or use the credit card. This prevents her from spending in the future. Her decision is sub-optimal compared to the option of paying off the card with her savings and not spending on her credit card, but she knows that in the future she will not be able to stick with that course of action. Although this comes at a cost - the higher interest rate - the commitment device was still the preferred option.

Yvonne may be a naive present-biased agent. She paid off her credit card debt with her savings, the apparently wise choice for that moment in time. However, Yvonne would not have foreseen that her present bias would lead her to spend in the future and make a time inconsistent decision. She would not foresee the need for a commitment device.

Alternatively, Yvonne may simply have a high exponential discount rate who prefers to spend at all times.

Chapter 56

Beliefs problems

56.1 Soccer

Your favourite soccer team is a terrible team that rarely wins. You know the following about your team:

- 5% of the time they win.
- When they win, they score at least two goals 60% of the time.
- When they lose, they score at least two goals 20% of the time.

a) You hear that they scored at least two goals in their most recent game. Use Bayes' rule to calculate the probability that they won.

Answer To answer this question, we need to use Bayes' rule. $P(\text{win}|2 \text{ goals}) = \frac{P(2 \text{ goals}|\text{win})P(\text{win})}{P(2 \text{ goals})}$ $= \frac{P(2 \text{ goals}|\text{win})P(\text{win})}{P(2 \text{ goals}|\text{win})P(\text{win}) + P(2 \text{ goals}|\text{lose})P(\text{lose})}$ $= \frac{0.6 \times 0.05}{0.6 \times 0.05 + 0.2 \times 0.95}$ = 0.136The team won with 13.6% probability.

b) Use natural frequencies to calculate the probability that they won.
Answer

Suppose there are 100 games.

That would mean there are 5 wins and 95 losses.

In those five wins, they would score at least two goals in three of them. In the 95 losses, they would score at least two goals in $0.2 \times 95 = 19$ of them. That means 3 of the 3 + 19 = 22 games in which they score two goals are wins. Therefore:

 $P(\text{win}|2 \text{ goals}) = \frac{\text{wins when they score two goals}}{\text{games where they score two goals}}$ $= \frac{3}{22}$

= 13.6%

56.2 Roulette

Imagine you are in a casino and observing a roulette wheel. The last five spins of the wheel have all resulted in the ball landing on a red number. Now, the wheel is about to spin again.

a) What is the probability that the next spin will be black?

? Answer

The probability that the next spin will be black is $\sim 50\%$. Each spin is independent of the previous spin. (Technically, the probability is slightly less than 50% because of the presence of the green 0 and, in the case of American roulette wheels, 00. I've ignored this for simplicity.)

b) Tommy is watching the same game of roulette. He suffers from the gambler's fallacy. What is he likely to believe? Give a potential explanation for his belief.

Answer

Tom is likely to believe that the next spin will be black. As he has the gambler's fallacy, he will believe that an outcome not recently realised in a sequence of independent draws - in this case, black - is more likely to occur on the next draw.

c) Sammy is at the next table watching a game of craps. He has seen a player win by rolling 7 three times in a row. The player makes the same bet and is about to roll again. Sammy suffers from the hot hand fallacy. What is he likely to believe? Why?

💡 Answer

Sammy is likely to believe that the next roll will be a 7. As he has the hot hand fallacy, he believes the streak will persist despite each outcome being independent of the last.

56.3 Buying a business

Dylan is considering buying a small business. The selling agent tells Dylan that the business is worth \$1.5 million. Dylan knows that the business is not worth that much.

a) Use concepts from the material on beliefs to explain why an agent might tell Dylan the business is worth so much.

? Answer

The agent is likely attempting to create an anchor. If Dylan uses the anchoring and adjustment heuristic in determining the price, he will start from that anchor and then inadequately adjust downwards. So despite knowing it is too high, it may still bias his estimate up.

b) To determine the value of the business, you estimate the revenue and profit the business will generate. Dylan accounts for his belief that he will be able to increase the revenue substantially from current levels. He determines a fair valuation is \$1.1 million.

He purchases the business. However, after running the business for several months, he does not realise the profit he anticipated.

What form of overconfidence is this? Explain. Under what conditions does this form of overconfidence tend to emerge?

? Answer

Dylan suffers from *overestimation*. Overestimation is the belief that we can perform at a level beyond that which we realistically can. In this case, Dylan did not generate the profit that he believed he could. People typically overestimate when attempting a difficult task.

c) When developing his valuation, Dylan determined there was a 95% probability that the business was worth between \$1.08 million and \$1.12 million.

The business ultimately sold for \$1 million.

What form of overconfidence is this? Explain.

Dylan suffers from *overprecision*. Overprecision is the tendency to believe that our predictions or estimates are more accurate than they are. Dylan's overprecision is indicated by the narrow confidence interval he gave for the value of the business.

56.4 A loaded dice

You have two six-sided dice.

- One die is fair with the numbers 1 through to 6 occurring with equal probability.
- The other die is loaded and always rolls an even number. It rolls a 2, 4 or 6 with equal probability.

You pull one die out of your pocket and roll it. You did not check which die it was before you rolled. (Assume you could have pulled either die out of your pocket with equal probability.)

a) The die shows a six. What is the probability that it is the loaded die? Calculate your answer using Bayes' rule.

$$P(D1 \text{ loaded}|6) = \frac{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}{P(6)}$$
$$= \frac{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}{P(6|D1 \text{ loaded})P(D1 \text{ loaded})}$$
$$= \frac{\frac{1}{3} \times 0.5}{\frac{1}{3} \times 0.5 + \frac{1}{6} \times 0.5}$$
$$= 0.67$$

The first die is the loaded die with 66.7% probability.

b) You pull the other die out of your pocket and roll it. It shows a 5. What is the updated probability that the first die you pulled out of your pocket is the loaded die? Calculate your answer using Bayes' rule.

$$P(D2 \text{ fair}|5) = \frac{P(5|D2 \text{ fair})P(D2 \text{ fair})}{P(5)}$$

= $\frac{P(5|D2 \text{ fair})P(D2 \text{ fair})}{P(5|D2 \text{ fair})P(D2 \text{ fair}) + P(5|D2 \text{ loaded})P(D2 \text{ loaded})}$
= $\frac{\frac{1}{6} \times \frac{2}{3}}{\frac{1}{6} \times \frac{2}{3} + 0 \times \frac{1}{3}}$
= 1

The first die is the loaded die with 100% probability. This should make intuitive sense. The second die showed a 5 and is, therefore, fair with 100% probability. The loaded die never shows a 5.

Chapter 57

Behavioural game theory problems

57.1 Airline competition

Quokka Airlines and Viper Air are competitor airlines. Each is considering how it should price its tickets. They have two options: a high price or a low price. Each must choose the price simultaneously.

If one airline offers a lower price than the other, they gain more market share but a lower profit margin. If both airlines offer the same price, Quokka takes more of the market as the incumbent airline.

The expected payoffs for each combination of actions are as follows, with the payoff (x, y) being the payoffs for Quokka and Viper respectively.

		Vip	ber
		High	Low
Quality	High	100, 20	40, 30
QUOKKA	Low	60, 0	50, 10

a) Are there any pure-strategy Nash equilibria? If so, what are they?

We determine the pure-strategy Nash equilibria by considering the best response of each player to each of the other player's strategies.

If Viper sets prices high, Quokka can choose high for a payoff of 100 or low for a payoff of 60. High is the best response.

If Viper sets prices low, Quokka can choose high for a payoff of 40 or low for a payoff of 50. Low is the best response.

If Quokka sets prices high, Viper can choose high for a payoff of 20 or low for a payoff of 30. Low is the best response.

If Quokka sets prices low, Viper can choose high for a payoff of 0 or low for a payoff of 10. Low is the best response.

The pure Nash equilibrium is therefore (low, low). Neither has an incentive to deviate.

		Viper		
		High	Low	
Quality	High	100 20	40,30	
QUOKKA	Low	60, 0	50,10	

b) Suppose the managers of Quokka and Viper are level-k thinkers.

If they were level-0, both would choose high or low with equal probability.

What would each player do if they were a level-1 thinker?

? Answer

A level-1 thinker assumes that the other player is a level-0 thinker. Each level-1 thinker plays the optimal strategy on this assumption.

A level-1 Quokka plays the optimal strategy against a level-0 Viper. A level-0 Viper plays high or low with equal probability. The payoffs to Quokka from each option are:

$$U_Q(\text{high}) = 0.5 \times 100 + 0.5 \times 40$$

= 70

$$U_Q(\text{low}) = 0.5 \times 60 + 0.5 \times 50$$

= 55

Quokka chooses high.

The payoffs to Viper from each option are:

$$U_V(\text{high}) = 0.5 \times 20 + 0.5 \times 0$$
$$= 10$$
$$U_V(\text{low}) = 0.5 \times 30 + 0.5 \times 10$$

$$= 20$$

Viper chooses low.

c) What would each player do if they were a level-2 thinker?

? Answer

A level-2 thinker assumes that the other player is a level-1 thinker. Each level-2 thinker plays the optimal strategy on this assumption.

A level-2 Quokka knows that the level-1 Viper will choose low. Therefore, the payoff from high is 40 and from low is 50. Quokka chooses low.

A level-2 Viper knows that the level-1 Quokka will choose high. Therefore, the payoff from high is 20 and from low is 30. Viper chooses low.

d) What would each player do if they were a level-3 or above thinker?

? Answer

At whatever level of thinking the players discover the Nash equilibrium, for any higher level of thinking they will remain at the Nash equilibrium. Therefore, at level-3 and above, both players will choose low.

e) Suppose now that the managers of Quokka and Viper are perfectly rational. Each is considering how they could make more profit.

A staff member in Viper suggests launching an advertising campaign with high prices before they and Quokka choose. Once Viper has advertised their prices, they will be constrained to setting them high.

Is this a good idea? What concept does this illustrate?

🔮 Answer

Suppose Viper advertises high prices and is committed to that choice. The game therefore becomes:



With this game, Viper plays high no matter what choice is made by Quokka. They effectively have no choice.

Quokka can choose high for a payoff of 100 or low for a payoff of 60. High is the best response.



The new Nash equilibrium is (high, high). Viper has increased its payoff from 10 to 20 by making the strategic move. The move is therefore a good idea. This example illustrates the concept of *commitment* via a strategic move.

57.2 Valuing a business

You are considering buying a small business. You research the business, attempting to estimate the revenue and profit it will generate. Based on your research, you determine a fair valuation is \$850,000.

The business ultimately sells for \$1.1 million to another person. You discover that many other parties were interested in buying the business.

You later hear that the purchaser is disappointed with the business and that it is worth less than they paid. What phenomenon could this be an example of? Explain.

This is an example of the *winner's curse*.

The winner's curse occurs when the winner of an auction is the bidder who most overestimates the value of the item. The winner therefore pays more than the item is worth.

57.3 Cafe

Two friends, Player 1 and Player 2, have arranged to meet at a cafe. Neither can remember which of the two cafes in town they had arranged to meet at.

Each player has a favourite cafe. However, they would prefer to go to their less-favourite cafe with a friend than go to their favourite cafe alone.

Each chooses a cafe and goes there. The payoffs for each combination of choices are in the table below, with the payoffs (x, y) being the payoffs for Player 1 and Player 2 respectively.

		Player 2	
		Café A	Café B
Disver 1	Café A	2, 1	0, 0
Player 1	Café B	0, 0	1, 3

a) What are the pure-strategy Nash equilbria of this game?

? Answer				
If player 2 of If player 2 of If player 1 of If player 1 of	chooses cafe chooses cafe chooses cafe chooses cafe	A, player 1 B, player 1 A, player 1 B, player 1	wants to ch wants to ch wants to ch wants to ch	oose cafe A. oose cafe B. oose cafe A. oose cafe B.
Player 2				
		Café A	Café B	
Diaman 1	Café A	2,1	0, 0	
Player 1	Café B	0, 0	13	
The two pu	re-strategy l	Nash equilib	ria are (cafe	A, cafe A) and (cafe B, cafe B).

b) What would Player 1 and Player 2 do if each was a level-1 thinker? Explain.

🔮 Answer

If Player 2 chooses randomly, Player 1's payoff's from each option are:

$$U_1(\text{cafe A}) = 0.5 \times 2 + 0.5 \times 0 = 1$$

$$U_1$$
(cafe B) = $0.5 \times 0 + 0.5 \times 1 = 0.5$

Player 1 chooses cafe A.

If Player 1 chooses randomly, Player 2's payoff's from each option are:

$$U_2(\text{cafe A}) = 0.5 \times 1 + 0.5 \times 0 = 0.5$$

$$U_2(\text{cafe B}) = 0.5 \times 0 + 0.5 \times 3 = 1.5$$

Player 2 chooses cafe B.

c) What would Player 1 and Player 2 do if each was a level-2 thinker? Explain.

? Answer

The level-2 Player 1 assumes that Player 2 is level-1. Therefore, they assume that Player 2 will choose cafe B. Accordingly, they choose cafe B. The level-2 Player 2 assumes that Player 1 is level-1. Therefore, they assume that Player 1 will choose cafe A. Accordingly, they choose cafe A.

57.4 A military attack

An army from the North is about to attack the South.

The North can attack one of two cities: Hobart or Launceston. Launceston is easier to attack as it is closer.

The South needs to decide which city it will plan to defend.

If the North attacks an undefended city, it will win. The South can repel any attack on a city it has chosen to defend.

The expected payoffs for each combination of actions are as follows, with the payoff (x, y) being the payoffs for the North and South respectively.

		So	uth
		Hobart	Launceston
North	Hobart	-1, 4	4, 0
NOTUT	Launceston	4, 0	0, 4

a) Are there any pure-strategy Nash equilibria? If so, what are they?

? Answer

We determine the pure-strategy Nash equilibria by considering the best response of each player to each of the other player's strategies.

If the South defends Hobart, North can choose Hobart for a payoff of -1 or Launceston for a payoff of 4. Launceston is the best response.

If the South defends Launceston, North can choose Hobart for a payoff of 4 or Launceston for a payoff of 0. Hobart is the best response.

If the North attacks Hobart, South can defend Hobart for a payoff of 4 or Launceston for a payoff of 0. Hobart is the best response.

If the North attacks Launceston, South can defend Hobart for a payoff of 0 or Launceston for a payoff of 4. Launceston is the best response.

There is no pure-strategy Nash equilibrium. For any combination of choices, one of the armies has an incentive to change their choice.

		South		
		Hobart	Launceston	
North	Hobart	-14	4,0	
North	Launceston	4,0	04	

b) Suppose the commanders of the North and South are level-k thinkers.

If they were level-0, both would choose Hobart or Launceston with equal probability. What would each player do if they were a level-1 thinker?

A level-1 thinker assumes that the other player is a level-0 thinker. Each level-1 thinker plays the optimal strategy on this assumption.

A level-1 North plays the optimal strategy against a level-0 South. A level-0 South plays Hobart or Launceston with equal probability. The payoffs to North from each option are:

$$U_N(\text{Hobart}) = 0.5 \times -1 + 0.5 \times 4 = 1.5$$

$$U_N$$
(Launceston) = $0.5 \times 4 + 0.5 \times 0 = 2$

North attacks Launceston.

$$U_{\rm S}({\rm Hobart}) = 0.5 \times 4 + 0.5 \times 0 = 2$$

$$U_{\rm S}({\rm Launceston}) = 0.5 \times 4 + 0.5 \times 0 = 2$$

South is indifferent between defending Hobart and Launceston. They can choose either.

c) What would each player do if they were a level-2 thinker?

Answer

A level-2 thinker assumes that the other player is a level-1 thinker. Each level-2 thinker plays the optimal strategy on this assumption.

A level-2 North knows that the level-1 South is indifferent between defending Hobart and Launceston. If North assumes that South will defend each with equal probability, the payoffs to North from each option are:

 $U_N(\text{Hobart}) = 0.5 \times -1 + 0.5 \times 4 = 1.5$

 $U_{\rm \textit{N}}({\rm Launceston}) = 0.5 \times 4 + 0.5 \times 0 = 2$

North attacks Launceston.

A level-2 South knows that the level-1 North attacks Launceston. The South defends Launceston for payoff of 4 (rather than Hobart for payoff of 0).

57.5 Acquiring a company

The following example draws on Bazerman and Moore (2013).

Company A is considering acquiring Company T.

The value of Company T depends on the outcome of an oil exploration project. If the project fails, the company under current management will be worth nothing (\$0 per share). If the project succeeds, the value of the company under current management could be as high as \$100 per share. All values between \$0 and \$100 are equally likely.

Company T will be worth 50 per cent more in the hands of Company A than under current management. If the project fails, the company will be worth \$0 per share under either management. If the exploration project generates a \$50 per share value under current management, the value under Company A will be \$75 per share. And so on.

Company A is considering what price per share they should offer. This offer must be made before Company A knows the outcome of the drilling project, but after Company T learns the result. Company T will accept any offer from Company A if it is profitable for them.

a) Show that for any offer above zero Company A expects to lose money.

Answer

If Company A offers x, Company T will accept x% of the time, whenever the firm is worth between \$0 and \$x. Since all those values are equally likely, the firm will be worth on average x/2 to company T when they accept. The shares will therefore be worth $1.5 \times x/2 = 3x/4$ on average for company A. That gives Company A profit of:

$$\begin{aligned} \pi_A &= \frac{3x}{4} - x \\ &= -\frac{x}{4} \end{aligned}$$

Any offer above 0 generates a negative expected return, a loss of 25% of the offer.

To give an example, if Company A offered \$60, it will be accepted 60% of the time - whenever the firm is worth between \$0 and \$60 for company T. Since all those values are equally likely, the firm will be worth on average \$30 to company T when they accept, meaning it will be worth \$45 on average for company A. A \$60 offer will result in an average loss of \$15.

b) People given this problem tend to bid between \$50 and \$75 per share (Samuelson and Bazerman (1985)). A typical explanation offered by these people is that the average outcome for Company T is \$50, making the value for Company A \$75. Any offer in the range between these two values would be agreeable to both parties.

Explain why a "cursed" player representing Company A might make a non-zero offer.

Answer

A "cursed" player representing company A does not think that company T's decision to sell depends on company T's knowledge of the oil exploration. As a result, they are likely to bid based on their unconditional expected value of the field, not the value conditional on acceptance.

This bidding approach leads to, on average, a loss. The cursed player underappreciates that company T is more likely to accept when company T's valuation is low.

57.6 Receiving a faulty product

A customer received a faulty product from a firm and requested a refund as per consumer law. The customer also threatened to complain to the Department of Fair Trading if they did not receive the refund. A customer complaint would be costly to the firm as they would be required to provide a refund in addition to incurring the cost of dealing with the complaint.

The firm offered a store credit instead, believing that the customer would not complain as the time and effort involved would not be worth the potential refund.

However, the customer still complained to the Department of Fair Trading.

a) Use concepts from game theory to explain why the firm might have held that belief.



b) Use concepts from behavioural game theory to explain why the firm's belief was ultimately incorrect.

? Answer

The customer's emotional response may lead them to complain. They might be angry or obtain satisfaction from seeing the firm punished. In that case, the

customer will complain even though it is not in their material best interest to do so. Emotionally, it is worthwhile. They incur the cost of complaining but get the benefit of both the refund and the satisfaction from punishing the firm.

References

- Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism*. The Quarterly Journal of Economics, 84(3), 488–500. https://doi. org/10.2307/1879431
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica*, 21(4), 503–546. https://doi.org/10.2307/1907921
- Andreoni, J., and Bernheim, B. D. (2009). Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica*, 77(5), 1607–1636. https://doi.org/10.3982/ECTA7384
- Ashraf, N., Karlan, D., and Yin, W. (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines^{*}. The Quarterly Journal of Economics, 121(2), 635–672. https://doi.org/10.1162/qjec.2006.121.2.635
- Ayton, P., and Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32(8), 1369–1378. https://doi.org/10.3758/BF03206327
- Barber, B. M., and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment*. The Quarterly Journal of Economics, 116(1), 261–292. https://doi.org/10.1162/003355301556400
- Barberis, N., Huang, M., and Thaler, R. H. (2006). Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing. American Economic Review, 96(4), 1069–1090. https://doi.org/10.1257/aer.96.4.1069
- Bar-Eli, M., Avugos, S., and Raab, M. (2006). Twenty years of "hot hand" research: Review and critique. *Psychology of Sport and Exercise*, 7(6), 525–553. https: //doi.org/10.1016/j.psychsport.2006.03.001
- Bazerman, M. H., and Moore, D. A. (2013). Judgment in managerial decision making (8. ed). Wiley.
- Beshears, J., Choi, J. J., Harris, C., Laibson, D., Madrian, B. C., and Sakong, J. (2020). Which early withdrawal penalty attracts the most deposits to a commitment savings account? *Journal of Public Economics*, 183, 104144. https: //doi.org/10.1016/j.jpubeco.2020.104144
- Camerer, C., Babcock, L., Loewenstein, G., and Thaler, R. (1997). Labor supply of new york city cabdrivers: One day at a time*. *The Quarterly Journal of Economics*, 112(2), 407–441. https://doi.org/10.1162/003355397555244
- Camerer, C., and Lovallo, D. (1999). Overconfidence and Excess Entry: An Experimental Approach. The American Economic Review, 89(1), 13.
- Card, D., Mas, A., Moretti, E., and Saez, E. (2012). Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. American Economic Review, 102(6), 2981–3003.

https://doi.org/10.1257/aer.102.6.2981

- Charness, G., and Rabin, M. (2002). Understanding social preferences with simple tests. The Quarterly Journal of Economics, 117(3), 817–869. https://www.jstor. org/stable/4132490
- Cosmides, L., and Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition*, 58(1), 1–73. https://doi.org/10.1016/0010-0277(95)00664-8
- Czerlinski, J., Gigerenzer, G., and Goldstein, D. G. (1999). How good are simple heuristics. In G. Gigerenzer, P. Todd, and The ABC Research Group (Eds.), Simple heuristics that make us smart. Oxford University Press.
- Davidson, D., McKinsey, J. C. C., and Suppes, P. (1955). Outlines of a Formal Theory of Value, I. Philosophy of Science, 22(2), 140–160. https://doi.org/10.1086/287412
- De Bondt, W. F. M., and Thaler, R. H. (1995). Chapter 13 Financial decision-making in markets and firms: A behavioral perspective. In R. A. Jarrow, V. Maksimovic, and W. T. Ziemba (Eds.), *Handbooks in Operations Research and Management Science* (Vol. 9, pp. 385–410). Elsevier. https://www.sciencedirect.com/science/ article/pii/S092705070580057X
- Engel, C. (2011). Dictator games: a meta study. *Experimental Economics*, 14(4), 583–610. https://doi.org/10.1007/s10683-011-9283-7
- Farber, H. S. (2005). Is tomorrow another day? The labor supply of new york city cabdrivers. Journal of Political Economy, 113(1), 46–82. https://doi.org/10.1086/ 426040
- Farber, H. S. (2008). Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers. American Economic Review, 98(3), 1069–1082. https://doi.org/10.1257/aer.98.3.1069
- Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. The Quarterly Journal of Economics, 130(4), 1975–2026. https://doi.org/10.1093/qje/qjv026
- Fehr, E., and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. The Quarterly Journal of Economics, 114(3), 817–868. https://www.jstor. org/stable/2586885
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. Journal of Economic Literature, 40(2), 351– 401. https://www.jstor.org/stable/2698382
- Genesove, D., and Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. The Quarterly Journal of Economics, 116(4), 1233–1260. https://www.jstor.org/stable/2696458
- Gigerenzer, G. (2011). What are natural frequencies? *BMJ*, 343, d6386. https://doi.org/10.1136/bmj.d6386
- Gigerenzer, G. (2021). Embodied heuristics. Frontiers in Psychology, 12. https://www.frontiersin.org/articles/10.3389/fpsyg.2021.711289
- Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3), 295–314. https://doi.org/10.1016/0010-0285(85)90010-6
- Giné, X., Karlan, D., and Zinman, J. (2010). Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation. American Economic Journal: Applied Economics, 2(4), 213–235. https://doi.org/10.1257/app.2.4.213
- Green, L., Fristoe, N., and Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review*, 1(3), 383–389. https://doi.org/10.3758/BF03213979

- Heath, C., Larrick, R. P., and Wu, G. (1999). Goals as Reference Points. Cognitive Psychology, 38(1), 79–109. https://doi.org/10.1006/cogp.1998.0708
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., and McElreath, R. (2001). In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. American Economic Review, 91(2), 73–78. https://doi.org/10.1257/aer. 91.2.73
- Hertwig, R., and Gigerenzer, G. (1999). The 'conjunction fallacy' revisited: how intelligent inferences look like reasoning errors. Journal of Behavioral Decision Making, 12(4), 275–305. https://doi.org/10.1002/(SICI)1099-0771(199912)12:4% 3C275::AID-BDM323%3E3.0.CO;2-M
- Hoffrage, U., and Gigerenzer, G. (1998). Using natural frequencies to improve diagnostic inferences. Academic Medicine, 73(5), 53840. https://doi.org/10.1097/ 00001888-199805000-00024
- Hoffrage, U., Krauss, S., Martignon, L., and Gigerenzer, G. (2015). Natural frequencies improve bayesian reasoning in simple and complex inference tasks. *Frontiers* in Psychology, 6, 1473. https://doi.org/10.3389/fpsyg.2015.01473
- Johnson, E. J., and Goldstein, D. (2003). Do Defaults Save Lives? Science, 302(5649), 1338–1339. https://doi.org/10.1126/science.1091721
- Johnson, N. D., and Mislin, A. A. (2011). Trust games: A meta-analysis. Journal of Economic Psychology, 32(5), 865–889. https://doi.org/10.1016/j.joep.2011.05.007
- Kahneman, D. (2011). Thinking, fast and slow (1st edition). Farrar, Straus; Giroux.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1991). Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *Journal of Economic Perspectives*, 5(1), 193–206. https://doi.org/10.1257/jep.5.1.193
- Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291. https://doi.org/10.2307/1914185
- Kahneman, D., and Tversky, A. (1984). Choices, values, and frames. American Psychologist, 39(4), 341–350. https://doi.org/10.1037/0003-066X.39.4.341
- Keynes, J. M. (1936). The general theory of employment, interest, and money. Macmillan. http://gutenberg.net.au/ebooks03/0300071h/printall.html
- Kirby, K. N., and Herrnstein, R. J. (1995). Preference Reversals Due to Myopic Discounting of Delayed Reward. *Psychological Science*, 6(2), 83–89. https://doi. org/10.1111/j.1467-9280.1995.tb00311.x
- Kirgios, E. L., Mandel, G. H., Park, Y., Milkman, K. L., Gromet, D. M., Kay, J. S., and Duckworth, A. L. (2020). Teaching temptation bundling to boost exercise: A field experiment. Organizational Behavior and Human Decision Processes, 161, 20–35. https://doi.org/10.1016/j.obhdp.2020.09.003
- Kőszegi, B., and Rabin, M. (2006). A model of reference-dependent preferences. The Quarterly Journal of Economics, 121(4), 1133–1165. https://www.jstor.org/ stable/25098823
- Martin, V. (2017). When to quit: Narrow bracketing and reference dependence in taxi drivers. Journal of Economic Behavior & Organization, 144, 166–187. https: //doi.org/10.1016/j.jebo.2017.09.024
- Miller, J. B., and Sanjurjo, A. (2018). Surprised by the Hot Hand Fallacy? A Truth in the Law of Small Numbers. *Econometrica*, 86(6), 2019–2047. https://doi.org/ 10.3982/ECTA14943
- Moore, D. A., Oesch, J. M., and Zietsma, C. (2007). What Competition? Myopic Self-Focus in Market-Entry Decisions. Organization Science, 18(3), 440–454. https: //doi.org/10.1287/orsc.1060.0243
- Moulin, H. (1986). Game theory for social sciences. New York Press.

- Nagel, R. (1995). Unraveling in guessing games: An experimental study. The American Economic Review, 85(5), 1313–1326. https://www.jstor.org/stable/2950991
- Newton, E. L. (1990). The rocky road from actions to intentions [PhD thesis]. https://www.proquest.com/openview/b740253d9b78599786f59d6b6055cc3b/1?pqorigsite=gscholar&cbl=18750&diss=y
- Page, L. (2022). Optimally Irrational: The Good Reasons We Behave the Way We Do. Cambridge University Press. https://www.cambridge.org/au/academic/subjects/economics/microeconomics/optim irrational-good-reasons-we-behave-way-we-do, https://www.cambridge.org/au/academic/subjects/economics/micro
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 66(3), 497–527. https://doi.org/10.2307/2998573
- Rabin, M. (2000). Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica*, 68(5), 1281–1292. http://www.jstor.org/stable/2999450
- Rabin, M. (2002). Inference by believers in the law of small numbers. The Quarterly Journal of Economics, 117(3), 775–816. https://doi.org/10.1162/ 003355302760193896
- Rabin, M., and Thaler, R. H. (2001). Anomalies: Risk Aversion. Journal of Economic Perspectives, 15(1), 219–232. https://doi.org/10.1257/jep.15.1.219
- Rabin, M., and Vayanos, D. (2010). The gambler's and hot-hand fallacies: Theory and applications. The Review of Economic Studies, 77(2), 730–778. https://doi. org/10.1111/j.1467-937X.2009.00582.x
- Rapoport, A., and Budescu, D. V. (1997). Randomization in individual choice behavior. Psychological Review, 104, 603–617. https://doi.org/10.1037/0033-295X.104. 3.603
- Read, D., and Leeuwen, B. van. (1998). Predicting Hunger: The Effects of Appetite and Delay on Choice. Organizational Behavior and Human Decision Processes, 76(2), 189–205. https://doi.org/10.1006/obhd.1998.2803
- Samuelson, W. F., and Bazerman, M. (1985). Negotiation under the winner's curse (V. Smith, Ed.). JAI Press.
- Savant, M. vos. (1990). Game show problem. *PARADE Magazine*, 16. https://web.archive.org/web/20130121183432/http://marilynvossavant.com/game-show-problem/
- Shefrin, H., and Statman, M. (1985). The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. The Journal of Finance, 40(3), 777–790. https://doi.org/10.1111/j.1540-6261.1985.tb05002.x
- Spiegelhalter, D., and Gage, J. (2015). What can education learn from real-world communication of risk and uncertainty? *The Mathematics Enthusiast*, 12(1), 4–10. https://doi.org/10.54870/1551-3440.1329
- Thaler, R. (1980). Toward a positive theory of consumer choice. Journal of Economic Behavior & Organization, 1(1), 39–60. https://doi.org/10.1016/0167-2681(80) 90051-7
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201–207. https://doi.org/10.1016/0165-1765(81)90067-7
- Thaler, Richard H., and Benartzi, S. (2004). Save More TomorrowTM: Using Behavioral Economics to Increase Employee Saving. *Journal of Political Economy*, 112(S1), S164–S187. https://doi.org/10.1086/380085
- Tversky, A., and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. Science, 185(4157), 1124–1131. https://doi.org/10.1126/science.185.4157. 1124
- Tversky, A., and Kahneman, D. (1982). Evidential impact of base rates. In A. Tversky, D. Kahneman, and P. Slovic (Eds.), Judgment under uncer-

tainty: Heuristics and biases (pp. 153–160). Cambridge University Press. https://www.cambridge.org/core/books/judgment-under-uncertainty/evidential-impact-of-base-rates/CC35C9E390727085713C4E6D0D1D4633

- Tversky, A., and Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315. https://doi.org/10.1037/0033-295X.90.4.293
- Tversky, A., and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 5, 297–323. https: //doi.org/10.1007/BF00122574
- Zelmer, J. (2003). Linear Public Goods Experiments: A Meta-Analysis. Experimental Economics, 6(3), 299–310. https://doi.org/10.1023/A:1026277420119