ST114 Decisions and Games

Adam M. Johansen

a.m.johansen@warwick.ac.uk

Based on an earlier version by Prof. Wilfrid Kendall

University of Warwick — Winter 2009

Administrative Details

Lecturer	Adam Johansen
email	`a.m.johansen@warwick.ac.uk`
Office	C0.20 (Maths and Statistics)
Office hours	Tuesday & Wednesday 11:30 – 12:30
Telephone	024 761 - 50919
Lectures	20 (approximately. . .) Tuesday 16:00 Friday 13:00
CATS	7.5
Assessment	100% Closed Book Examination

Aims

- To give an introduction into how the use of probabilistic and mathematical ideas can enhance decision making by providing a framework in which actions can be judged as sensible or irrational.
- Examples will be given both of games against nature and games against other rational opponents.

Objectives

- The student will be taught some of the arguments underpinning the use of rationality and a definition of subjective probability.
- They will be taught how to use the simpler tools of decision analysis as a framework to discover sensible decision rules which balance quantified uncertainties and payoffs.
- The course will explain and illustrate some of the issues of rationality as they apply to games and techniques will be given which will enable the student to solve some simple zero sum games.

Syllabus

Ideas to be presented will include:

- The quantification of subjective belief through probability.
- The EMV decision rule.
- The quantification of subjective preferences.
- The concept of a rational opponent in a two player game.

The course aims to

- Provide an insight into various applications of mathematical concepts.
- Inform students how they might ensure that their own decision-making is coherent and rational.

Detailed Syllabus

1. Introduction
2. Axiomatic Probability
3. What is Probability
4. Conditional Probability
5. Decisions
6. Preferences and Objectives
7. Games

Books

There are a great number of books of the subjects of this course…

- You don't *need* to buy any of them.
- Many are available in the library.
- Jim Smith has kindly made copies of his "Decision Analysis: A Bayesian Approach" available at cost price ($\sim$ £3.50) from Hilda Cooper's office.
- James Berger's "Statistical Decision Theory and Bayesian Analysis" is a good reference but goes way beyond the scope of this course.
- Dover republishes many classics, including:
  - Thomas' "Games, Theory and Applications"
  - Luce and Raiffa's "Games and Decisions"

Introduction

The Problem of the Decision Analyst

This stylised scenario embodies the core problems of decision analysis:

- You have a client[1].
- The client must choose one action from a set of possibilities.
- This client is uncertain about many things, including:
  - Her priorities.
    Conflicting requirements can be difficult to resolve.
  - What might happen.
    Fundamental uncertainty – things not within her control.
  - How other people may act.
    Other interested parties might influence the outcome.
- You must advise this client on the best course of action.

[1] This may be yourself, but it is useful to separate the two rôles.

A problem of two parts

- Elicitation: Obtain precise answers to several questions:
  - What is the client's problem?
  - what does she believe?
  - What does she want?
- Calculation: Given this information
  - What are its logical implications?
  - What should our client do?

Elicitation $\longrightarrow$ Calculation $\longrightarrow$ Elicitation $\longrightarrow$ Calculation $\longrightarrow \dots$

What does she really want?

Example (Advising a university undergraduate)

What is their objective?

- Getting the best possible degree?
- Trying to get a particular job after university?
- Learning for its own sake?
- Having as much fun as possible?
- A combination of the above?

Example (A small business owner)

What is their objective?

- Staying in business?
- Making £X of profit in as short a time as possible?
- Making as much profit as possible in time T?
- Eliminating competition?
- Maximising growth?

What does she know?

As well as knowing what our client *wants* we need to know what they *know*:

- What are their options?
- What are the possible consequences of these actions?
- How are the consequences related to the action taken?
- Are any other parties involved? If so, what are their objectives?

Example (Marketing)

- How can we advertise?
- What are the *costs* of different approaches?
- What are the *effects* of these approaches?
- What volume of production is possible?
- What competition do we have?

Example (Insurance)

Insurance against a particular type of loss...

- Probability of the loss occurring is $p \ll 1$.
- Cost of that lost would be, say, $£5,000$.
- Insurance premium is $£10$.

Why are both parties happy with this?

Example (A Simple Lottery)

- $\mathbb{P}(\{\text{Win}\}) = 1/10,000$
- $Value(\text{Win}) = £5,000$
- Ticket price $£1$.

Why is this acceptable? What about simple variations?

Is that *really* what she believes?

It is important to distinguish between that which is *believed* from that which is *hoped*, *feared* or simply asserted.

Example (Economic forecasting)

Recent forecasts of British GDP growth in 2009:

- -0.1% – International Monetary Fund
- -0.75– -1.25% British Government
- -1.1% Organisation for Economic Co-operation and...
- -1.7% Confederation of British Industry
- -2.9% Centre for Economics and Business research

Each organisation has different objectives & knowledge.
Are they necessarily reliable indications of the underlying beliefs of these organisations[2]?

[2] We will put aside the philosophical questions raised by this concept...

Quantification of Subjective Knowledge

Our client has beliefs and some idea about her objective. She probably isn't a mathematician. We have to codify things in a rigorous mathematical framework.

In particular, we must be able to encode:

- Beliefs about what can happen and how likely those things are to happen.
- The cost or reward of particular outcomes.
- In the case of games: What any other interest parties want and how they are likely to react.

Having done this, we must use our mathematical skills to work out how to advise our client.

Some Terminology

Before considering details, we should make sure we agree about terminology.

- In a *decision problem* we have:
  - A (random) source of uncertainty.
  - A collection of possible *actions*.
  - A collection of *outcomes*.

  and we wish to choose the action to obtain a favourable outcome.
- A *game* is a similar problem in which the uncertainty arises from the behaviour of a (rational) opponent.

From Questions to Answers

Now we need to answer some questions:

1. How can be elicit and quantify beliefs?
2. How can we represent their particular problem mathematically?
3. How do we represent her objectives quantitatively?
4. What should we advise our client to do?
5. What can we do if other rational agents are involved?

We will begin by answering question 1: we can use probability.

Probability

Foundations of An Axiomatic Theory of Probability

The *Russian school* of probability is based on axioms.
The abstract specification of probability requires three things:

1. A set of all possible outcomes, Ω.

   The *sample space* containing elementary events.

2. A collection of subsets of Ω, $\mathcal{F}$.

   Outcomes of interest.

3. A function which assigns a probability to our events: $\mathbb{P} : \mathcal{F} \to [0, 1]$

   The probability itself.

Example (Simple Coin-Tossing)

- All possible outcomes might be:

$$\Omega = \{H, T\}.$$

- And we might be interested in all possible subsets of these outcomes:

$$\mathcal{F} = \{\emptyset, \{H\}, \{T\}, \Omega\}.$$

- In which case, under reasonable assumptions:

$$\mathbb{P}(\emptyset) = 0 \qquad \mathbb{P}(\{H\}) = \frac{1}{2}$$
$$\mathbb{P}(\{T\}) = \frac{1}{2} \qquad \mathbb{P}(\{H, T\}) = 1$$

Example (A Tetrahedral (4-faced) Die)

- The possible outcomes are: $\Omega = \{1, 2, 3, 4\}$
- And we might again consider all possible subsets:

$$
\begin{aligned}
\mathcal{F} = \{ \quad & \emptyset, & \{1\}, && \{2\}, && \{3\}, \\
& \{4\}, & \{1,2\}, && \{1,3\}, && \{1,4\}, \\
& \{2,3\}, & \{2,4\}, && \{3,4\}, && \{1,2,3\}, \\
& \{1,2,4\}, & \{1,3,4\}, && \{2,3,4\}, && \{1,2,3,4\}\}
\end{aligned}
$$

- In this case, we might think that, for any $A \in \mathcal{F}$:

$$\mathbb{P}(A) = |A|/|\Omega| = \frac{\text{Number of values in } A}{4}$$

Example (The National Lottery)

- $\Omega = \{$All unordered sets of 6 numbers from$\{1, \ldots, 49\}\}$
- $\mathcal{F} =$ All subsets of Ω
- Again, we can construct $\mathbb{P}$ from expected uniformity.

- But there are $\binom{49}{6} = 13983816$ elements of Ω and consequently $2^{13983816} \approx 6 \times 10^{6000000}$ subsets!
- Even this simple discrete problem has produced an object of incomprehensible vastness.
- What would we do if $\Omega = \mathbb{R}$?
- It's often easier not to work with *all* of the subsets of Ω.

Algebras of Sets

Given Ω, $\mathcal{F}$ must satisfy certain conditions.

1. $\Omega \in \mathcal{F}$

   The event "something happening" is in our set.

2. If $A \in \mathcal{F}$, then

$$\Omega \setminus A = \{x \in \Omega : x \notin A\} \in \mathcal{F}$$

   If A happening is in our set then A not happening is too.

3. If $A, B \in \mathcal{F}$ then

$$A \cup B \in \mathcal{F}$$

   If event A and event B are both in our set then an event corresponding to either A or B happening is too.

A set that satisfies these conditions is called an *algebra* (over Ω).

σ-Algebras of Sets

If, in addition to meeting the conditions to be an algebra, $\mathcal{F}$ is such that:

- If $A_1, A_2, \dots \in \mathcal{F}$ then $\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}$

  If any countable sequence of events is in our set, then the event corresponding to any one of those events happening is too.

then $\mathcal{F}$ is known as a σ-algebra.

Example (Selling a house)

- You wish to sell a house, for at least £250,000.
- On Monday you receive an offer of X.
- You must accept or decline this offer immediately.
- On Tuesday you will receive an offer of Y.
- What should you do?

- $\Omega = \{(x, y) : x, y \geq £100,000\}$
- But, we only care about events of the form:

$$\{(i, j) : i < j\} \text{ and } \{(i, j) : i > j\}$$

- Including some others ensures that we have an algebra:

$$\{(i, j) : i = j\} \quad \{(i, j) : i \neq j\} \quad \{(i, j) : i \leq j\} \quad \{(i, j) : i \geq j\} \quad \emptyset \quad \Omega$$

Atoms

Some events are *indivisible* and somehow fundamental: An event $E \in \mathcal{F}$ is said to be an atom of $\mathcal{F}$ if:

1. $E \neq \emptyset$
2. $\forall A \in \mathcal{F}$:

$$E \cap A = \begin{cases} \emptyset \\ \text{or } E \end{cases}$$

Any element of $\mathcal{F}$ contains all of E or none of E.

If $\mathcal{F}$ is finite then any $A \in \mathcal{F}$, we can write:

$$A = \bigcup_{i=1}^{n} E_i$$

for some finite number, n, and atoms E_i of $\mathcal{F}$.

We can represent any event as a combination of atoms.

Example (Selling a house. . .)

Here, our algebra contained:

$$\begin{array}{cccc} \{(i,j) : i < j\} & \{(i,j) : i > j\} & \{(i,j) : i \neq j\} & \emptyset \\ \{(i,j) : i \leq j\} & \{(i,j) : i \geq j\} & \{(i,j) : i = j\} & \Omega \end{array}$$

Which of these sets are atoms?

- $\{(i,j) : i < j\}$ **is**
- $\{(i,j) : i > j\}$ **is**
- $\{(i,j) : i \neq j\}$ **is not** – it's the union of two atoms
- $\emptyset$ **is not** $\emptyset$ is never an atom
- $\{(i,j) : i = j\}$ **is**
- $\{(i,j) : i \leq j\}$ **is not** – it's the union of two atoms
- $\{(i,j) : i \geq j\}$ **is not** – it's the union of two atoms
- Ω **is not** – it's the union of three atoms

The Axioms of Probability – Finite Spaces

$\mathbb{P} : \mathcal{F} \to \mathbb{R}$ is a probability measure over $(\Omega, \mathcal{F})$ iff:

1. For any $A \in \mathcal{F}$:
$$\mathbb{P}(A) \geq 0$$
All probabilities are positive.
2. $$\mathbb{P}(\Omega) = 1$$
Something certainly happens.
3. For any[3] $A, B \in \mathcal{F}$ such that $A \cap B = \emptyset$:
$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$$
Probabilities are (sub)additive.

[3] This is sufficient if Ω is finite; we need a slightly stronger property in general.

The Axioms of Probability – General Spaces [see ST213]

$\mathbb{P} : \mathcal{F} \to \mathbb{R}$ is a probability measure over $(\Omega, \mathcal{F})$ iff:

1. For any $A \in \mathcal{F}$:
$$\mathbb{P}(A) \geq 0$$
All probabilities are positive.
2. $$\mathbb{P}(\Omega) = 1$$
Something certainly happens.
3. For any $A_1, A_2, \dots \in \mathcal{F}$ such that $\forall i \neq j : A_i \cap A_j = \emptyset$:
$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i).$$
Probabilities are countably (sub)additive.

Measures and Masses

- A *measure* tells us "how big" a set is [see MA359/ST213].
- A *probability measure* tells us "how big" an event is in terms of the likelihood that it happens [see ST213/ST318].
- In discrete spaces probability mass functions are often used.

Definition (Probability Mass Function)

If $\mathcal{F}$ is an algebra containing finitely many atoms $E_1, \ldots, E_n$. A *probability mass function*, f, is a function defined for every atom as $f(E_i) = p_i$ with:

- $p_i \in [0, 1]$
- and $\sum_{i=1}^{n} p_i = 1$.

Masses to Measures

- Let $S = \{A_1, \ldots, A_n\}$ be such that:
  - $\forall i \neq j : A_i \cap A_j = \emptyset$

    The elements of S are disjoint.
  - $\cup_{i=1}^{n} A_i = \Omega$

    S *covers* Ω.
- We can construct a finite algebra, $\mathcal{F}$ which contains the 2^n sets obtained as finite unions of elements of S.

  This algebra is *generated* by S.
- The atoms of the generated algebra are the elements of S.
- A mass function f on the elements of S defines a probability measure on $(\Omega, \mathcal{F})$:

$$\mathbb{P}(B) = \sum f(A_i)$$

(the sum runs over those atoms A_i which are contained in B).

So what?

So far we've seen:

- A mathematical framework for dealing with probabilities.
- A way to construct probability measures from the probabilities of every elementary event in a discrete problem.
- A way to construct probability measures from the probability mass function of a complete set of atoms.

But this doesn't tell us:

- What probabilities really mean.
- How to assign probabilities to *real* events... dice aren't everything!
- Why we should use probability to make decisions.

Geometry, Symmetry and Probability

- If probabilities have a geometric interpretation, we can often deduce probabilities from symmetries.

Example (Coin Tossing Again)

- Here, $\Omega = \{H, T\}$ and $\mathcal{F} = \{\emptyset, \{H\}, \{T\}, \{H, T\}\}$
- Axiomatically: $\mathbb{P}(\Omega) = P(\{H, T\}) = 1$.
- The atoms are $\{H\}$ and $\{T\}$.
- Symmetry arguments suggest that $\mathbb{P}(\{H\}) = \mathbb{P}(\{T\})$.

  Implicitly, we are assuming that the symbol on the face of a coin does not influence its final orientation.

- Axiomatically: $\mathbb{P}(\{H, T\}) = \mathbb{P}(\{H\}) + \mathbb{P}(\{T\})$.
- Therefore: $\mathbb{P}(\{H\}) = \mathbb{P}(\{T\}) = 1/2$.

Example (Tetrahedral Dice Again)

- Here, $\Omega = \{1, 2, 3, 4\}$ and $\mathcal{F}$ is the set of all subsets of Ω.
- The atoms in this case are $\{1\}, \{2\}, \{3\}$ and $\{4\}$.
- Physical symmetry suggests that:

$$\mathbb{P}(\{1\}) = \mathbb{P}(\{2\}) = \mathbb{P}(\{3\}) = \mathbb{P}(\{4\})$$

- Axiomatically, $1 = \mathbb{P}(\{1, 2, 3, 4\}) = \sum_{i=1}^{4} \mathbb{P}(\{i\}) = 4\mathbb{P}(\{1\})$.
- And we again end up with the expected result $\mathbb{P}(\{i\}) = 1/4$ for all $i \in \Omega$.

Example (Lotteries Again)

- $\Omega = \{\text{All unordered sets of 6 numbers from}\{1, \ldots, 49\}\}$
- $\mathcal{F} =$ All subsets of Ω
- Atoms are once again the sets containing a single element of Ω.

This is usual when $|\Omega| < \infty \ldots$

- As $|\Omega| = 13983816$, we have that many atoms.
- Each atom corresponds to drawing one unique subset of 6 balls.
- We might assume that each subset has equal probability... in which case:

$$\mathbb{P}(\{< i_1, i_2, i_3, i_4, i_5, i_6 >\}) = 1/13983816$$

for any valid set of numbers $< i_1, \ldots, i_6 >$.

Complete Spatial Randomness and π

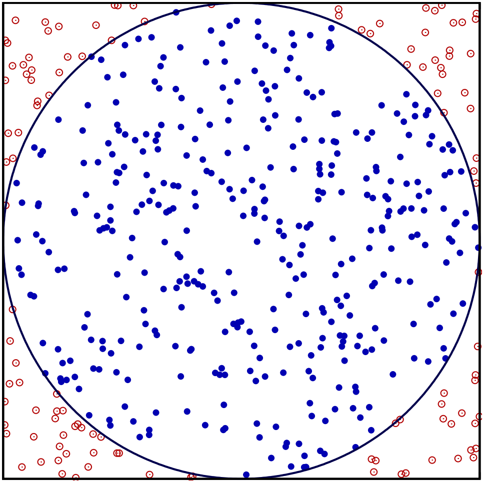

- Let (X, Y) be uniform over the centred unit square.
- Define

$$E = \left\{(x, y) : x^2 + y^2 \leq \frac{1}{4}\right\}$$

- Now

$$\begin{aligned}\mathbb{P}((X, Y) \in E) =& A_{\text{circle}}/A_{\text{square}} \\ =& \pi \times (1/2)^2/1^2 \\ =& \pi/4\end{aligned}$$

Balls in Urns

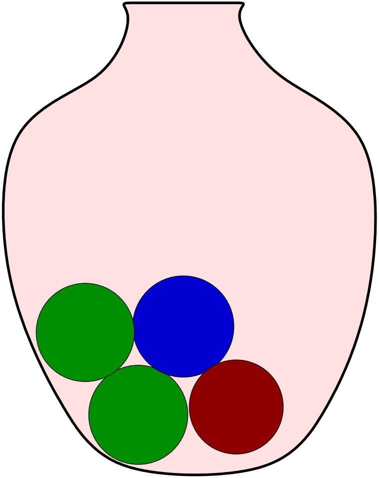

- Let $\mathcal{I}$ be (discrete) a set of colours.
- An urn contains n_i balls of colour i.
- The probability that a drawn ball has colour i is:

$$\frac{n_i}{\sum_{j \in \mathcal{I}} n_j}$$

We assume that the colour of the ball does not influence its probability of selection.

Spinners

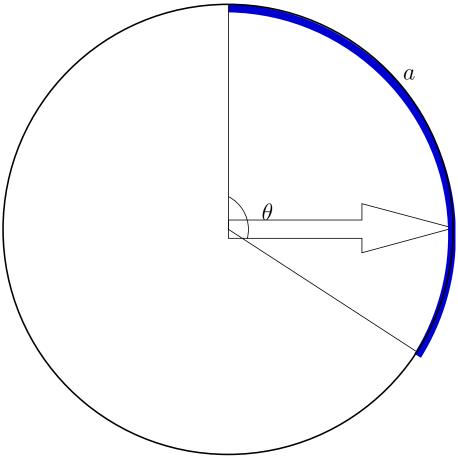

- $\mathbb{P}$[Stops in purple] $= a$
- Really a statement about physics.
- What do we mean by probability?

A Frequency Interpretation

A classical *objective* interpretation of probabilities.

Consider repeating an experiment, with possible outcomes Ω, n times.

- Let $X_1, \ldots, X_n$ denote the results of each experiment.
- Let $A \subset \Omega$ denote an event of interest ($A \in \mathcal{F}$).
- If we say $\mathbb{P}(A) = p_A$ we mean:

$$\lim_{n \to \infty} \frac{\sum_{i=1}^{n} \mathbb{I}_A(X_i)}{n} = p_A$$

where

$$\mathbb{I}_A(X_i) = \begin{cases} 1 & \text{if } X_i \in A \\ 0 & \text{otherwise} \end{cases}$$

Probabilities are relative frequencies of occurrence.

Subjective Probability

What is the probability of a nuclear war occurring next year?

- First, we must be precise about the question.
- We can't appeal to symmetry of geometry.
- We can't appeal meaningful to an infinite ensemble of experiments.
- We *can* form an individual, *subjective* opinion.

If we adopt this subjective view, difficulties emerge:

- How can we quantify degree of belief?
- Will the resulting system be internally consistent?
- What does our calculations actually tell us?

Bayesian/Behavioural/Subjective Probability

- All uncertainty can be represented via probabilities.
- Inference can be conducted using Bayes rule:

$$\mathbb{P}(\theta|y) = \frac{\mathbb{P}(y|\theta)\mathbb{P}(\theta)}{\mathbb{P}(y)}$$

- Later [Bruno de Finetti et al.]: Probability is *personalistic* and *subjective*.

Rev. Thomas Bayes, "An Essay towards solving a Problem in the Doctrine of Chances", Philosophical Transactions of the Royal Society of London (1763). Reprinted as Biometrika 45:293–315 (1958).

http://www.stat.ucla.edu/history/essay.pdf

A Behavioural Definition of Probability

- Consider a *bet*, $b(M, A)$, which pays a reward M if A happens and nothing if A does not happen.
- Let $m(M, A)$ denote the maximum that *You* would be prepared to pay for that bet.
- Two events A_1 and A_2 are equally probable if $m(M, A_1) = m(M, A_2)$.
- Equivalently $m(M, A)$ is the minimum that *You* would accept to offer the bet.
- A value for $m(M, \Omega \setminus A)$ is implied for a rational being...

Personal probability must be a matter of action!

A Bayesian View of Symmetry

- If $A_1, \ldots, A_k$ are *disjoint/mutually exclusive*, *equally likely* and *exhaustive*

$$\Omega = A_1 \cup \cdots \cup A_k,$$

- then, for any i,

$$\mathbb{P}(A_i) = \frac{1}{k}.$$

- Think of the examples we saw before...

Discretised Spinners

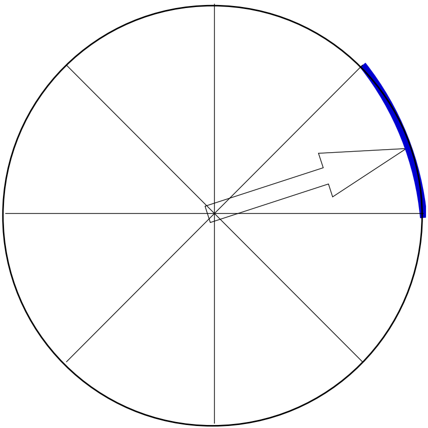

- Each of k segments is equally likely:

$$\mathbb{P}[\text{Stops in purple}] = 1/k$$

- k may be very large.
- Combinations of arcs give rational lengths.
- Limiting approximations give real lengths.
- We can describe *most* subsets this way [ST213].

Example (House selling again)

- The three atoms in this case were:

$$\{(i,j) : i > j\} \qquad \{(i,j) : i = j\} \qquad \{(i,j) : i < j\}$$

- No reason to suppose all three are equally likely.
- If our bidders are believed to be *exchangeable*

$$\mathbb{P}(\{(i,j) : i > j\}) = \mathbb{P}(\{(i,j) : i < j\})$$

- So we arrive at the conclusion that:

$$\mathbb{P}(\{(i,j) : i > j\}) = \mathbb{P}(\{(i,j) : i < j\}) \leq \frac{1}{2}$$
$$\mathbb{P}(\{(i,j) : i = j\}) \geq 0$$

- One strategy would be to accept the first offer if $i > k$...

Elicitation

What probabilities does someone assign to a complex event?

- We can use our behavioural definition of probability.
- The *urn* and *spinner* we introduced before have probabilities which we all agree on.
- We can use these to *calibrate* our personal probabilities.
- When does an *urn* or *spinner* bet have the same value as one of interest.
- There are some difficulties with this approach, but it's a starting point.

A First Look At Coherence

- Consider a collection of events $A_1, \ldots, A_n$.
- If
  - the elements of this collection are disjoint: $\forall i \neq j : A_i \cap A_j = \emptyset$
  - the collection is exhaustive: $\cup_{i=1}^{n} A_i = \Omega$

  then a collection of probabilities $p_1, \ldots, p_n$ for these events is *coherent* if:
  - $\forall i \in \{1, \ldots, n\} : p_i \in [0, 1]$
  - $\sum_{i=1}^{n} p_i = 1$

Assertion: A *rational being* will adjust their personal probabilities until they are coherent.

Dutch Books

- A collection of bets which:
  - definitely won't lead to a loss, and
  - might make a profit

  is known as a Dutch book.

  A rational being would not accept such a collection of bets.
- If a collection of probabilities is incoherent, then a Dutch book can be constructed.

  A rational being must have coherent personal probabilities.

Example (Trivial Dutch Books)

- Consider two cases of incoherent beliefs in the coin-tossing experiment:

  Case 1 $P(\{H\}) = 0.4, P(\{T\}) = 0.4.$
  Case 2 $P(\{H\}) = 0.6, P(\{T\}) = 0.6.$

- To exploit our good fortune, in case 1:
  - *Place* a bet of £X on *both* possible outcomes.
  - Stake is £$2X$; we win £$X/\frac{2}{5} =$ £$5X/2$.
  - Profit is £$(5/2 - 2)X = X/2$.
- In case 2:
  - *Accept* a bet of £X on *both* possible outcomes.
  - Stake is £$2X$; we lose £$X/\frac{3}{5} =$ £$5X/3$.
  - Profit is £$(2 - 5/3)X = X/3$.

Example (A Gambling Example)

Consider a horse race with the following odds:

Horse	**Odds**
Padwaa	7-1
Nutsy May Morris	5-1
Fudge Nibbles	11-1
Go Lightning	10-1
The Coaster	11-1
G-Nut	5-1
My Bell	10-1
Fluffy Hickey	15-1

If you had £100 available, how would you bet?

Example

My own collection of bets looked like this:

Horse	Odds	Stake
Padwaa	7-1	£14.38
Nutsy May Morris	5-1	£19.17
Fudge Nibbles	11-1	£9.58
Go Lightning	10-1	£10.46
The Coaster	11-1	£9.58
G-Nut	5-1	£19.17
My Bell	10-1	£10.45
Fluffy Hickey	15-1	£7.19

Outcome: profit of

$$16 \times \pounds 7.19 - \pounds 99.99 = \pounds(115.04 - 99.99) = \pounds(15.05)$$

Example

My own collection of bets looked like this:

Horse	**Odds**	**Implicit P.**	**Stake**
Padwaa	7-1	0.125	£14.38
Nutsy May Morris	5-1	0.167	£19.17
Fudge Nibbles	11-1	0.083	£9.58
Go Lightning	10-1	0.091	£10.46
The Coaster	11-1	0.083	£9.58
G-Nut	5-1	0.167	£19.17
My Bell	10-1	0.091	£10.45
Fluffy Hickey	15-1	0.063	£7.19

Outcome: profit of

$$16 \times \pounds 7.19 - \pounds 99.99 = \pounds(115.04 - 99.99) = \pounds(15.05)$$

Example

My own collection of bets looked like this:

Horse	**Odds**	**Implicit P.**	**Stake**	S/P
Padwaa	7-1	0.125	£14.38	£115.04
Nutsy May Morris	5-1	0.167	£19.17	£115.02
Fudge Nibbles	11-1	0.083	£9.58	£114.96
Go Lightning	10-1	0.091	£10.46	£115.06
The Coaster	11-1	0.083	£9.58	£114.96
G-Nut	5-1	0.167	£19.17	£115.02
My Bell	10-1	0.091	£10.45	£115.06
Fluffy Hickey	15-1	0.063	£7.19	£115.04

Outcome: profit of

$$16 \times £7.19 - £99.99 = £(115.04 - 99.99) = £(15.05)$$

Efficient Markets and Arbitrage

- The *efficient market hypothesis* states that the prices at which instruments are traded reflects all available information.
- In the world of economics a Dutch book would be referred to as an arbitrage opportunity: a risk-free collection of transactions which guarantee a profit.
- The *no arbitrage principle* states that there are no arbitrage opportunities in an efficient market at equilibrium.
- The collective probabilities implied by instrument prices are coherent.

Elicitation

What does she believe?

We need to obtain and quantify our clients beliefs.
Asking for a direct statement about personal probabilities doesn't usual work:

- $\mathbb{P}(A) + \mathbb{P}(A^c) \neq 1$
- Recall the British economy: people confuse belief with desire.

A better approach uses *calibration*: comparison with a standard.

Example (General Election Results)

Which party you think will win most seats in the next general election?

- Conservative
- Labour
- Liberal Democrat
- Green
- Monster-Raving Loony

Consider the bet $b(£1, \text{Conservative Victory})$:

- You win £1 if the Conservative party wins.
- You win nothing otherwise.

Behavioural Approach to Elicitation

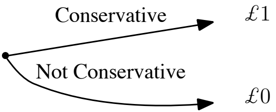

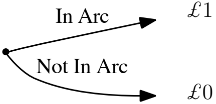

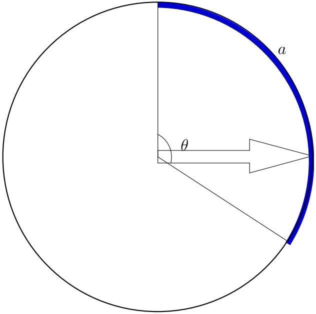

- We said that A_1 and A_2 are *equally probable* if $m(M, A_1) = m(M, A_2)$.
- The probability of a Conservative victory is the same as the probability of a spinner bet of the same value.
- What must a be for us to prefer the spinner bet to the political one?

Eliciting With Urns Full of Balls

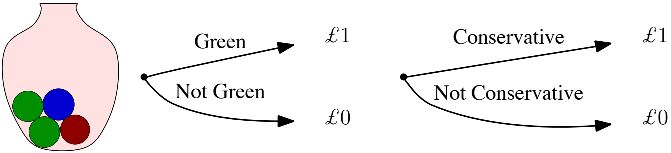

- If the urn contains:
  - n balls
  - g of which are green
- Increase g from 0 to n...
- Let $g^\star$ be such that
  - The real bet is preferred when $g = g^\star$.
  - The urn bet is preferred when $g = g^\star + 1$.

- This tells us that:
  - $\mathbb{P}(\text{C.}) \geq g^\star/n$
  - $\mathbb{P}(\text{C.}) \leq (g^\star + 1)/n$
- Nominal accuracy of $1/n$.

Why should subjective probabilities behave in the same way as our axiomatic system requires?

- We began with axiomatic probability.
- We introduce a subjective interpretation of probability.
- We wish to combine both aspects…

- We briefly looked at "coherence" previously.
- Now, we will formalise this notion.

Coherence Revisited

Definition

Coherence An individual, $\mathcal{I}$, may be termed *coherent* if her probability assignments to an algebra of events obey the probability axioms.

Assertion

A rational individual must be coherent.

A Dutch book argument in support of this assertion follows.

Theorem

Any rational individual, $\mathcal{I}$, must have $\mathbb{P}(A) + \mathbb{P}(A^c) = 1$.

Proof: Case 1: $\mathbb{P}(A) + \mathbb{P}(A^c) < 1$
Consider an urn bet with n balls.

- Let $g^\star(A)$ and $g^\star(A^c)$ be preferred to bets on A and A^c.
- As $\mathbb{P}(A) + \mathbb{P}(A^c)$, for large enough n and $k > 0$:

$$g^\star(A) + g^\star(A^c) = n - k.$$

- (Think of an urn with *three* types of ball).
- Let $b^u(n, k)$ pay £1 if a "k from n" urn-draw wins.
- Bet $b(A)$ pay £1 if event A happens.
- Consider two systems of bets…

- System 1: $S_1^u = [b^u(n, g^\star(A)), b^u(n, g^\star(A^c) + k)]$

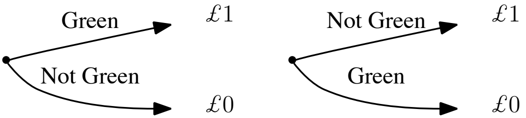

- System 2: $S_1^e = [b(A), b(A^c)]$

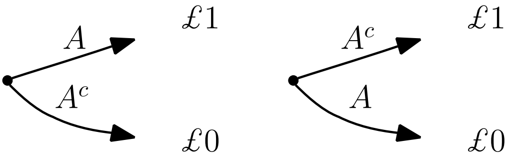

- $\mathcal{I}$ prefers S_1^u to S_1^e and so should pay to win on S_1^u and lose of S_1^e... but everything cancels!

Case2: $\mathbb{P}(A) + \mathbb{P}(A^c) > 1$

- Now, our elicited urn-bets must have

$$g^\star(A) + g^\star(A^c) = n + k$$

- Consider an urn with $g^\star(A)$ green balls and $g^\star(A^c) - k$ blue.
- This time, consider two other systems of bets:

$$S_2^u = [b^u(n, g^\star(A)), b^u(n, g^\star(A^c) - k)]$$

$$S_2^e = [A, A^c]$$

- The stated probabilities mean, $\mathcal{I}$ will pay £c to win on S_2^e and lose on S_2^u.
- Again, everything cancels.

A rational individual won't pay for a bet which certainly returns £0. So $\mathbb{P}(A) + \mathbb{P}(A^c) = 1$.

Theorem

A rational individual, $\mathcal{I}$, must set

$$\mathbb{P}(A) + \mathbb{P}(B) = \mathbb{P}(A \cup B)$$

for any $A, B \in \mathcal{F}$ with $A \cap B = \emptyset$.

Proof: Case 1 $\mathbb{P}(A) + \mathbb{P}(B) < \mathbb{P}(A \cup B)$

- Urn probabilities must be such that:

$$g^\star(A) + g^\star(B) = g^\star(A \cup B) - k$$

- Let

$$s_3^e = [b(A), b(B)]$$

and

$$S_3^u = [b^u(n, g^\star(A)), b^u(n, g^\star(B) + k)]$$

- $\mathcal{I}$ will pay £c to win with S_u^3 which they consider equivalent to $b(\{A \cup B\}$ and lose with S_e^3 ...
- Hence they will pay to win and lose on equivalent events!

Example (Football betting)

- Football team C is to play AV.
- A friend says:

$$\mathbb{P}(C) = \qquad \mathbb{P}(C \text{ wins}) = \frac{7}{8}$$
$$\mathbb{P}(A) = \qquad \mathbb{P}(AV \text{ wins}) = \frac{1}{3}$$

- This is vexatious. Your revenge is as follows:
- Consider an urn containing 7 balls; 6 are green...
- and the "sure-thing" system of bets:

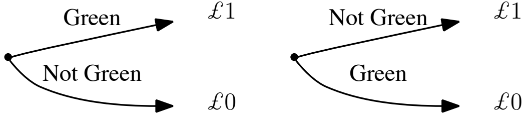

Example (continued)

- The two urn bets are inferior to $b(C)$ and $b(A)$, respectively.
- Your friend should pay £c to win on $[b(A), b(C)]$ but lose on the urn system.
- But logically, $b(C)$ and $b(A)$ are not exhaustive (there may be a draw).
- So your friend should pay a little to switch back.
- Iterate until your point has been made.
- If your friend refuses argue that their "probabilities" are meaningless.

The Cox-Jaynes Axioms

Another view: if we want the following to hold

- Degrees of plausibility can be represented by real numbers, B.
- Mathematical reasoning should show a qualitative correspondence with common sense.
- If a conclusion can be reasoned out in more than one way, then every possible way must lead to the same result.

Then, up to an arbitrary rescaling, B, must satisfy our probability axioms.
See "Probability Theory: The Logic of Science" by E. T. Jaynes for a recent summary of these results.

Caveat Mathematicus

There are several points to remember:

- Subjective probabilities are subjective.

  People need not agree.

- Elicited probabilities should be coherent.

  The decision analyst must ensure this.

- Temporal coherence is not assumed or assured.

  You are permitted to change your mind.

The latter is re-assuring, but how *should* we update our beliefs?

Conditions

Conditional Probabilities

- The probability of one event occurring *given* that another has occurred is critical to Bayesian inference and decision theory.
- If A and B are events and $\mathbb{P}(B) > 0$, then the *conditional probability* of A given B (i.e. conditional upon the fact that B is known to occur) is:

$$\mathbb{P}(A|B) = \mathbb{P}(A \cap B)/\mathbb{P}(B)$$

- This amounts to taking the restriction of $\mathbb{P}$ to B and renormalizing.

Example (Cards)

- Consider a standard deck of 52 cards which is well shuffled.
- Let A be the event "drawing an ace".
- Let B be the event "drawing a spade".
- If we believe that each card is equally probable:

$$\begin{aligned}
\mathbb{P}(A) =& 4/52 = 1/13 \\
\mathbb{P}(B) =& 13/52 = 1/4 \\
\mathbb{P}(A|B) =& \mathbb{P}(A \cap B)/\mathbb{P}(B) \\
=& 1/52/13/52 = 1/13
\end{aligned}$$

- Knowing that a card is a spade doesn't influence the probability that it is an ace.

Example (Cards Again)

- Consider a standard deck of 52 cards which is well shuffled.
- Let A' be the event "drawing the ace of spades".
- Let B be the event "drawing a spade".
- If we believe that each card is equally probable:

$$
\begin{aligned}
\mathbb{P}(A') &= 1/52 \\
\mathbb{P}(B) &= 13/52 = 1/4 \\
\mathbb{P}(A'|B) &= \mathbb{P}(A' \cap B)/\mathbb{P}(B) \\
&= 1/52/13/52 = 1/13
\end{aligned}
$$

- Knowing that a card is a spade does influence the probability that it is the ace of spades.

Called-off Bets

- We must justify the interpretation of conditional probabilities within a subjective framework.
- Consider a *called-off bet* $b(A|B)$ which pays
  - £1 if A happens and B happens,
  - nothing if B happens but A does not
  - nothing and is called off (stake is returned) if B does not happen.

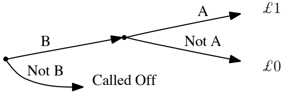

- How would a rational being value such a bet?

Theorem (Conditional Probability and Called-Off Bets)

A rational individual, $\mathcal{I}$, with subjective probability measure $\mathbb{P}$ must assess the called-off bet $b(A|B)$ as having the same value as a simple bet on an event with probability $\mathbb{P}(A|B)$.

Outline of proof:

- Consider a simple bet with 4 possible outcomes ($A \cap B, A \cap B^c, A^c \cap B$ and $A^c \cap B^c$).
- Given an urn containing n balls, let n_{AB} be red, n_{AB^c} be blue, n_{A^cB} be green and $n_{A^cB^c}$ be yellow.
- Choose that $\mathcal{I}$ is indifferent to bets on the four outcomes and the four colours of ball.

- Logically, a bet on B or B^c is of the same value as one on (red or blue) or on (green or yellow)
- Consider a second bet: B occurs. What are the probabilities $\mathcal{I}$ attaches to A and A^c conditional upon this?
- Given an urn with m balls, let m_A and m_{A^c} be the number of red and blue balls.
- Let m_A and m_{A^c} be chosen such that $\mathcal{I}$ is indifferent to the two bets.
- By equivalence/symmetry arguments, we may deduce that:

$$\frac{n_{AB} + n_{A^cB}}{n} \times \frac{m_A}{m} = \frac{n_{AB}}{n}$$

- Hence

$$\frac{m_A}{m} = \frac{n_{AB}}{n_{AB} + n_{A^cB}} = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(A \cap B) + \mathbb{P}(A \cap B^c)}$$

Independence

Some events are unrelated to one another. That is, sometimes knowing that an event B occurs tells us nothing about how probable it is that a second event, A, also occurs.

Definition (Independence)

Events A and B are *independent* if:

$$\mathbb{P}(A \cap B) = \mathbb{P}(A) \times \mathbb{P}(B)$$

and this can be written as $A \perp\!\!\!\perp B$.

If A and B are independent and of positive probability, then:

$$\mathbb{P}(A|B) = \mathbb{P}(A)$$
$$\mathbb{P}(B|A) = \mathbb{P}(B)$$

Learning about one doesn't influence our beliefs about the other.

The Law of Total Probability

- Let $B_1, \ldots, B_n$ *partition* the space:

$$\bigcup_{i=1}^{n} B_i = \Omega$$

$$B_i \cap B_j = \emptyset \qquad \forall i \neq j$$

- Let A be another event.
- It is simple to verify that:

$$A = \bigcup_{i=1}^{n} (B_i \cap A)$$

- And hence that:

$$\mathbb{P}(A) = \sum_{i=1}^{n} \mathbb{P}(A \cap B_i)$$

- This is sometimes termed the law of total probability.

The Partition Formula

Theorem (The Partition Formula)

If $B_1, \ldots, B_n$ partition Ω, then:

$$\mathbb{P}(A) = \sum_{i=1}^{n} \mathbb{P}(A|B_i)\mathbb{P}(B_i)$$

Proof:
By the law of total probability:

$$\mathbb{P}(A) = \sum_{i=1}^{n} \mathbb{P}(A \cap B_i)$$

and $\mathbb{P}(A \cap B_i) = \mathbb{P}(A|B_i)\mathbb{P}(B_i)$ by definition of $\mathbb{P}(A|B_i)$.

Example (Buying a house)

- Your client wishes to decide whether to buy a house.
- If $A =$ [Making a loss when buying the house.]
- It might be easier to elicit probabilities for component events:

$$\mathbb{P}(A) = \sum_i \mathbb{P}(A|B_i)\mathbb{P}(B_i)$$

where

$$\begin{aligned} E_1 =& \text{[Inflation is low.]} \\ E_2 =& \text{[Inflation is high; salary rises]} \\ E_1 =& \text{[Inflation is high; salary doesn't rise]} \end{aligned}$$

Bayes' Rule

The core of Bayesian analysis is the following elementary result:

Theorem

If A and B are events of positive probability, then:

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(A)\mathbb{P}(B|A)}{\mathbb{P}(B)} = \frac{\mathbb{P}(A)\mathbb{P}(B|A)}{\mathbb{P}(A)\mathbb{P}(B|A) + \mathbb{P}(A^c)\mathbb{P}(B|A^c)}$$

Proof: This follows directly from the definition of conditional probability:

$$\mathbb{P}(A|B)\mathbb{P}(B) = \mathbb{P}(A \cap B) = \mathbb{P}(B|A)\mathbb{P}(A)$$

This allows us to update our beliefs.

Example (Disease Screening)

Consider screening a rare disease.

$$
\begin{aligned}
A =& [\text{Subject has disease.}] \\
B =& [\text{Screening indicates disease.}]
\end{aligned}
$$

If $\mathbb{P}(A) = 0.001$, $\mathbb{P}(B|A) = 0.9$ and $\mathbb{P}(B|A^c) = 0.1$ then:

$$
\begin{aligned}
\mathbb{P}(A|B) =& \frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(B|A)\mathbb{P}(A) + \mathbb{P}(B|A^c)\mathbb{P}(A^c)} \\
=& \frac{0.9 \times 0.001}{0.9 \times 0.001 + 0.1 \times 0.999} \\
=& 0.0089
\end{aligned}
$$

Think about what this *means*. . .

screening requires small $\mathbb{P}(B|A^c)$

Some Bayesian Terminology

- In the previous example $\mathbb{P}(A)$ is the *prior* probability of the subject carrying the disease.

  That is, the probability assigned to the event before the observation of data.

- Given that event B is observed, $\mathbb{P}(A|B)$ is termed the *posterior* probability of A.

  That is, the probability assigned to the event after the observation of data.

- Note that these aren't absolute terms: in a sequence of experiments the posterior distribution from one stage may serve as the prior distribution for the next.

Random Variables

- So far we have talked only about events.
- It is useful to think of *random variables* in the same language.
- Let X be a "measurement" which can take values $x_1, \ldots, x_n$.
- let $\mathcal{F}$ be the algebra generated by $\mathcal{X}$.
- If we have a probability measure, $\mathbb{P}$, over $\mathcal{F}$ then X is a random variable with law $\mathbb{P}$.
- A probability mass function is sufficient to specify $\mathbb{P}$.

Example (Roulette)

- Consider spinning a roulette wheel with $n(r) = n(b) = 18$ red/black spots and $n(g) = 1$ green one.
- Set X to 1 if the ball stops in a red region, 2 for a black one and 20 for a green.
- Under a suitable assumption of symmetry, the probability mass function is:

$$\begin{aligned} \mathbb{P}[X = 1] =& n(r)/n \\ \mathbb{P}[X = 2] =& n(g)/n \\ \mathbb{P}[X = 20] =& n(b)/n \end{aligned}$$

where $n = n(r) + n(g) + n(b) = 37$ normalises the distribution.

Independence of Random Variables

As you might expect, the concept of independence can also be applied to random variables.

Definition

Random variables, X and Y, are independent if for all possible x_i, y_j:

$$\mathbb{P}[X = x_i, Y = y_j] = \mathbb{P}[X = x_i]\mathbb{P}[Y = y_j]$$

[Mathematical] Expectation

It is useful to have a mathematical idea of the *expected value* of a random variable: a weighted average of its possible values that behaves as a “centre of probability mass”.

Definition

The expectation of a random variable, X, is:

$$\mathbb{E}\left[X\right] = \sum_i x_i \times \mathbb{P}[X = x_i]$$

where the sum is taken over all possible values.

Useful Properties of Expectations

- Expectation is linear:

$$\mathbb{E}\left[aX + bY + c\right] = a\mathbb{E}\left[X\right] + b\mathbb{E}\left[Y\right] + c$$

- The expectation of a *function* of a random variable is:

$$\mathbb{E}\left[f(X)\right] = \sum_i f(x_i) \times \mathbb{P}[X = x_i]$$

  where the sum is over all possible values.
- One interpretation: a function of a random variable is itself a random variable.
- If X takes values in $x_i \in \Omega$ with probabilities $\mathbb{P}[X = x_i]$ then $f(X)$ takes values $f(x_i)$ in $f(\Omega)$:

$$\mathbb{P}[f(X) = f(x_i)] = \mathbb{P}[X = x_i].$$

Example (Die Rolling)

Consider rolling a six-sided die:

- $\Omega = \{1, 2, 3, 4, 5, 6\}$
- Let X be the number rolled.
- Under a symmetry assumption:

$$\forall x \in \Omega : \qquad \mathbb{P}[X = x] = 1/6$$

- Hence, the expectation is:

$$\begin{aligned}\mathbb{E}\left[X\right] &= \sum_{x \in \Omega} x \mathbb{P}[X = x] \\ &= \sum_{x=1}^{6} x \mathbb{P}[X = x] \\ &= 21 \times 1/6 = 7/2\end{aligned}$$

Example (A Roulette Wheel Again)

- Recall the roulette random variable introduced earlier.

$$\begin{aligned}\mathbb{E}[X] =& \sum_{x_i} x_i \times \mathbb{P}[X = x_i] \\ =& 1 \times \mathbb{P}[X = 1] + 2 \times \mathbb{P}[X = 2] + 20 \times \mathbb{P}[X = 20] \\ =& 1 \times n(r)/n + 2 \times n(b)/n + 20 \times n(g)/n \\ =& (n(r) + 2 \times n(b) + 20 \times n(g))/n = (18 + 36 + 20)/37 = 2\end{aligned}$$

- Whilst, considering $f(x) = x^2$ we have:

$$\begin{aligned}\mathbb{E}\left[X^2\right] =& \mathbb{E}[f(X)] \\ =& 1^2 \times \mathbb{P}[X = 1] + 2^2 \times \mathbb{P}[X = 2] + 20^2 \times \mathbb{P}[X = 20] \\ =& (n(r) + 4 \times n(b) + 400 \times n(g))/n = 490/37\end{aligned}$$

Decisions

Decision Ingredients

The basic components of a decision analysis are:

- A space of possible decisions, D.
- A set of possible outcomes, $\mathcal{X}$.

By choosing an element of D you exert some influence over which of the outcomes occurs.

Definition (Loss Function)

A loss function, $L : D \times \mathcal{X} \to \mathbb{R}$ relates decisions and outcomes. $L(d, x)$ quantifies the amount of loss incurred if decision d is made and outcome x then occurs.

An algorithm for choosing d is a *decision rule*.

Example (Insurance)

- You must decide whether to pay c to insure your possessions of value v against theft for the next year:

$$d = \{\text{Buy Insurance}, \text{Don't Buy Insurance}\}$$

- Three events are considered possible over that period:

$x_1 = $ {No thefts.} $x_2 = $ {Small theft, loss $0.1v$}
$x_3 = $ {Serious burglary, loss v}

- Our loss function may be tabulated:

$L(d,x)$	x_1	x_2	x_3
Buy	c	c	c
Don't Buy	0	$0.1v$	v

Uncertainty in Simple Decision Problems

- As well as knowing how desirable action/outcome pairs are, we need to know how probable the various possible outcomes are.
- We will assume that the underlying system is independent of our decision.
- Work with a probability space $\Omega = \mathcal{X}$ and the algebra generated by the collection of single elements of $\mathcal{X}$.
- It suffices to specify a probability mass function for the elements of $\mathcal{X}$.
- One way to address uncertainty is to work with expectations.

Example (Insurance Continued)

- There are 25 million occupied homes in the UK (2001 Census).
- Approximately 280,000 domestic burglaries are carried out each year (2007/08 Crime Report)
- Approximately 1.07 million acts of "theft from the house" were carried out.
- We might naïvely assess our pmf as:

$$
\begin{aligned}
p(x_1) =& \frac{25 - 1.07 - 0.28}{25} && = 0.946 \\
p(x_2) =& \frac{1.07}{25} && = 0.043 \\
p(x_3) =& \frac{0.28}{25} && = 0.011
\end{aligned}
$$

The EMV Decision Rule

- If we calculate the expected loss for each decision, we obtain a function of our decision:

$$\bar{L}(d) = \mathbb{E}\left[L(d, X)\right] = \sum_{x \in \mathcal{X}} L(d, x) \times p(x)$$

- The *expected monetary value* strategy is to choose $d^\star$, the decision which minimises this expected loss:

$$d^\star = \arg\min_{d \in D} \bar{L}(d)$$

- This is sometimes known as a *Bayesian decision.*
- A justification: If you make a lot of decisions in this way the you might expect an averaging effect...

Example (Still insurance)

- Here, we had a loss function:

$L(d, x)$	x_1	x_2	x_3
Buy	c	c	c
Don't Buy	0	$0.1v$	v

- And a pmf:

$$p(x_1) = 0.946 \qquad p(x_2) = 0.043 \qquad p(x_3) = 0.011$$

- Which give us an expected loss of:

$$\begin{aligned} \bar{L}(\text{Buy}) &= 0.946c + 0.043c + 0.011c && = c \\ \bar{L}(\text{Don't Buy}) &= 0.946 \times 0 + 0.0043v + 0.011v && = 0.0153v \end{aligned}$$

- Our decision should, of course, depend upon c and v.
- If $c < 0.0153v$ then the EMV decision is to buy insurance:

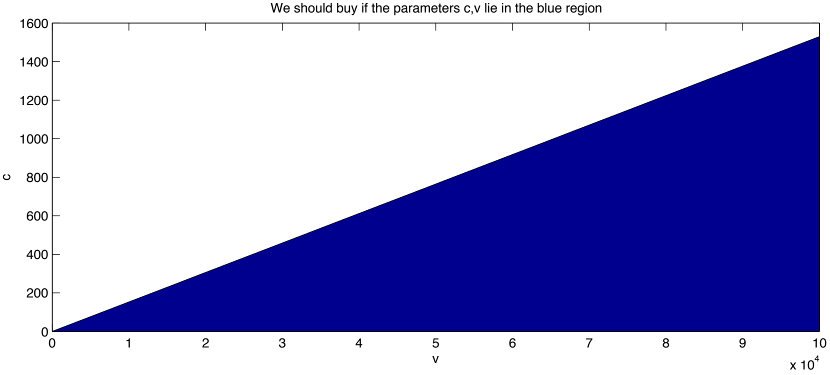

Optimistic EMV

- We can be more optimistic in our approach.
- Rather than defining a *loss function*, we could work with a *reward function*:

$$R(d, x) = -L(d, x)$$

- Leading to an expected reward:

$$\bar{R}(d) = \mathbb{E}\left[R(d, \cdot)\right] = -\mathbb{E}\left[L(d, \cdot)\right] = -\bar{L}(d)$$

- And the EMV rule becomes choose

$$d^{\star} = \arg\max_{d \in D} \bar{R}(d)$$

Desiderata

- We need a convenient notation to encode the entire decision problem.
- It must represent all possible outcomes for all possible decision paths.
- It must encode the possible outcomes and their probabilities given each set of decisions.
- It must allow us to calculate the EMV decision for a problem. . .

and ultimately, other "optimal" decisions.

Ch.2 of Jim Smith's "Decision Analysis" covers this material in detail.

Graphical Representation: Decision Trees

Drawing a decision tree:

1. Find a large piece of paper.
2. Starting at the left side of the page and working chronologically to the right...
   - 2.1 Indicate decisions with a □.
   - 2.2 Draw forks from decision *nodes* labelled with the decisions.
   - 2.3 Indicate sets of random outcomes with a ◯.
   - 2.4 Draw edges from random event *nodes* labelled with their (conditional) probabilities.
   - 2.5 Continue iteratively until all decisions and random variables are shown.
   - 2.6 At the right hand end of each path indicate the *loss/reward*.

In the case of the insurance example, start with the first possible decision and we obtain:

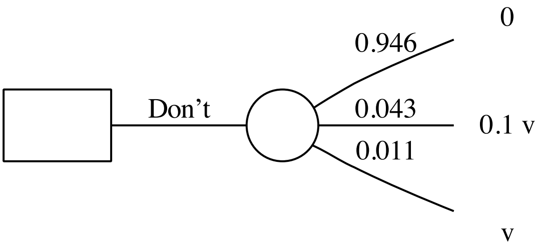

Doing this for all of the decisions and combining them:

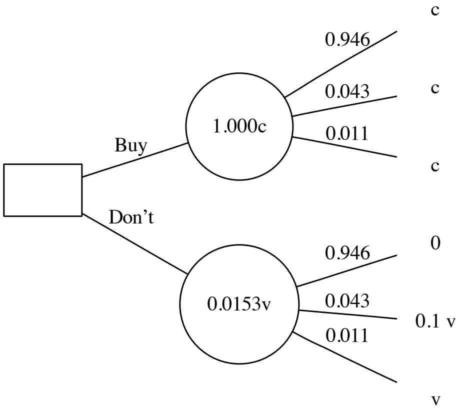

We've worked backwards from the RHS filling in the expected loses associated with each decision.

But we didn't need to make things that complicated... there is only one outcome if we buy insurance:

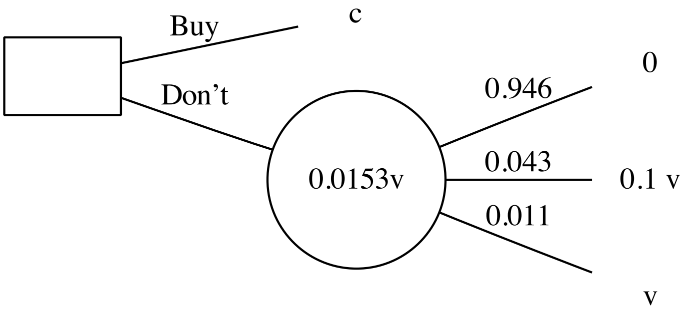

In more complex examples, we should label the random events (say N for no robbery, T for small theft and B for burglary…

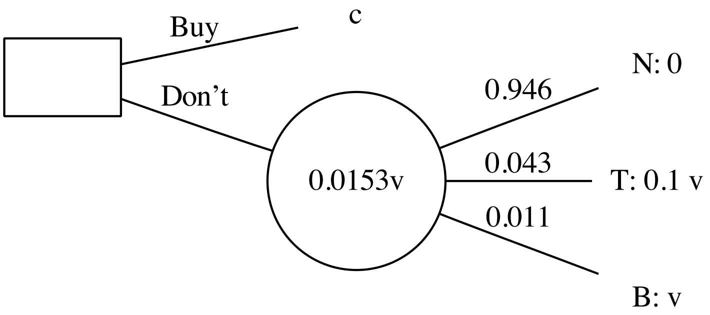

Calculation and Decision Trees

First, we fill in the expected loss associated with decisions:

- starting at the RHS of the graph, trace paths back to ◯ nodes.
- Fill in the rightmost ◯ nodes with the (conditional[4]) expected losses (the probabilities and losses are indicated at the edges and ends of the edges).
- For each decision node which now has values at the end of each branch, find the branch with the largest value.
- Eliminate all of the others.
- This produces a reduced decision tree.
- Iterate.
- When left with one path, this is the EMV decision!

[4] On all earlier events – i.e. ones to the left

Do Not Laugh at Notations[5]

- At this point you may be thinking that this is a silly picture and that you'd rather just calculate things.
- That's all very well…
- but it gets harder and harder as decisions become more complicated.
- This graphical representation provides an easy to implement recursive algorithm and a convenient representation.
- This lends itself to automatic implementation as well as manual calculation.

[5] Invent them, for they are powerful. *RP Feynman.*

More Complicated Cases

Consider this case:

- You may drill (at a cost of £31M) in one of two sites: field A and field B.
  - If there is oil in site A it will be worth £77M.
  - If there is oil in site B it will be worth £195M.
- Or you may conduct preliminary trials in either field at a cost of £6M.
- Or you can do nothing. This is free.

This gives a set of 5 decisions to make immediately. If you investigate site A or B you must then, further, decide whether to drill there, in the other site or not at all (we'll make things simpler by neglecting the possibility of investigating both).

Your Knowledge

- The probability that there is oil in field A is 0.4.
- The probability that there is oil in field B is 0.2.
- If oil is present in a field, investigation will advise drilling with probability 0.8.
- If oil is not present, investigation will advise drilling with probability 0.2.
- The presence of oil and investigation results in one field provides no information about the other field.

What do you know – formally?

Let A be the event that there is oil in site A and let B be the event that there is oil in site B. Let a be the event that investigation suggests there is oil in site a and let b be the event that investigation suggests that there is oil in site b.
The information on the previous page becomes:

- $\mathbb{P}(A) = 0.4$
- $\mathbb{P}(B) = 0.2$
- $\mathbb{P}(a|A) = \mathbb{P}(b|B) = 0.8$
- $\mathbb{P}(a|A^c) = \mathbb{P}(b|B^c) = 0.2$

Bayes Rule is Needed

We really need to know the probability that oil is present in a field given that investigation indicates that there is (we know the converse).

$$\begin{aligned}\mathbb{P}(A|a) =& \frac{\mathbb{P}(a|A)\mathbb{P}(A)}{\mathbb{P}(a|A)\mathbb{P}(A) + \mathbb{P}(a|A^c)\mathbb{P}(A^c)} \\ =& \frac{0.8 \times 0.4}{0.8 \times 0.4 + 0.2 \times 0.6} = 0.727\end{aligned}$$

$$\begin{aligned}\mathbb{P}(B|b) =& \frac{\mathbb{P}(b|B)\mathbb{P}(B)}{\mathbb{P}(b|B)\mathbb{P}(B) + \mathbb{P}(b|B^c)\mathbb{P}(B^c)} \\ =& \frac{0.8 \times 0.2}{0.8 \times 0.2 + 0.2 \times 0.8} = 0.500\end{aligned}$$

Decision Trees — Example

We begin by constructing the tree without probabilities.

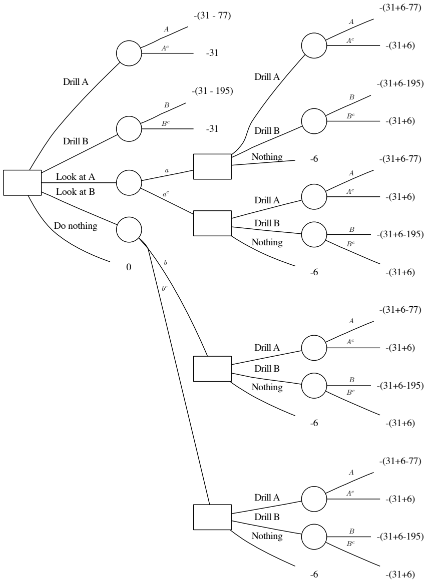

Decision Trees — Example

We begin by constructing the tree without probabilities. Then work out what each probability should be.

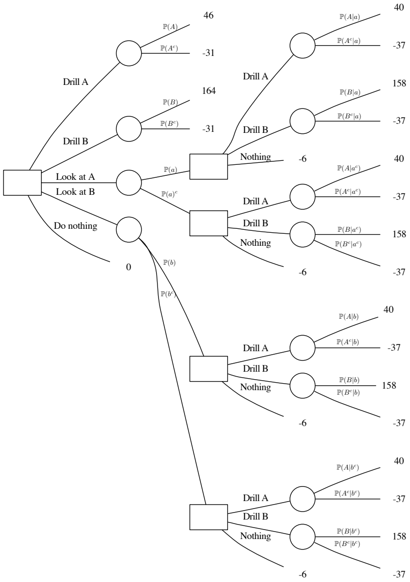

Decision Trees — Example

We begin by constructing the tree without probabilities. Then work out what each probability should be numerically.

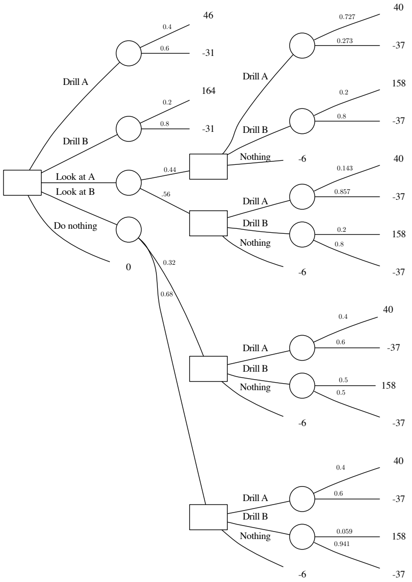

Decision Trees — Example

We begin by constructing the tree without probabilities. Then work out what each probability should be numerically. Then starting at the RHS calculate expectations and make optimal decisions to determine the solution.

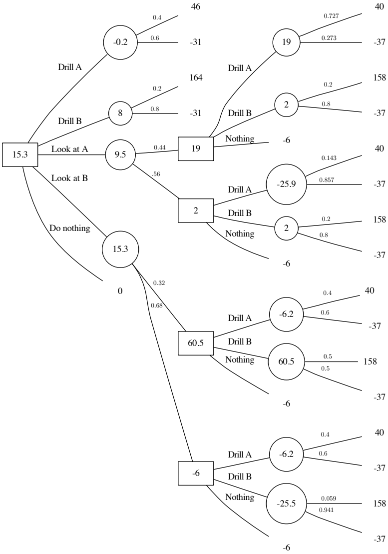

Perfect Information

- How useful would it be to know in advance what value all relevant random variables take?

  If we know everything in advance, how well would we do?

- Expected Value of Perfect Information: the difference in the expected value of a decision problem in which decisions are made with full knowledge of the outcome of chance events and one in which no additional knowledge is available.

Preferences

Example (The Farmer's Trilemma)

A farmer must decide which crop to plant; profit depends upon the weather:

Weather:	Good	Fair	Bad
Crop A	11	1	-3
Crop B	7	5	0
Crop C	2	2	2

- Which crop should he plant?
- Thus far, we've considered EMV decisions.
- What else could we do?

Maximin Decisions

- One farmer believes that the weather will do whatever makes things worst, whatever decision he makes.
- He's either pessimistic or paranoid.
- He maximise his worst case return.
- The worst behaviour of crop A is -3, that of crop B is 0 and that of crop C is 2.
- He consequently sows crop C.
- This is known as a *maximin* decision: it maximises the minimum reward.

Maximax Decisions

- One farmer believes that the weather will do whatever makes things best, whatever decision he makes.
- He's either optimistic or feeling lucky.
- He maximise his best case return.
- The best behaviour of crop A is 11, that of crop B is 7 and that of crop C is 2.
- He consequently sows crop A.
- This is known as a *maximax* decision: it maximises the maximum reward.

The Hazards of Extremism

- Maximin and maximax solutions may sometimes be acceptable.
- But they aren't stable: what if you introduce another possible outcome with probability $\epsilon \ll 1$?
- However small ϵ is, this outcome could be the only one you base you decision upon.
- But, in decision problems, you work with an idealisation in which you haven't really considered *every* possible outcome.
- This seems rather inconsistent.

Paradoxes in St. Petersburg

- How much is the following bet worth?
- The prize is initial £1.
- A fair coin is tossed until a tail is shown.
- The prize is doubled every time a head is shown.
- You win the prize when the first tail arrives.

St Petersburg: Expected Monetary Value

- The expected value of the decision to play this game is:

$$\begin{aligned}\bar{R}(\text{“play”}) &= \sum_{n=1}^{\infty} R(\text{“play”}, n)p(n) \\ &= \sum_{n=1}^{\infty} 2^{n-1}2^{-n} \\ &= \sum_{n=1}^{\infty} \frac{1}{2} = \infty\end{aligned}$$

- So a choice between receiving a reward $\bar{R}(\text{“don't”}) < \infty$ or playing this game should, by EMV, always be resolved by playing.
- Would you rather play this game of have £1,000,000?

Utility of Opportunity / Certain Monetary Equivalence

- If there is a problem with using EMV it is this: it assumes that we value a probability p of receiving some reward r as being of the same value as receiving a reward pr with certainty.
- Would you rather have £10^8 with certainty or a probability of 10^{-9} of having £10^{17}?
- We see that EMV might make sense for moderate probabilities and moderate sums, but it doesn't match our real preferences in general.
- It is useful to think how much a probability p of receiving a reward r is *worth* to us: we call this the *utility* of such a bet.

Some Notation

- Let A, B and C be random outcomes (i.e. particular rewards with some probability or nothing otherwise).
- Write $A \succ B$ if A is preferred to B.
- Write $A \sim B$ if A and B are equally preferable.
- Write $A \succeq B$ if A is at least as good as B.
- For some $t \in (0, 1)$, let $tA + (1-t)B$ denote outcome A occurring with probability t and B with probability $1-t$.

Axioms of Preference

If a collection of preferences obey the following:

1. Completeness: For any A, B one of the following holds:
$$A \succ B \qquad A \sim B \qquad A \prec B$$
2. Transitivity:
$$A \succeq B, B \succeq C \Rightarrow A \succeq C$$
3. Independence: if $A \succ B$ then, for any $t \in [0, 1)$:
$$(1-t)A + tC \succ (1-t)B + tC$$
4. Continuity: If $A \succ B \succ C$, there exists $\rho \in (0, 1)$ such that:
$$\rho A + (1-\rho)C \sim B$$

Then that collection of preferences is considered rational.

Utility Functions

- If the axioms from the previous slide are satisfied…
- The preferences can be encoded in a *utility function*, U.
- This function maps the (monetary) value of each outcome to a real number.
- Maximising the *expectation* of the utility in a decision problem makes decisions compatible with the preferences.

It's outside the scope of this course to prove this…but it will become apparent that it is reasonable from the next few slides.

Eliciting Utilities

If preferences are to be represented by utilities, we must be able to determine utility functions.

- What m would you accept not to benefit from the bet shown?
- This is a function of α.
- The utility of m is $U(m) = f^{-1}(m)$.

This bet:

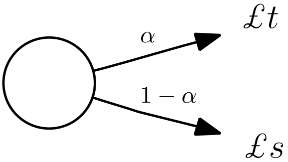

has CME value

$$m = f(\alpha).$$

A Family of Utilities

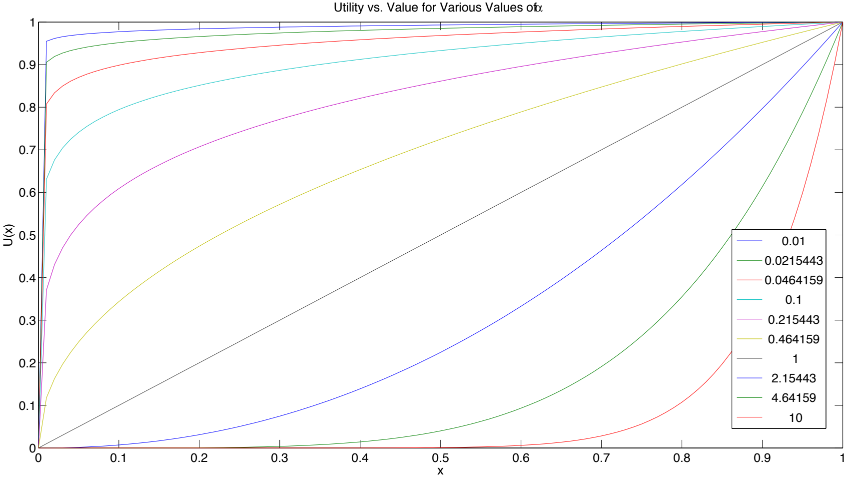

$$U(x) = x^{\alpha} \qquad \alpha > 0$$

Example (The Utility of Insurance)

EMV:

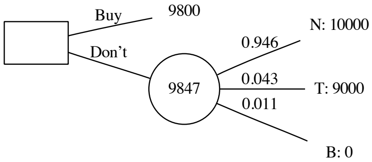

EMU:

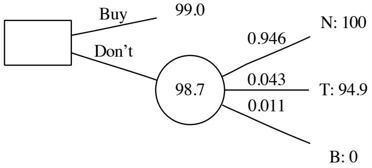

- Consider the insurance example.
- The first figure shows the EMV position: the insurer would prefer you to insure; you'd prefer not to.
- The second shows the EMU position with

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

  You prefer to insure.
- EMV makes sense for the insurer; EMU for you.

Example (The Value of Money)

- Consider a lottery which pays a reward £X where X is a random number distributed uniformly over $[0, 4]$.
- An individual with utility function $U_\alpha(x) = x^\alpha$ considers buying a ticket.
- How much would they be prepared to pay for a ticket?
- The expected utility of the lottery is:

$$\mathbb{E}\left[U_\alpha(X)\right] = \int_0^4 \frac{1}{4} x^\alpha dx = \frac{4^\alpha}{\alpha + 1}$$

- The fair price, x_f is such that

$$U_\alpha(x) = \mathbb{E}\left[U_\alpha(X)\right]$$

Example

- The fair price is the solution of the equation:

$$U(x_f) = x_f^\alpha = \frac{4^\alpha}{x+1}$$
$$x_f = \frac{4}{(x+1)^{1/\alpha}}$$

- For various values of α:

α	0.5	1.0	1.5	2.0
x_f	1.78	2	2.17	2.31

- Notice that for $\alpha < 1$ the "fair price" of the game is less than its expected value; for $\alpha = 1$ the price and expected value coincide and for $\alpha > 1$ a price above the expected value is considered fair.

Making Decisions

We've covered the making of decisions:

1. Determine possible chance events and elicit probabilities.
2. Enumerate the possible actions.
3. Determine preferences via utility.
4. Choose actions to maximise expected utility.
5. Return to elicitation if necessary.

Now, we move on to games…

Games

What is a Game

A *game* in mathematics is, roughly speaking, a problem in which:

- Several *agents* or *players* make 1 or more decisions.
- Each player has an objective / set of preferences.
- The outcome is influenced by the set of decisions.
- There may be additional non-deterministic uncertainty.
- The players may be in competition or they may be cooperating.
- Examples include: chess, poker, bridge, rock-paper-scissors and many others.

However, we will stick to simple two player games with each player simultaneously making a single decision.

Simple Two Player Games

- Player 1 chooses a move for a set $D = \{d_1, \ldots, d_n\}$.
- Plater 2 chooses a move from a set $\Delta = \{\delta_1, \ldots, \delta_m\}$.
- Each player has a *payoff function*.
- If the players choose moves d_i and δ_j, then:
  - Player 1 receives reward $R(d_i, \delta_j)$.
  - Player 2 receives reward $S(d_i, \delta_j)$.
- The relationship between decisions and rewards is often shown in a payoff matrix:

	δ_1	$\ldots$	δ_m
d_1	$(R(d_1, \delta_1), S(d_1, \delta_1))$	$\ldots$	$(R(d_1, \delta_m), S(d_1, \delta_m))$
$\vdots$			$\vdots$
d_n	$(R(d_n, \delta_1), S(d_n, \delta_1))$	$\ldots$	$(R(d_n, \delta_m), S(d_n, \delta_m))$

Payoff Matrices Again

It's sometimes useful to consider a single player's payoff as a function of the possible decisions.

Player 1 and player 2 have these payoff matrices:

	δ_1	$\dots$	δ_m
d_1	$R(d_1, \delta_1)$	$\dots$	$R(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$R(d_n, \delta_1)$	$\dots$	$R(d_n, \delta_m)$

	δ_1	$\dots$	δ_m
d_1	$S(d_1, \delta_1)$	$\dots$	$S(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$S(d_n, \delta_1)$	$\dots$	$S(d_n, \delta_m)$

Example (Rock-Paper-Scissors)

- Each player picks from the same set of decisions:

$$D = \Delta = \{R, P, S\}$$

- R beats S; S beats P and P beats R
- One possible payoff matrix is:

	R	P	S
R	(0,0)	(-1,1)	(1,-1)
P	(1,-1)	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Example (The Prisoner's Dilemma)

- Again, each player picks from the same set of decisions:

$$D = \Delta = \{\text{Stay Silent}, \text{Betray Partner}\}$$

- If they both stay silent they will receive a short sentence; if they both betray one another they will get a long sentence; if only one betrays the other the traitor will be released and the other will get a long sentence.
- One possible payoff matrix is:

	S	B
S	(1,1)	(5,0)
B	(0,5)	(4,4)

- Notice that each player wishes to minimise this payoff!

Example (Love Story)

- A boy and a girl must go to either of:

$$D = \Delta = \{\text{Football}, \text{Opera}\}$$

- They both wish to meet one another most of all.
- If they don't meet, the boy would rather see the football; the girl, the opera.
- A possible payoff matrix might be:

	F	O
F	(100,100)	(50,50)
O	(0,0)	(100,100)

Some Features of these Examples

- The rock-paper-scissors game is *purely competitive*: any gain by one player is matched by a loss by the other player.
- The RPS and PD problems are symmetric:

$$R(d, \delta) = S(\delta, d)$$

  [Note that this makes sense as $D = \Delta$]
- $D = \Delta$ in all three of these examples, but it isn't always the case.

Uncertainty in Games

As the players don't know what action the other will take, there is uncertainty.

- Thankfully, the Bayesian interpretation of probability allows them to encode their uncertainty in a probability distribution.
- Player 1 has a probability mass function p over the actions that player 2 can take, Δ.
- Player 2 has a probability mass function q over the actions that player 1 can take, denoted D.

Expected Rewards

Just as in a decision problem, we can think about expected rewards:

- For player 1, the expected reward of move d_i is:

$$\begin{aligned}\bar{R}(d_i) =& \mathbb{E}\left[R(d_i, \delta_j)\right] \\ =& \sum_{j=1}^{m} q(\delta_j) R(d_i, \delta_j)\end{aligned}$$

- Whilst, for player 2, we have

$$\begin{aligned}\bar{S}(\delta_j) =& \mathbb{E}\left[S(d_i, \delta_j)\right] \\ =& \sum_{i=1}^{n} p(d_i) S(d_i, \delta_j)\end{aligned}$$

Some Interesting Questions

- When can a player act without considering what the opponent will do? i.e. When is player 1's strategy independent of p or player 2's of q?
- When p or q is important, how can rationality of the opponent help us to elicit them?
- What are the implications of this?

Separable Games

If we can decompose the rewards appropriately, then there is no interaction between the players' decisions:

- A game is *separable* if:

$$R(d,\delta) = r_1(d) + r_2(\delta)$$
$$S(d,\delta) = s_1(d) + s_2(\delta)$$

- Here, the effect of the other player's act on a player's reward doesn't depend on their own decision:

$$\bar{R}(d_i) = r_1(d_i) + \sum_{j=1}^{m} q(\delta_j) r_2(\delta_j)$$

$$\bar{S}(\delta_j) = \sum_{i=1}^{n} p(d_i) r_1(d_i) + r_2(\delta_j)$$

Strategy in Separable Games

- Player 1's strategy should depend only upon r_1 as the decision they make doesn't alter the reward from r_2.
- Player 2's strategy should depend only upon s_2 as the decision they make doesn't alter the reward from s_1.
- So, player 1 should choose a strategy from the set:

$$D^\star = \{d^\star : r_1(d^\star) \geq r_1(d_i) \quad i = 1, \ldots, n\}$$

- And player 2 from:

$$\Delta^\star = \{\delta^\star : s_2(\delta^\star) \geq s_2(\delta_j) \quad j = 1, \ldots, m\}$$

The Prisoner's Dilemma is a Separable Game

- Let $r_1(S) = 0$ and $r_1(B) = 1$.
- Let $r_2(S) = -1$ and $r_2(B) = -5$.
- Now, $R(d, \delta) = r_1(d) + r_2(\delta)$.
- And $D^\star = \{B\}$.
- Similarly for the second player, $\Delta^\star = \{B\}$.
- This is the so-called paradox of the prisoner's dilemma: both players acting rationally and independently leads to the worst possible solution!

Rationality and Games

As in decision theory, a rational player should maximise their expected utility. We will generally assume that utility is equal to payoff; no greater complications arise if this is not the case.

- For a given pmf q, player 1 has:

$$\bar{R}(d_i) = \sum_{j=1}^{m} R(d_i, \delta_j) q(\delta_j)$$

- Whilst for given p, player 2 has:

$$\bar{S}(\delta_j) = \sum_{i=1}^{n} S(d_i, \delta_j) p(d_i)$$

- We want p and q to be consistent with the assumption that the opponent is rational.
- We assume, that *rationality of all players is common knowledge.*

Common Knowledge: A Psychological Infinite Regress

In the theory of games the phrase *common knowledge* has a very specific meaning.

- Common knowledge is known by all players.
- That common knowledge is known by all players is known by all players.
- That common knowledge is common to all players is known by all players

  ⋮
- More compactly: common knowledge is something that is known by all players and the fact that this thing is known by all players is itself common knowledge.
- This is an example of an infinite regress.

Domination

- A move $d^\star$ is said to dominate all other strategies if:

$$\forall d_i \neq d^\star, j : \qquad R(d^\star, \delta^j) \geq R(d_i, \delta_j)$$

- It is said to *strictly dominate* those strategies if:

$$\forall d_i \neq d^\star, j : \qquad R(d^\star, \delta^j) > R(d_i, \delta_j)$$

- A move d' is said to be *dominated* if:

$$\exists i \text{ such that } d_i \neq d' \text{ and } \forall j : R(d', \delta_j) \leq R(d_i, \delta_j)$$

- It is said to be *strictly dominated* if:

$$\exists i \text{ such that } d_i \neq d' \text{ and } \forall j : R(d', \delta_j) < R(d_i, \delta_j)$$

Theorem (Dominant Moves Should be Played)

If a game has a payoff matrix such that player 1 has a dominant strategy, $d^\star$ then the optimal move for player 1 is $d^\star$ irrespective of q.

Proof:

- *Player 1 is rational and hence seeks the d_i which maximises*

$$\sum_j R(d_i, \delta_j) q(d_j)$$

- *Domination tells us that* $\forall i, j : \quad R(d^\star, \delta_j) \geq R(d_i, \delta_j)$
- *And hence, that:*

$$\sum_j R(d^\star, \delta_j) q(d_j) \geq \sum_j R(d_i, \delta_j) q(d_j)$$

A similar results holds for player 2

Rationality and Domination

If rationality is common knowledge and $d^\star$ is a strictly dominant strategy for player 1 then:

- Player 1, being rational, plays move $d^\star$.
- Player 2, knows that player 1 is rational, and hence knows that he will play move $d^\star$.
- Player 2 can exploit this knowledge to play the optimal move *given that player 1 will play* $d^\star$.
- Player 2 plays moves $\delta^\star$ with $\delta^\star$ such that:

$$\forall j : S(d^\star, \delta^\star) \geq S(d^\star, \delta_j)$$

- If there are several possible $\delta^\star$ then one may be chosen arbitrarily.

Example (A game with a dominant strategy)

Consider the following payoff matrix:

	δ_1	δ_2	δ_3	δ_4
d_1	(2,-2)	(1,-1)	(10,-10)	(11,-11)
d_2	(0,0)	(-1,1)	(1,-1)	(2,-2)
d_3	(-3,3)	(-5,5)	(-1,1)	(1,-1)

- If rational, player 1 must choose d_1.
- Player 2 knows that player 1 will choose d_1.
- Consequently, player 2 will choose δ_2.
- (d_1, δ_2) is known as a discriminating solution.

Iterated Strict Domination

1. Let $D_0 = D$ and $\Delta_0 = 0$. Let t = 1
2. Player 1 checks D_{t-1} to see if it contains one or more strictly dominated moves. Let D'_t be the set of such moves.
3. Let $D_t = D_{t-1} \setminus D'_t$.
4. Player 1 checks D_{t-1} to see if it contains one or more strictly dominated strategies given that player 2 must choose a move from Δ_{t-1}. Let D'_t be the set of these strategies. Let $D_t = D_{t-1} \setminus D'_t$.
5. Player 2 updates Δ_{t-1} in the same way noting that player 1 must choose a move from D_t.
6. If $|D_t| = |\Delta_t| = 1$ then the game is solved.
7. If $|D_t| < |D_{t-1}|$ or $|\Delta_t| < |\Delta_{t-1}|$ let $t = t + 1$ and goto 2.
8. Otherwise, we have reduced the game to the simplest form we can by this method.

Example (Iterated Elimination of Dominated Strategies)

Consider a game with the following payoff matrix:

	L	C	R
T	(4,3)	(5,1)	(6,2)
M	(2,1)	(8,4)	(3,6)
B	(3,0)	(9,6)	(2,8)

Look first at player 2's strategies...

Example (Iterated Elimination of Dominated Strategies)

C is strictly dominated by R, leading to:

	L	R
T	(4,3)	(6,2)
M	(2,1)	(3,6)
B	(3,0)	(2,8)

Player 1 *knows* that player 2 won't play C…

Example (Iterated Elimination of Dominated Strategies)

Conditionally, both M and B are dominated by T:

	L	R
T	(4,3)	(6,2)

Player 2 *knows* that player 1 will play T and so, they play L.
Again, we have a deterministic “solution”.

Purely Competitive Games

- In a purely competitive game, one players reward is improved only at the cost of the other player.
- This means, that if $R(d', \delta) = R(d, \delta) + x$ then $S(d', \delta) = S(d, \delta) - x$.
- Hence $R(d', \delta) + S(d', \delta) = R(d, \delta) + S(d, \delta)$.
- The sum over all players' rewards is the same for all sets of moves.
- It doesn't change the domination structure or the ordering of expected rewards if we add a constant to all rewards.
- Hence, any purely competitive game is equivalent to a game in which:

$$\forall \delta \in \Delta, d \in D : R(d, \delta) + S(d, \delta) = 0$$

a *zero-sum game*.

Payoff and Zero-Sum Games

- In a zero-sum game:

$$S(d_i, \delta_j) = -R(d_i, \delta_j)$$

- Hence, we need specify only one payoff.
- Payoff matrices may be simplified to specify only one reward[6]

Example (Rock-Paper-Scissors is a zero-sum game)

	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

- It can be convenient to use standard matrix notation, with $M = (m_{ij})$ and $R(d_i, \delta_j) = m_{ij}$.

[6] In the two player case, at least.

What if no move is dominant?

- In the RPS game, like many others, no move is dominant (or dominated) for either player.
- If either player commits themself to playing a particular move, the other play can exploit that commitment (if they knew what it was, that is).
- We need a strategy for dealing with such games.
- Perhaps the maximin approach might be useful here...

Maximin Strategies in Zero-Sum Games

- If a player adopts a maximin strategy, he believes that the opponent will always correctly predict their move.
- This means, the opponent will choose their best possible action based upon the player's act.
- In this case, player 1's expected payoff is:

$$R_{\text{maximin}}(d_i) = \min_j R(d_i, \delta_j)$$

- If this is the case, then player 2's payoff is:

$$-R_{\text{maximin}}(d_i) = \max_j -R(d_i, \delta_j)$$

- Hence $P1$ should play $d^{\star}_{\text{maximin}} = \arg\max_{d_i} \min_j R(d_i, \delta_j)$.
- One could swap the two players to obtain a maximin strategy for player 2.

Example (RPS and Maximin)

- Let $M = (m_{ij})$ denote the payoff matrix for the RPS game.
- Then, $\min_j R(d_i, \delta_j) = \min_j m_{ij} = -1$ for all i.
- Thus any move is maximin for player 1.
- Player 1 expects to receive a payout of -1 whatever he does.
- If both players adopt a maximin view, then player 2 has the same expectation (by symmetry).
- How can we resolve this paradox?

What's Gone Wrong?

- The players aren't using all of the information available.
- They haven't used the fact that it is a zero sum game.
- They don't have compatible beliefs:
  - If P1 believes P2 can predict their move and P2 believes that P1 can predict their move then things inevitably go wrong.
  - It cannot be common knowledge that *both* players will adopt a maximin strategy!
- If a player really believes their opponent can predict their move then they can use randomization to make their action less predictable. . .

Mixed Strategies

- A *mixed strategy* for player 1 is a probability distribution over D.
- If a player has mixed strategy $\mathbf{x} = (x_1, \ldots, x_n)$ then they will play move d_i with probability x_i.
- This can be achieved using a randomization device such as a spinner to select a move.
- A *pure* strategy is a mixed strategy in which exactly one of the x_i is non-zero (and is therefore equal to 1).
- A similar definition applies when considering player 2.

Expected Rewards and Mixed Strategies

What is player 1's expected reward if...

- Player 1 has mixed strategy $\underline{x}$ and player 2 plays pure strategy δ_j?
- Player 1 has pure strategy d_i and player 2 plays mixed strategy $\underline{y}$?
- Player 1 has mixed strategy $\underline{x}$ and player 2 has mixed strategy $\underline{y}$?

In the first case, the uncertainty is player 1's own move, and his expectation is:

$$\sum_{i=1}^{n} x_i R(d_i, \delta_j)$$

In the second case, the uncertainty comes from player 2:

$$\sum_{j=1}^{m} y_j R(d_i, \delta_j)$$

Whilst both provide (independent) uncertainty in the third case:

$$\sum_{i=1}^{n} \sum_{j=1}^{m} x_i R(d_i, \delta_j) y_j = \underline{x}^\mathsf{T} M \underline{y}$$

Maximin Revisited

- Player 1's maximin *mixed* strategy is the $\underline{x}$ which minimises:

$$V_1 = \max_{\underline{x}} \min_{\underline{y}} \sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

- Player 2's maximin *mixed* strategy is the $\underline{y}$ which minimises:

$$\max_{\underline{y}} \min_{\underline{x}} - \sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

$$= \min_{\underline{y}} \max_{\underline{x}} \sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

- Which leads to a payoff for player 1 of:

$$V_2 = \min_{\underline{y}} \max_{\underline{x}} \sum_i \sum_j x_i R(d_i, \delta_j) y_j$$

Theorem (Fundamental Theorem of Zero Sum Two Player Games)

V_1 and V_2 as defined before satisfy:

$$V_1 = V_2$$

The unique value, $V = V_1 = V_2$ is known as the value of the game.

- The strategies $\underline{x}$ and $\underline{y}$ which achieve this value may not be unique.
- How can we find suitable strategies in general?

Example (Maximin in a Simple Game)

- Consider a zero sum two player game with the following payoff matrix:

	δ_1	δ_2
d_1	1	3
d_2	4	2

- With a pure strategy maximin approach:
  - P1 plays d_2 expecting P2 to play δ_2.
  - P2 plays δ_2 expecting P1 to play d_1.
  - P1 expects to gain 2; P2 expects to lose 3.
  - This is not consistent.

Example

- Consider, instead, a mixed strategy maximin approach:
  - P1 plays a strategy $(x, 1-x)$ and player 2 plays $(y, 1-y)$.
  - Player 1's expected payoff is:

$$[x \quad 1-x]\begin{bmatrix} 1 & 3 \\ 4 & 2 \end{bmatrix}\begin{bmatrix} y \\ 1-y \end{bmatrix} = -4(x-\frac{1}{2})(y-\frac{1}{4})+\frac{5}{2}$$

  - Player 1 seeks to maximise this for the worst possible y.
  - As the 2nd player can control the sign of the first term, his optimal strategy is to make it vanish by choosing $x = \frac{1}{2}$.
  - Similarly, the 2nd player wants to prevent the first player from exploiting the first term and chooses $y = \frac{1}{4}$.
  - Now, the expected reward for the first player is, consistently, 2.5 as both expect the same maximin strategies to be played.
  - *Both* players have a higher expected return than they would playing pure strategies.

How do we determine maximin mixed strategies?

- We need a general strategy for determining strategies $\underline{x}^\star$ and $\underline{y}^\star$ which achieve the common maximin return for player 1.
- It's straightforward (if possibly tedious) to calculate, for payoff matrix M the expected return for player 1 as a function of the strategies:

$$V(\underline{x}, \underline{y}) = \underline{x}^\top M \underline{y}$$

- We then seek to obtain $\underline{x}^\star, \underline{y}^\star$ such that:

$$V(\underline{x}^\star, \underline{y}^\star) = \max_{\underline{x}} \min_{\underline{y}} V(\underline{x}, \underline{y})$$

- In general, this is a problem which can be efficiently addressed by linear programming.
- If one player has only two possible decisions, however, a simple graphical method can be employed.

Graphical Solution, Part 1: Player 1's approach

- Consider a two player zero sum game with payoff matrix:

$$M = \begin{bmatrix} 2 & 3 & 11 \\ 7 & 5 & 2 \end{bmatrix}$$

- Consider a mixed strategy $(x, 1-x)$ for player 1.
- For the three pure strategies available to player 2, player 1 has expected reward:
  - $\delta_1 : 2x + 7(1-x) = 7 - 5x$
  - $\delta_2 : 3x + 5(1-x) = 5 - 2x$
  - $\delta_3 : 11x + 2(1-x) = 2 + 9x$
- For each value of x, the worst case response of player 2 is the one for which the expected reward of player 1 is minimised.
- Plotting the three lines as a function of x...

Zero-Sum Games

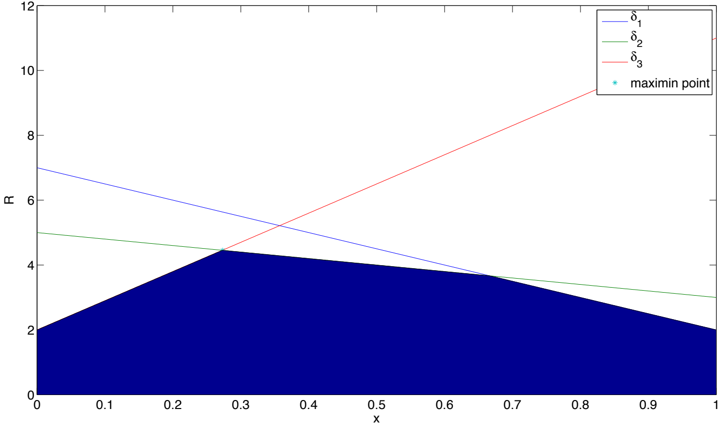

- The maximin response maximises the return in the worst case.
- In terms of our graph, this means we choose x to maximise the distance between the lowest of the lines and the ordinate axis.
- This is at the point where the lines associated with δ_2 and δ_3 intersect, at $x^\star$ which solves:

$$\begin{aligned} 5 - 2x =& 2 + 9x \\ 11x =& 3 \Rightarrow x^\star = 3/11 \end{aligned}$$

- Hence player 1's maximin mixed strategy is $(3/11, 8/11)$.
- Playing this, his expected return is:

$$V_1 = 2 + 9 \times 3/11 = 49/11 = \quad 5 - 2 \times 3/11 = 49/11$$

Graphical Solution, Part 2: Player 2's approach

- Player 2 only needs to consider the moves which optimally oppose player 1's maximin strategy (δ_2 and δ_3).
- They may consider a mixed strategy $(0, y, 1-y)$.
- By the fundamental theorem, player 2's maximn strategy leads to the same expected payoff for player 1 as his own maximin strategy:

$$V_2 = V_1 = 49/11.$$

- They should play $y^\star$ to solve:

$$\begin{aligned} V_2 = 3y + 11(1-y) =& 49/11 \\ 8y =& (121-49)/11 = 72/11 \Rightarrow y^\star = 9/11 \end{aligned}$$

- Leading to a mixed strategy $(0, 9/11, 2/11)$.

Example (Spy Game)

- A spy has escaped and must choose to flee down a *river* or through a *forest*. Their guard must choose to chasse them using a *helicopter*, a pack of *dogs* or a *jeep*.
- They agree that the probabilties of escape are as given in this payoff matrix:

	H	D	J
R	0.1	0.8	0.4
F	0.9	0.1	0.6

- Both players wish to adopt maximin strategies.

Example

- The spy plays strategy $(x, 1-x)$: with probability x they escape via the river; with probability $1-x$ they run through the forest.
- For given x, their probabilities of escaping for each of the guard's possible actions are:

$$\begin{aligned} p_H &= 0.1x + 0.9(1-x) \\ &= \frac{9-8x}{10} \\ p_J &= 0.4x + 0.6(1-x) \\ &= \frac{6-2x}{10} \end{aligned} \qquad \begin{aligned} p_D &= 0.8x + 0.1(1-x) \\ &= \frac{1+7x}{10} \end{aligned}$$

- Plotting these three lines as a function of x we obtain the following figure:

Zero-Sum Games

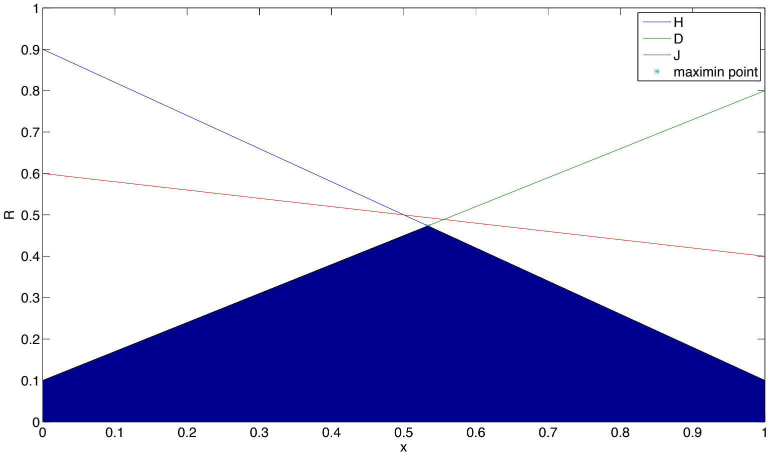

Example

- The maximin solution is the interesection of the lines for strategies D and H.
- This occurs at the solution, $x^\star$ of:

$$\begin{aligned} p_H = p_D \Rightarrow 9 - 8x =& 1 + 7x \\ 8 =& 15x \qquad \Rightarrow x^\star = 8/15 \end{aligned}$$

- The value of the game is: $V = V_1 = \frac{9-8x^\star}{10} = 71/150$

Example

- By the fundamental theorem of zero sum two player games, the guard needs to consider only H and D.
- Otherwise the spy's chance of escape will be better than V_1 if he plays his own maximin strategy.
- Consider a strategy $(y, 1-y, 0)$.
- By the same theorem, $V_2 = V = V_1$, so:

$$\begin{aligned} V_2 = 0.1y^\star + 0.8(1-y^\star) &= 71/150 \\ 8 - 7y^\star &= 71/15 \\ y^\star &= 7/15 \end{aligned}$$

On Zero Sum Two Player Games

- The "fundamental theorem" does not generalise to games of more than two players.
- The "fundamental theorem" does not generalise to non-zero-sum games.
- Games with an element of co-operation are much more interesting.

A Few Useful Concepts from Game Theory

- Maximin pairs provide a "solution" concept for zero-sum games.
- Some problems arise considering non-zero-sum games:
  - Maximin pairs don't necessarily make sense any more.
  - It's not obvious what properties a solution should have.
- In general, we consider ideas of equilbrium and stability.
- Notions of optimality and equilibrium:
  - Pareto optimality.
  - Nash equilibrium.

Pareto Optimality

- A collection of strategies (one per player) in a game is *(strongly) Pareto optimal/efficient* if no change can be made which will improve one players reward without harming any other player.
- A collection of strategies is *weakly Pareto optimal* if no change can be made which will improve all players' rewards.
- If a collection of strategies is not Pareto optimal then at least one player could obtain a better outcome with a different collection.
- In a game of pure conflict, all sets of pure strategies are Pareto optimal.

Nash Equilibrium

- A collection of strategies (one per player) in a game is a *Nash equilibrium* if no player can improve their reward by unilaterally changing their strategy.
- In the two-player case, mixed strategies $\underline{x}$ and $\underline{y}$ comprise a Nash equilibrium if:

$$\begin{aligned} \forall \underline{x}' : \quad & \bar{R}(\underline{x}, \underline{y}) \geq \bar{R}(\underline{x}', \underline{y}) \\ \forall \underline{y}' : \quad & \bar{S}(\underline{x}, \underline{y}) \geq \bar{S}(\underline{x}, \underline{y}') \end{aligned}$$

  where

$$\bar{R}(\underline{x}, \underline{y}) = \sum_{i=1}^{n} \sum_{j=1}^{m} x_i R(d_i, \delta_j) y_j \quad \bar{S}(\underline{x}, \underline{y}) = \sum_{i=1}^{n} \sum_{j=1}^{m} x_i S(d_i, \delta_j) y_j$$

- If the inequality holds strictly we have a *strict Nash equilibrium*.

Nash Equilibria in 2 Player Zero Sum Games

- Maximin pairs are equivalent to Nash equilibria: if $\underline{x}^\star$ and $\underline{y}^\star$ are maximin, then, by definition:

$$\begin{aligned} \forall \underline{x}' : \quad & \bar{R}(\underline{x}^\star, \underline{y}^\star) \geq \bar{R}(\underline{x}', \underline{y}^\star) \\ \forall \underline{y}' : \quad & \bar{S}(\underline{x}^\star, \underline{y}^\star) \geq \bar{S}(\underline{x}^\star, \underline{y}') \end{aligned}$$

  A similar argument holds in the reverse direction.
- All equilibria have the same expected payoff (this follows from the fact that $S = -R$).
- These properties do not extend to non zero-sum games.

Nash Equilibria and the Prisoner's Dilemma

- Recall the prisoner's dilemma:

	S	B
S	(-1,-1)	(-5,0)
B	(0,-5)	(-4,-4)

- (B, B): both players betraying one another is a pure-strategy Nash equilibrium.
- (S, S): both players remaining silent is Pareto optimal: no change can be made which leads to improvement for one player and no worsening of the other player's situation.
- The (S, S) strategy set is not stable: it is not an equilibrium as either player can unilateral improve their own reward.

Solutions I: The Nash Sense

- Two pairs $(\underline{x}, \underline{y})$ and $(\underline{x}', \underline{y}')$ are interchangeable with respect to some property if $(\underline{x}', \underline{y})$ and $(\underline{x}, \underline{y}')$ have the same property.
- A game is *Nash solvable* if all equilibrium pairs are interchangeable (with respect to being equilibrium pairs).
- All zero-sum games are Nash solvable.
- Not many other games are.

Solutions II: The Strict Sense

- A game is *solvable in the strict sense* if:
  - Amongst the Pareto optimal pairs there is at least one equilibrium pair.
  - The equilibrium Pareto optimal pairs are interchangeable.
- The solution to such a game is the set of equilibrium Pareto optimal pairs.
- In a zero sum game, all strategies are Pareto optimal and so this reduces to the notion of Nash solvability: all zero sum games are solvable in the strict sense.

Solutions III: The Completely Weak Sense

- A game is *solvable in the completely weak sense* if after iterated elimination of dominated strategies, the reduced game is solvable in the strict sense.
- The solution is then the strict solution of the reduced game.
- In a zero sum game no strategies are dominated and so this reduces to the notion of solvability in the strict sense: all zero sum games are solvable in the completely weak sense.

Solutions and the Prisoner's Dilemma

- The only equilibrium pair of this game is (B, B).
- The only Pareto optimal strategy is (S, S).
- The game is Nash Solvable, with solution (B, B).
- The game is not solvable in the strict sense: no Pareto efficient pair of strategies is an equilibrium pair.
- The game is solvable in the completely weak sense:
  - S is a dominated strategy for both players.
  - The reduced game after IEDS has a single strategy (B) for each player.
  - The strategy (B, B) is Pareto efficient in the reduced game (no other strategy exists).
  - (B, B) is an equilibrium pair in the reduced game.
  - The solution set is (B, B).