Ben Graham

ST114 Decisions and Games

University of Warwick — Winter 2012/13

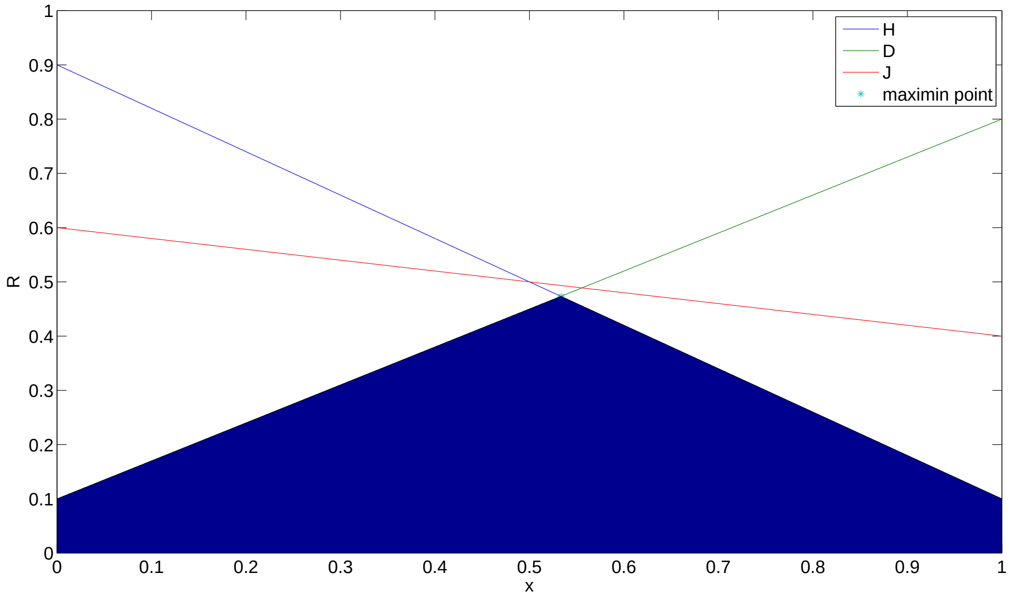

Warwick Statistics

Administrative Details

Lecturer	Dr Ben Graham (these typed lecture notes a small perturbation of the lecture notes of Dr Adam Johansen from when he taught the course)
Email	`b.graham@warwick.ac.uk`
Lectures	20
	Tuesday 16:00
	Friday 13:00
CATS	7.5
Assessment	100% Closed Book Examination
Exercise sheets	`http://www2.warwick.ac.uk/fac/sci/statistics/modules/st1/st114/resources`

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.

Books. There are a great number of books on the subjects of this course...

– You don't *need* to buy any of them.
– Many are available in the library.
– Three are particularly well suited to the present course:

  - Although not a traditional maths book, Körner's recent "Naïve Decision Making" quite closely matches the syllabus of this course and is extremely readable.
  - Peterson's recent "Introduction to Decision Theory" matches the course syllabus quite closely. It is written by a professional philosopher and has a rather different style.
  - Jim Smith's "Bayesian Decision Analysis: Principles and Practice" is excellent and was published last year. It covers much of the material in this course and much more besides. It would be a good investment if you plan to study further Bayesian inference/decision theory courses in later years.
- 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"

Assessment. A few points about the assessment of this module. There will be one 90 minute examination in which you must answer two questions from three. These questions may not have the precisely same form as those asked in previous years and may not follow the same division between parts of the course as they have in previous years. A good understanding of the course will be required to do well in the exam: please don't mistake a course with relatively little technical mathematical content for one which requires no effort.

Always practice good exam technique. The following may seem obvious, but experience shows that the following points are worth making:

- Show *all* rough work. Indicate it as such if you wish, but don't erase it or scribble it out.
- Always explain what you are doing.
- Explain why things are true: don't just write a sequence of equations.
- Symbolic logic is not preferable to a lucid explanation.
- If asked to *prove* or *show* something then do it rigorously.
- Don't write text in pencil. If you wish to draw diagrams in pencil then make sure they are clear.

Calculators are permitted (and may be needed) in the examination. Any calculator used must comply with the statistics departments' standing rule on calculators, reproduced here from the 2012/13 MORSE handbook:

```
Concerning the use of calculators in examinations the Department of Statistics follows
the University rule which states that except for the display of error or function
messages, calculators with non-numeric displays are not allowed. In other words
prohibited calculators are those which can accept alphabetical data. Note that this
includes most graphical calculators of the type acceptable in GCSE and A-level
examinations. It is your responsibility to ensure that your calculator fulfils the
Universitys criterion and that your calculator is not of the prohibited type. Otherwise
you may find yourself denied the use of your calculator and be involved in disciplinary
proceedings.

Suggested suitable calculators for incoming students which are in line with
recommendations from the Computer Science Department are Casio fx 82, fx83 or fx85. All
of these are available from SU and from well known retailers. They are also reasonably
priced.
```

Content of these Notes

1. Introduction 5
1.1 The basis of decision analysis 5

2. Probability 10
2.1 Axiomatic Probability 10
2.2 What do we mean by probability... Objectively? 16
2.2.1 Probability and Relative Frequencies 19
2.3 What do we mean by probability... Subjectively? 19

3. Elicitation 26
3.1 Elicitation of Personal Beliefs 26
3.2 Axiomatic and Subjective Probability Combined 28

4. Conditions 33
4.1 Conditional Probability 33
4.2 Useful Probability Formulæ 35
4.3 Random Variables and Expectations 37

5. Making Decisions 42
5.1 Decision Problems 42
5.2 Decision Trees 45
5.3 Decision Trees — Example 48

6. Preferences 55
6.1 The Trouble With Money 55
6.2 Utility 57

7. Games 63
7.1 What is a Game? 63
7.2 Separability and Domination 66
7.3 Zero-Sum Games 71
7.4 Selected Game-Theoretic Concepts (Not Examinable) 78

8. Closing Comments 81

1. Introduction

We begin by considering what a decision analyst does before going on to develop the necessary tools.

1.1 The basis of decision analysis

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: there may be some 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.

Essentially all problems of decision analysis can be cast within this framework. The objective in such problems is invariably to try to make the *best* decision, in some sense, given all of these difficulties. That is, the decision which is likely to achieve a good outcome and which balances the available risks and rewards taking into account the client's knowledge, objectives and preferences.

A problem of two parts. Such decision problems really contain two sub-problems. The first is finding out what the client knows and wants. This may not be a completely trivial exercise if you are dealing with someone who isn't particularly mathematically literate. This phase is known as *elicitation*: the drawing out of that knowledge which the client has in some form and rendering it precise in a mathematical formalism.

Elicitation involves obtaining precise answers to several questions:

- What is the client's problem?
- What does she believe?
- What does she want?

The second phase is then one of more mundane calculation. Given all of the information obtained from the client, we must work out two things:

- What are its logical implications?

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

– What should our client do?

In answering this final question, we have solved the decision problem. The other steps are essential precursors and, in some settings, may yield additional information which is of value in itself but the real task of the decision analyst is to provide the client with rigorously-justified advice on how they should act. Likely consequences and risks can also be identified at this stage.

Pragmatically, it may be necessary to apply this procedure iteratively: eliciting the client's beliefs and performing a calculation based upon them would be enough in an idealisation of the problem, but in practice your client may wish to change their beliefs once they have seen what consequences they have. This can lead to an iterative procedure of elicitation and calculation:

$$\text{Elicitation} \longrightarrow \text{Calculation} \longrightarrow \text{Elicitation} \longrightarrow \text{Calculation} \longrightarrow \dots$$

Care should be exercised when doing this, of course. It isn't a good idea to keep updating the input you provide to a calculation until the output corresponds to that which was desired before it was done! If that's what's happening then there's no point in carrying out a decision analysis. However, it may be reasonable for a client to want to modify the statement of their beliefs if the decision analysis reveals that their original statement does not really encode what they believe at all.

What does she really want?. Finding out what a client wants may seem to be a straightforward exercise. However, it is not. The following example illustrates some of the difficulties which might have to be faced:

Example 1.1 (Advising a university undergraduate). If you were attempting to advise a new undergraduate student on how they should act at university you should, ideally, know what it is that they are trying to achieve by being there. What constitutes a reasonable balance between academic and other considerations will depend upon the outcome that they are hoping for. So it is necessary to work out their objective, is it:

– 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?

◁

One problem is clear here: most — but not all — students would attach at least some weight to each of the first four criteria. In order to make progress, it will be necessary to establish how important each of these things is to at least a semi-quantitative level. This isn't straightforward: the assignment is a personal and subjective one, but the client is not likely to have a set of numbers associated with each of these aspects of the experience and there may well be other considerations for each individual.

It might seem that many of the problems faced by businesses will be easier to address than this one and, to some extent, this is true as businesses tend to have more clearly defined objectives and their owners are likely to be able to express reasonably concisely what it is that they want. Of course, this isn't always the case and it isn't always as simple as it might at first seem. Consider the following example:

Example 1.2 (A small business owner). Imagine advising a small business owner on the pricing policy they should adopt and what proportion of their profits should be reinvested. 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?

◁

Even in this simple setting in which it might at first have seemed that the objective was clearly "to make as much profit as possible" it quickly becomes clear that their are complications which have to be addressed. These arise depending upon several factors: the timescale on which that profit is to be made; whether making money fast is more valuable than making more money in the medium term and whether there are other parties involved who will have their own pricing policy.

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?

Each of these four questions can be turned into a formal mathematical one. However, it's usually necessary for the decision analyst to extract this information from their client in an inexact and incomplete form and to turn that information into a formal specification. It's useful to consider another example to see what sort of information we might realistically be able to extract in a real problem.

Example 1.3 (Marketing). If advising a marketing executive on the advertising strategy which they should adopt for a new product, it is unlikely that you will be provided with the information you need in the form that you want it. However, there are numerous questions which the executive is likely to be able to answer for you:

– 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?

◁

In this example, the first questions tells us what advertising strategies are possible. The next four are all related to the collection of possible outcomes and the relationship between these and the decision we make. The last question also tells us something about the interaction between our client and other involved parties but in reality we would probably also have to consider people other than direct competitors in a detailed analysis.

Two more examples show that it is possible for different parties to reach different conclusions when presented with essentially the same decision problem. It will later become clear that this can happen even when both parties have the same information available to them and are acting within the framework of decision theory. The examples also show that this is perhaps a desirable property of any theory which attempts to codify the way in which humans both do, and wish to, make decisions.

Example 1.4 (Insurance). Consider advising a client on whether to insure against a particular sort of loss. You might also be called upon by the insurance company to advise them whether they should be prepared to offer insurance on similar grounds. Of course, you shouldn't be advising both parties at the same time!

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

◁

If p is in a broad range of values, it is likely that the client will take out insurance and the insurer will be happy to offer it. One question which will be addressed later is this: Why are both parties happy with this? In order to answer this question, we will have to establish how we can determine the value that different parties might attach to the same outcome. People want different things and value things differently; organizations have different priorities again.

Example 1.5 (A Simple Lottery). – $\mathbb{P}(\{\text{Win}\}) = 1/10,000$
– $Value\,(\text{Win}) = £5,000$
– Ticket price $£1$.

That is, a ticket can be purchased for $£1$ and each ticket has an equal chance of winning (1/10,000). The prize of $£5,000$ is less than the ratio of the cost of the ticket to the probability of winning, as is usually the case with lotteries. ◁

Why is this acceptable to the customer? It is also acceptable to the provider of the lottery but for different reasons. Simple variations of the lottery could be considered in which, say, the ticket price remains $£1$ but $\mathbb{P}(\text{Win}) = 1$ and $Value(\text{Win}) = £0.5$. Whilst the lottery organiser might be equally – or more – happy with this lottery than the other, it is unlikely that they would find many customers. This example suggests that people treat a small chance of a large reward differently to a small reward which they are much more likely to receive.

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 1.6 (Economic forecasting). Forecasts of British GDP growth in 2009 given at the end of 2008:

– -0.1% – International Monetary Fund
– -0.75– -1.25% British Government
– -1.1% Organisation for Economic Co-operation and Development
– -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 (We will put aside the philosophical questions raised by this concept. . .)? Actually, the figure was worse than any of these predictions at -3.7%. ◁

A problem which you are more likely to encounter as a decision analyst is that many people, even experts, are prone to confusing what they really believe will happen with what they want to happen. This can significantly distort the estimates they provide to a decision analyst and it's important to be aware of this potential problem.

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. The choice of which action to make is the decision.
- A *game* is a similar problem in which at least some of the uncertainty arises from the behaviour of a (rational) opponent.

From Questions to Answers. Having established what types of problems decision theory attempts to address, we will need to answer the following questions before we can make much progress:

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.

2. Probability

Before we can quantify risks, beliefs and rewards, we need a formal framework for the description of uncertainty. We also need to know what that framework means and how to interpret it. Section 2.1 will introduce a rigorous mathematical framework for the description of probabilistic events. This framework will be based upon a minimal collection of axiomatic statements. Inevitably, this will lead to some duplication of material covered in other first year courses. Whilst this duplication will be kept to a minimum it is to some extent unavoidable and experience from previous years indicates that it may not be undesirable. The remainder of this chapter is dedicated to a discussion of the possible interpretations of probability and and proving that probability, as we shall interpret it, must behave in the same way as probability as defined by our axiomatic framework. This is a subtle point, but an important one.

2.1 Axiomatic Probability

We want a theory of probability which allows us to perform calculations in the presence of uncertainty. In a purely-mathematical setting it is possible to specify a few desiderata and from these to obtain a theory of probability. This theory is a mathematical abstraction and it will remain to be shown that it corresponds in a meaningful way to our everyday experiences.

Foundations of An Axiomatic Theory of Probability. The *Russian school* of probability is based on axioms. These were first formulated by A. A. Kolmogorov who laid many of the foundations of a beautiful and powerful area of mathematics. Before it is possible to state those axioms, it is necessary to define some quantities and to determine what elements a theory of probability must contain.

The abstract specification of probability requires three things:

1. A set of all possible outcomes, Ω. This is sometimes termed the *sample space*. For $\omega \in \Omega$, the singleton set $\{\omega\}$ is called an *elementary event*.
2. A collection of subsets of Ω, $\mathcal{F}$. In some sense, these subsets of Ω describe the outcomes of interest: we may not be particularly interested in the particular value within Ω which corresponds to an event which occurs, but only in a coarse property of the event and a proper specification of $\mathcal{F}$ can dramatically simplify the problem.
3. Finally, we will require a function which assigns a probability to our events:

$$\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$$

Notice that this function acts not on the points of Ω, but on the elements of $\mathcal{F}$. These are *subsets* of Ω.

When a mathematician refers to the probability of an event, in the notation used here, the event is an element, E, of $\mathcal{F}$ and the probability is $\mathbb{P}(E)$. The statement is not meaningful unless Ω, $\mathcal{F}$ and $\mathbb{P}$ are specified — although in many real world examples there will be a natural choice. When we do know what all of these things are, the meaning is intended to be the following: the "probability" of an event $E \in \mathcal{F}$, $\mathbb{P}(E)$ is a measure of how likely it is that, on conducting a random experiment described by the collection we have specified, the outcome which actually occurs (an element of Ω) will be in E.

It is useful to think about a few examples in which probability might be a useful concept and to consider what these objects might be in these settings.

Example 2.1 (Simple Coin-Tossing). Consider the probabilistic experiment in which a standard coin is tossed.

– All possible outcomes are "heads" and "tails" which can be abbreviated, giving us:

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

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

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

– In which case, under reasonable assumptions we might suppose that:

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

◁

Notice that at this stage we have said nothing about the properties which $\mathcal{F}$ or $\mathbb{P}$ must have and so this specification is rather arbitrary.

Example 2.2 (A Tetrahedral (4-faced) Die). Rolling a die with 4 faces.

– 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 $A \in \mathcal{F}$:

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

◁

Even in this small example, we see that $\mathcal{F}$ is becoming larger and the specification of $\mathbb{P}$ looks a little esoteric. Finally, let's consider a more complicated example:

Example 2.3 (The National Lottery). – $\Omega = \{\text{All unordered sets of 6 numbers from}\{1, \ldots, 49\}\}$
– $\mathcal{F}$ = All subsets of Ω
– Again, we can construct $\mathbb{P}$ from expected uniformity.

◁

But in this example 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. In this case, it would still be possible to work with this object in a mechanical way, and it wouldn't produce any fundamental difficulties, but what would we do if $\Omega = \mathbb{R}$?

It's often easier not to work with *all* of the subsets of Ω. But if we want to work with a smaller collection of subsets of Ω than the set of all it's subsets then it's necessary to make sure that our collection is big enough to allow for meaningful statements about probability to be made and to avoid leading to any contradictions. Remember, we have said so far that a probability function maps only elements of this collection to real numbers: we cannot obtain the probability of any event which does not lie in $\mathcal{F}$.

We will need to consider particular types of sets of subsets of Ω; those which contain Ω itself as an element and which are closed under certain important operations.

Definition 2.1 (Algebras of Sets). *Given Ω, $\mathcal{F}$ must satisfy certain conditions.*

1. $\Omega \in \mathcal{F}$ *In terms of probabilities, this means that the event "something happening" is in our set and can be ascribed a probability.*
2. *If $A \in \mathcal{F}$, then*
$$\Omega \setminus A = \{x \in \Omega : x \notin A\} \in \mathcal{F}$$
Hence, if A happening is in our set then $\Omega \setminus A$ happening (that is, 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).*

It might seem that we need to include some other sets as well in order for this to allow us to describe all events of interest, but with one exception the above definition is a complete one. We don't, for example, need to specify that events corresponding to both A and B happening must lie in the set if both A and B do as this follows from the conditions specified already. However, if Ω is not a finite set then a little more care is required and, in fact one additional condition must be met by $\mathcal{F}$ in order to avoid some technical problems involving limits.

Definition 2.2 (σ-Algebras of Sets). *If, in addition to meeting the conditions to be an algebra, $\mathcal{F}$ is such that:*

– *If $A_1, A_2, \cdots \in \mathcal{F}$ 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.

It's useful to look at another example at this stage.

Example 2.4 (Selling a house). – You wish to sell a house, and you hope to make at least £250,000.

– The estate agent has promised to filter out any offer below £100,000 to avoid wasting our time
– 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?

A good starting point would be to specify the sample space – the collection of possible outcomes of the two offers being made. In the situation described above, this is simply all pairs of numbers exceeding $£100,000$.

$$\Omega = \{(x,y) : x, y \geq £100,000\}$$

But, we only care about the much coarser events corresponding to one or other of the two offers being larger; that is, sets:

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

It's important to include some others to ensure that we have an algebra:

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

Thus in an experiment in which the number of possible distinct outcomes is rather enormous, we have reduced the problem to one with just a handful of events to which it is necessary to ascribe probabilities.

As an aside, if we require that X and Y take only integer values then it would, at least, have been possible to assign probabilities to any subset of the outcome space and to work with a very large collection — in such a case we have simplified the problem but not changed its character in a qualitative way. However, had x and y been allowed to take real values it would not even be possible to do this — the introduction of the σ-algebra is essential to allow for a meaningful collection of probabilities to be established. This will be covered in detail in ST213: Mathematics of Random Events. ◁

In this example, an algebra was constructed by first considering the particular events of interest (that one or other bid was strictly larger than the other) and then augmenting these sets with the others which are required to provide us with an algebra. One question is whether any of the elements of a general algebra are distinguished in a similar sort of way.

Atoms. The answer to that question is very often yes. Some events are *indivisible* and somehow fundamental:

Definition 2.3 (Atoms). *An event $E \in \mathcal{F}$ is said to be an* atom *of $\mathcal{F}$ if:*

1. *$E \neq \emptyset$. The empty set is not an atom.*
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. This is the property of indivisibility that makes atoms interesting and gives them their (physically inaccurate) name.

If $\mathcal{F}$ is finite (that is, it contains only finitely many subsets of Ω) then for 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 $\mathcal{F}$. That is, we can represent any event as a combination of atoms.

Returning to the algebra of the house-selling example,

$\{(i,j):i<j\}$	$\{(i,j):i>j\}$	$\{(i,j):i\neq j\}$	$\emptyset$
$\{(i,j):i\leq j\}$	$\{(i,j):i\geq j\}$	$\{(i,j):i=j\}$	Ω

It's interesting to consider 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. Given a set of possible outcomes, Ω, and a σ-algebra of its subsets, $\mathcal{F}$, the final ingredient required to describe the system probabilistically is a function which maps the elements of $\mathcal{F}$ to probabilities. Such a function must have certain properties in order to correspond to the way that we expect probability to behave.

$\mathbb{P}:\mathcal{F}\to\mathbb{R}$ is a *probability measure* over $(\Omega,\mathcal{F})$ iff[1]:

1. For any $A\in\mathcal{F}$:
$$\mathbb{P}(A)\geq 0$$
All probabilities are positive.
2. $$\mathbb{P}(\Omega)=1$$
Something certainly happens. Remember that Ω is the collection of all possible outcomes and the probability of Ω is the "chance" that the outcome which actually occurs lies in Ω. By convention, a probability of 1 is interpreted as certainty.
3. For any[2] $A,B\in\mathcal{F}$ such that $A\cap B=\emptyset$:
$$\mathbb{P}(A\cup B)=\mathbb{P}(A)+\mathbb{P}(B)$$
Probabilities are (sub)additive. If the outcome of the experiment can lie in at most one of A and B then the chance that it lies in either A or B is simply the chance that it lies in A added to the chance that it lies in B.

Again, if we are dealing with more complicated spaces then it can be necessary to introduce an additional technical condition to ensure that limits behave as we would expect them to.

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.

[1] Iff is a convenient shorthand for "if and only if".
[2] This is sufficient if Ω is finite; we need a slightly stronger property in general.

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. Rather than considering only pairs of disjoint events, we consider any sequence of such events (which is at most countably infinite).

It's useful to think a little about probability measures: they are complicated objects. First, why are they so-named and what are they intended to achieve?

- 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].

So, for example, the *Lebesgue measure* over (a particular class of) subsets of the real numbers tells us how long an interval is and has generalisations to higher dimensional spaces corresponding to area, volume and so on. This is outside the scope of this course. In advanced probability texts, the triple $(\Omega, \mathcal{F}, \mathbb{P})$ comprising a sample space, a σ-algebra of subsets of that space and a probability measure associated with that σ-algebra is termed a probability space.

It's potentially difficult to develop a probability measure (which must assign a probability to all events within the σ-algebra of our probability space directly. A number of techniques exist for constructing probability measures (or probability distributions, as they are often called) from simpler objects. In discrete algebras (those which can be described in terms of their atoms) probability mass functions are often used.

Definition 2.4 (Probability Mass Function). *Let $\mathcal{F}$ be an algebra containing finitely many atoms $E_1, \dots, E_n$. 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* $\sum_{i=1}^{n} p_i = 1$.

The following technique allows us to take a probability mass function and turn it into a full probability measure (which is unique and consistent with that probability mass function).

Proposition 2.1 (From masses to measures). *Let $S = \{A_1, \dots, A_n\}$ be such that:*

1. *The elements of S are disjoint: $\forall i \neq j : A_i \cap A_j = \emptyset$*
2. *S* covers *Ω: $\cup_{i=1}^{n} A_i = \Omega$*

Then:

1. *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.*
2. *The atoms of the generated algebra are the elements of S.*
3. *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).

Proof of this proposition is left as an exercise.

2.2 What do we mean by probability… Objectively?

So what?. It is all very well to define a mathematical theory of probability in this way, but if it is to be used to make decisions then it is necessary for probabilities to have some meaning and for the *events* of our theory to have some connection with events in the world at large. 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 several other important aspects of the problem haven't been addressed at all. The following points are critical if probabilities are to be used to describe uncertainty in decision problems:

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

The remainder of this chapter seeks to address these problems. In particular, two ways to interpret probabilities are given and it is shown that within these interpretations it is not just reasonable but necessary that the probabilities of real events should behave in the same way as the theoretical probabilities described in the previous section.

It is very convenient to have an objective interpretation of probability and this is a good starting point. Before justifying the use of the axioms of probability when such an interpretation exists, and ascribing a precise interpretation to probabilities in this setting, it is convenient to consider a class of problems in which there is a natural probability measure (if it is assumed, for the time being, that the axioms of probability hold; we will subsequently demonstrate that this must be so under the interpretation of probability which we will use).

Geometry, Symmetry and Probability. If probabilities have a geometric interpretation, we can often deduce probabilities from symmetries. That is, if all of the elementary outcomes of an experiment should be equally likely, and we want probability to obey our mathematical axioms then we can assign probabilities to these atoms by exploiting that symmetry.

Example 2.5 (Coin Tossing Again). Return to the problem of tossing an unbiased coin.

– Here, $\Omega = \{H, T\}$ and $\mathcal{F} = \{\emptyset, \{H\}, \{T\}, \{H, T\}\}$
– Axiomatically: $\mathbb{P}(\Omega) = P(\{H, T\}) = 1$. A head or tail is obtained with probability 1 (i.e. certainty).
– 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: $\mathbb{P}(\{H\}) = \mathbb{P}(\{T\}) = \frac{1}{2}\mathbb{P}(\{H, T\}) = \frac{1}{2}$.

◁

Example 2.6 (Tetrahedral Dice Again). Similar reasoning can be applied to the four-faced die considered earlier.

– 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 2.7 (Lotteries Again). Symmetry arguments can often be applied to more complicated problems, such as that of the national lottery.

– $\Omega = \{$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. In fact, this is the case when $\mathcal{F}$ is the set of all subsets of Ω.
– 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 $< i_1, \ldots, i_6 >$.

◁

In all three of these examples, we identified a finite collection of equally-probable atoms and developed a probability mass function for that collection of atoms. This mass function can then be extended to give us a full probability measure by proposition 2.1. We will revisit these systems later.

Complete Spatial Randomness and π. Other arguments based upon geometry and symmetry can also lead to natural probability distributions in some circumstances. The following is a famous pedagogical example of something known as the "Monte Carlo method" in which random samples from a probability distribution are used to provide estimates of a quantity of interest, in this case π. Of course, this example isn't really of practical usefulness, but it provides a useful illustration.

– Let (X, Y) be the coordinates of a point which is equally likely to occupy any position within a unit square with centre (0,0).
– Define

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

– Now

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

What this description really does is assign an informal notion of "uniform over the unit square" or something which is often referred as *complete spatial randomness*. If all points are equally probably then the probability that a sampled point lies in any subset of the space is proportional to its area. The details

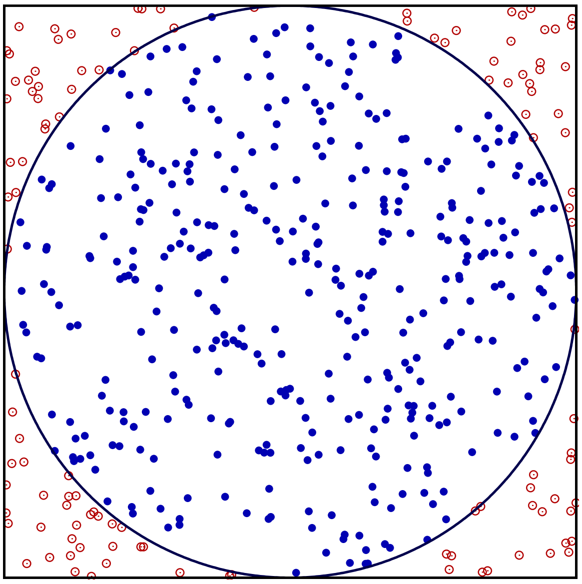

Fig. 2.1. A primitive stochastic estimation technique allows the approximation of π via this technique: consider the proportion of "uniform raindrops" landing within the circle... as the number of raindrops goes to infinity, this proportion converges (with probability 1, it turns out) to $\pi/4$.

of this kind of argument are a little subtle as they rely on probability distributions over *uncountably infinite* spaces and hence on a more complicated specification of the σ-algebra than has been necessary in the discrete cases covered here. This is not within the scope of this course but is fascinating and covered in detail in ST213.

Balls in Urns. Two canonical model systems are useful when considering probabilities: these two systems provide flexible abstract systems in which we can mimic almost any probability distribution in which we are likely to be interested.

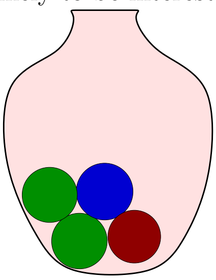

Let $\mathcal{I}$ be a (discrete) set of colours.
An urn contains n_i balls of colour i for each $i \in \mathcal{I}$.
The probability that a drawn ball has colour i is: $n_i \Big/ \sum_{j \in \mathcal{I}} n_j$.

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

Spinners. The other system which will be useful is the spinner.

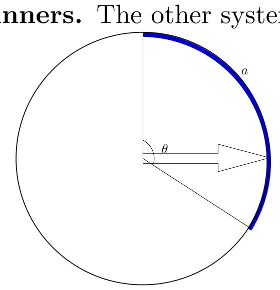

$\mathbb{P}[\text{Stops in arc}] = a$ (choose a radius of $1/2\pi$)
Really a statement about physics.
What do we mean by probability?

Both of these model systems will be useful in what follows as concrete examples of random objects which take values from discrete or continuous collections. We will later establish that the probabilities ascribed to these systems here by assuming that the axioms of probability hold are, essentially, the only probabilities which could sensibly be assigned to them.

2.2.1 Probability and Relative Frequencies

A Frequency Interpretation. A classical *objective* interpretation of probabilities is provided by the *frequentist* paradigm, which you may have encountered previously. Within this framework, it is only possible to consider the probability of events which occur as the outcome of an experiment which could, at least hypothetically, be repeated an arbitrary number of times.

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 the *indicator function* (which is sometimes referred to as the characteristic function by some analysts) is defined

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

- Equivalently, if $n_A(n)$ is the number of times A occurs during the first n repetitions of the experiment, then

$$\lim_{n\to\infty} \frac{n_A(n)}{n} = p_A$$

Thus the frequentist framework ascribes a precise meaning to probabilities: they correspond to long-run frequencies of occurrence. Although other interpretations of probability might share this property within the frequentist formalism this is the *definition* of probability.

2.3 What do we mean by probability... Subjectively?

Subjective Probability. What is the probability of a nuclear war occurring next year? Answering this question is clearly not simple, but it's not clear how we could possibly interpret this "probability" in the frequentist setting:

- First, we must be precise about the question: how exactly is nuclear war defined; when is next year?
- We can't appeal to symmetry or geometry. There is no obvious geometric formulation of our problem and there is no natural symmetry.
- We can't appeal meaningfully to an infinite ensemble of experiments. There's no way in which one can work out what proportion of worlds within such an ensemble would succumb to nuclear war.
- We *can* form an individual, *subjective* opinion. As with many real-world problems, most people have views about this issue and these personal opinions are based upon varying degrees of knowledge and expertise.

If we adopt this subjective view, difficulties emerge. Two of these are essential procedural difficulties: if one is to quantify a degree of belief, how can it be done and, furthermore, if we do this, can we be sure that we will end up with an internally consistent approach. Another, more philosophical, question but one which clearly has a great deal of impact on whether it is reasonable to make use of the conclusions of such a formulation is simply this: what do the calculations we make in this way actually tell us?

The remainder of this chapter will be dedicated to partially answering these questions. It should be noted, that opinions differ on this issue and there do exist forms of decision theory which rely only upon frequentist probability. That isn't the approach which will be followed in this course; the much greater breadth of problems that can be addressed within the subjective framework is one of many arguments that can be raised in favour of the approach which we will consider.

Bayesian/Behavioural/Subjective Probability. Rather than restricting ourselves to objective probabilities obtained by considering those things for which it is easy to quantify probabilities, one could begin by asking what one wants probabilities for and how they are to be used.

Taking this approach might lead to the conclusion that it would be very useful if all uncertainty could be represented using probability. In fact, this approach dates back to the eighteenth century:

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

Within this framework there is a natural way to frame and address all problems of inference and decision. This approach does not directly specify what probabilities are or what they tell us. In the last century, some Bayesians made some much bolder statements. Bruno de Finetti et al. believed that: Probability is *personalistic* and *subjective*. That is, not only is it reasonable that different people might attribute different values to the probability of a particular event but the mechanism by which they choose those values may not be objective. In order to take this view, and we shall do so in this course as it is a convenient and pragmatic one when dealing with problems of inference and decision making although it is not universally accepted, one must first establish what is it that probabilities mean and then that these probabilities should behave as described by the formal axioms of probability.

A Behavioural Definition of Probability. de Finetti et al. proposed the following device for attributing meaning to subjective probabilities:

- 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.
- Equivalently $m(M, A)$ is the *minimum* that *You* would accept to offer the bet.
- Requiring that you are prepared to both place or accept the bet at this price ensures that it is a *fair* price for the bet (clearly, it is necessary that M is sufficient small that you would be prepared to place or accept a bet at a fair price).
- Two events A_1 and A_2 are equally probable if $m(M, A_1) = m(M, A_2)$.
- A value for $m(M, \Omega \setminus A)$ is implied for a rational being... as we shall see.

Note that in this interpretation of probability it is the equality of bets on events which allows us to assign specific values to personal, subjective probabilities in a meaningful way. The additional structure imposed by this mechanism leads almost inevitably to adopting the axiomatic formalism of probability within the subjective framework.

The behavioural approach to probability (or something like it) is required because if the numbers we assign to probabilities are subjective then their must be some mechanism by which we can relate them to the way in which we act or they are essentially meaningless. That is:subjective probabilities are only meaningful if they characterise the behaviour of their holder.

A Bayesian View of Symmetry. If we return to the symmetric events we considered earlier, we can consider a simple abstraction:

- 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, j, we should set $m(M, A_i) = m(M, A_j)$ and so

$$\mathbb{P}(A_i) = c$$

  for some constant c.
- We would like the subjective probabilities of our Bayesian interpretation to coincide with the frequentist probabilities which we saw previously suggesting that $c = 1/k$.
- We can assign values to bets on this system of urns using a combination of logic and the assumption that the each event is equally likely.
  - Consider a bet, $b(M, \cup_{i=1}^k A_i)$ which pays M if one of the exhaustive outcomes occurs. This certainly happens and so we must value the bet and the reward equally: $m(M, \cup_{i=1}^k A_i) = M$.
  - Conversely, a bet, $b(M, \emptyset)$ which pays M if an impossible event occurs. This never happens and so the bet is worthless: $m(M, \emptyset) = 0$.
  - Logically, a bet on any event A in the algebra generated by $A_1, \ldots, A_k$ must have a value $m(M, A) \in [0, A]$ as there is no better bet with reward M than one which is won with certainty and no worse bet than one which is certainly lost.
  - We can compare the collection of bets $b(M, A_1), b(M, A_2), \ldots, b(M, A_k)$ with $b(M, \cup_{i=1}^k A_i)$ and observe that the first will pay M for each of the events A_i which occur; the second pays M if at least one of the A_i occur. As the A_i are mutually exclusive the two are *logically equivalent* and so:

$$\sum_{i=1}^{k} m(M, A_i) = m(A, \cup_{i=1}^k A_i) = k \Rightarrow m(M, A_i) = 1/k.$$

  - The same logical equivalence argument allows us to assign values to bets which pay on any event which may be expressed as a union of the A_i.
  - Thus we have demonstrated that the $m(M, A)$ viewed as a function of A obey the second and third of Kolmogorov's axioms.
  - We have the freedom to choose any one-to-one function to relate P to m.
  - If we allow ourselves to identify $P(A) = m(1, A)$ then we establish that $P(A) = m(1, A) = m(M, A)/M$ obeys Kolmogorov's axioms. Although we'd be free to choose any other relationship between P and m, as the behaviour of m is fully constrained, we'd arrive at an exactly equivalent formulation and there's no point in making life more complicated than it needs to be.
- Think of the examples we saw before...

Discretised Spinners. In the case of the spinner:

- Each of k segments is equally likely:

$$\mathbb{P}[\text{Stops in shaded arc}] = 1/k.$$

- k may be very large.
- Combinations of arcs give all rational lengths (length p/q can be obtained by dividing the spinner into q segments and then combining p of them).

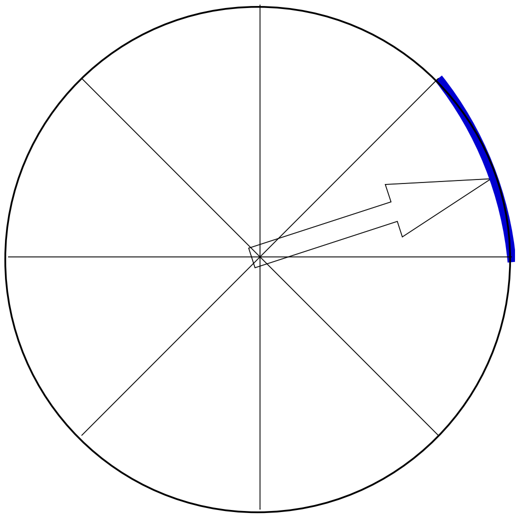

Fig. 2.2. We can divide the circumference of the spinner into n arcs of equal length to produce a discrete approximation of arbitrary accuracy.

– Limiting approximations give real lengths.
– We can describe *most* subsets this way [ST213].

In this case, the notion of symmetry corresponds to a statement of belief about the physics governing the system: essentially, if we accept that the spinner is equally likely to stop in any of the k segments, then it is because we believe that there is no way to predict where it will stop.

Returning to another of our examples, we can see an example of a concept slightly weaker than that of all outcomes having equal probability.

Example 2.8 (House selling again). – The three atoms in this case were:

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

– There is no reason to suppose all three are equally likely.
– However, 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 $i > k$ for some threshold k...

◁

Elicitation. What probabilities does someone assign to a complex event? In order for subjective probability to be useful, it is necessary that we can assign a value to an individual's personal probabilities. One way to do this is by comparison with a reference event.

In fact, we can use our behavioural definition of probability. The *urn* and *spinner* we introduced before have probabilities which we all agree on — those assigned by symmetry arguments. Even if probability is subjective, there are some simple situations in which almost everyone is likely to arrive at the same conclusions. We can use these to *calibrate* our personal probabilities: When does an *urn* or *spinner* bet have the same value as one of interest? And, hence, when do those events have the same probability as the event of interest? There are some difficulties with this approach, but it's a starting point.

A First Look At Coherence. This section informally motivates the ideas which we will make more rigorous in what follows.

Why should personal, subjective probabilities behave in the same way as those mathematical objects which we call probabilities that arise from Kolmogorov's axioms?

Consider a collection of events $A_1, \dots, A_n$. 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 $p_1 = \mathbb{P}(A_1), \dots, p_n = \mathbb{P}(A_n)$ for these events is *coherent* if:

- $\forall i \in \{1, \dots, n\} : p_i \in [0, 1]$
- $\sum_{i=1}^{n} p_i = 1$

Assertion 2.1. A *rational being* will adjust their personal probabilities until they are coherent.

We will look at this and related concepts in some detail in the following chapter, but for now we can get some flavour of the concept by considering the following argument.

Definition 2.5 (Dutch Book). *A collection of bets for which one party cannot win (or alternatively, the other cannot lose) is known as a Dutch book.*

For now, assume that a "rational being" would not be prepared to place a collection of bets as a result of which they cannot win. This should not seem unreasonable, but we will make this concept precise slightly later. If a collection of probabilities is incoherent, then a Dutch book can be constructed. Consequently, a rational being must have coherent personal probabilities.

In support of these assertions, which will ultimately be proved, we can consider some simple examples. We will use the identity $m(1, A) = P(A)$ during this argument.

Example 2.9 (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 $£(2 - 5/3)X = X/3$.

◁

This isn't a very interesting example, though, and it's perhaps not very plausible that such a situation could arise. Here's a more complicated situation in which the same idea works.

Example 2.10 (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? In the lecture one strategy will be investigated, but try to devise your own strategy and complete the following table:

Horse	Odds	Stake
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	

◁

The following example of a Dutch book was constructed in the lecture:

Example 2.11. 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. These ideas, as with many concepts from the race course, have close parallels in the world of economics. Unfortunately, you'll probably find it rather difficult to take this idea and use it to make your fortune.

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 in such a market are coherent.

3. Elicitation

3.1 Elicitation of Personal Beliefs

One of the questions which was raised at the beginning of these notes was, how can we assign numerical probabilities to our client's subjective beliefs. In this section some simple possibilities are presented. The extraction and quantification of such subjective beliefs is known in the Bayesian literature as *elicitation*; it is a much more complicated problem than it may at first appear and it should be realised that the approaches presented here are the simplest ones likely to lead to reasonable results. It may be necessary to employ more sophisticated techniques in some problems.

What does she believe?. If probability is the method which we wish to use to describe uncertainty and to encode the beliefs of our client, then we need to obtain and quantify our clients beliefs as statements about numerical probabilities.

Asking for a direct statement about personal probabilities doesn't usual work. There are numerous problems that are likely to occur, for example, if we ask what the probability of A happening is and then go on to ask a number of other questions which give us our client's personal probability for A not happening then we may find that $\mathbb{P}(A)+\mathbb{P}(A^c) \neq 1$ And, if you recall the British economy example, people confuse belief with desire (or in that case, perhaps, what they want others to believe that they themselves believe).

Behavioural Approach to Elicitation. A better approach uses comparison with a standard. To see how this might work, it's useful to consider a simple example in which obtaining numbers directly is far from simple.

Example 3.1 (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.

And also consider a bet which pays £1 if a spinner stops in an arc of length a and nothing otherwise.

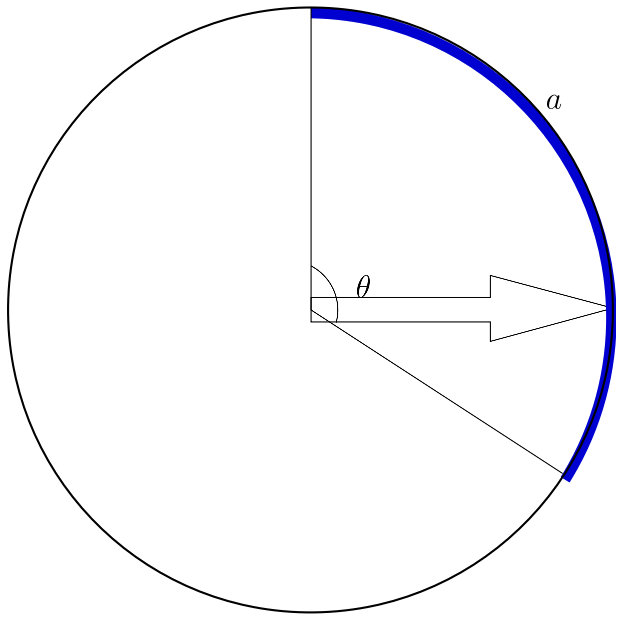

These two bets are illustrated in figure 3.1.

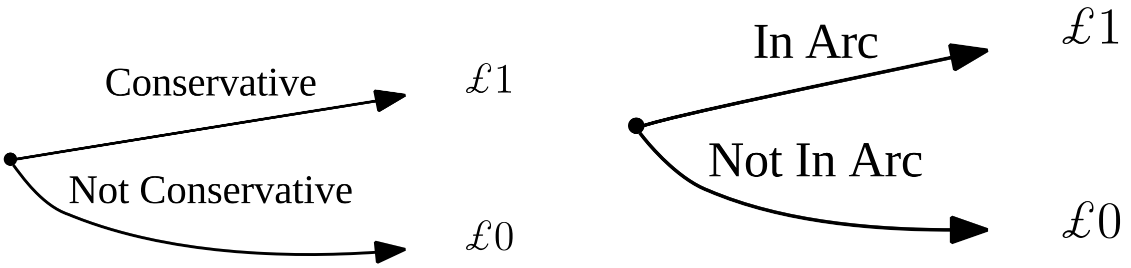

Fig. 3.1. Comparable bets: for what value of a are these two bets of the same value?

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. There must be an arc length above which we prefer the spinner bet and below which we prefer the political one. If we can determine this length and we accept the symmetry argument for the probabilities associated with the spinner, then knowing this arc length will allow us to calculate the subjective probability we associate with a Conservative victory at the next general election. ◁

Eliciting With Urns Full of Balls. Using a spinner gives us one way of eliciting probabilities. However, it may prove difficult to decide exactly how long the arc needs to be for us to have no preference between the two bets. A simple alternative is provided by the urn of balls.

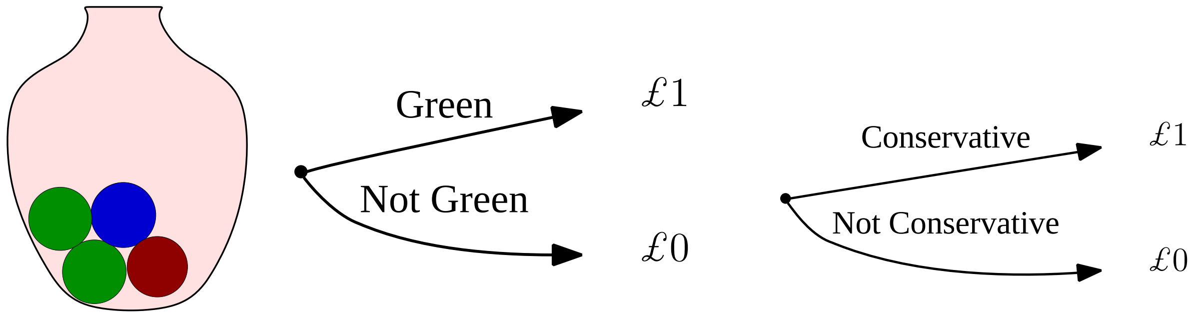

If an urn contains a fixed number of balls, n, how many of them must be green in order for us to prefer a bet on drawing a green ball from the urn to the political bet? One procedure for finding out is, allowing g to denote the number of green balls in the urn: Increase g from 0 to n and let $g^\star$ 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:

– $\mathbb{P}(\text{C.}) \geq \mathbb{P}\left(\text{Green ball drawn from } n \text{ when } g^\star \text{ are green}\right) = g^\star/n$
– $\mathbb{P}(\text{C.}) \leq \mathbb{P}\left(\text{Green ball drawn from } n \text{ when } g^\star + 1 \text{ are green}\right) = (g^\star + 1)/n$

We can elicit any probability using this technique and obtain a nominal accuracy of $1/n$ (and as we can choose n this means we can in principle obtain an estimate to as great an accuracy as we wish although it is sensible not to make n larger than it needs to be in order to obtain the required degree of accuracy).

3.2 Axiomatic and Subjective Probability Combined

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

We began with axiomatic probability. We then introduced a subjective interpretation of probability. We wish to combine both aspects so that we have a rigorous mathematical framework in which we can manipulate probabilities and perform calculations while at the same time having a way to relate the results to the real world in general and decision problems in particular.

Coherence Revisited. The idea of coherence was introduced briefly in order to support the idea that combining axiomatic and subjective probability was a reasonable thing to do. It is possible to make this concept precise.

Definition 3.1. *Coherence An individual, $\mathcal{I}$, may be termed* coherent *if her probability assignments to an algebra of events obey the probability axioms.*

Assertion 3.1. A rational individual must be coherent.

A Dutch book argument in support of this assertion follows. The simplest way to show that this is true is to show that any violation of these axioms will lead to a situation in which anyone accepting the behavioural view of probability (i.e. that one should act as dictated by ones personal probabilities) will act irrationally in that they will accept a collection of bets that will certainly lose them money.

Theorem 3.1. *For any event, A, any rational individual, $\mathcal{I}$, must have* $\mathbb{P}(A) + \mathbb{P}(A^c) = 1$.

Proof. There are two ways in which this can fail to hold.
Case 1: $\mathbb{P}(A) + \mathbb{P}(A^c) < 1$

Consider an urn containing n balls.

– Let $b^u(n,k)$ pay £1 if a green ball is drawn from an urn containing n balls of which k are green.
– Bet $b(A)$ pays £1 if event A happens.

Let $n^\star(A)$ and $n^\star(A^c)$ be the smallest integers such that bets $b^u(n, n^\star(A))$ and $b^u(n, n^\star(A^c))$ are preferred to bets $b(A)$ and $b(A^c)$, respectively.

As $\mathbb{P}(A) + \mathbb{P}(A^c) < 1$, for large enough n there exists some $k > 0$ such that:

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

for some $k > 0$. Think of an urn containing balls of *three* colours. The specified probability assignments imply that drawing the first colour is considered more likely than A occurring and drawing the second is considered more likely than A^c occurring. The fallacy of this position is clear as one of A and A^c must occur whilst it's possible that neither of the first two colours of ball will be drawn. To make this more precise, let the first set of balls be green and consider these two systems of bets:

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

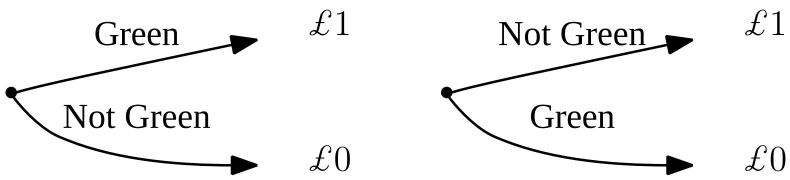

– System 2: $S_1^e = [b(A), b(A^c)]$

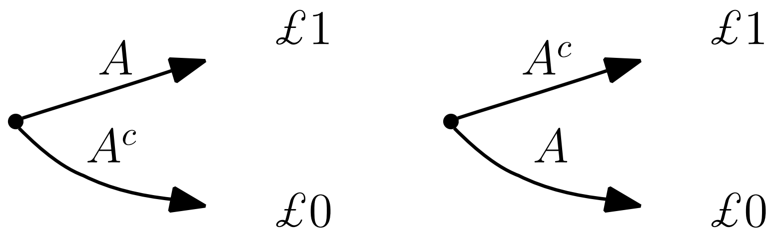

As $\mathcal{I}$ is indifferent to which of the left hand bets they make and strictly prefers the right hand bet of S_1^u to that of S_1^e (in the sense that they consider it to have greater value), $\mathcal{I}$ prefers S_1^u to S_1^e. Thus $\mathcal{I}$ would pay more for a bet in which they win if any of the bets of S_1^u win than they would charge someone for a bet in which $\mathcal{I}$ pays out if any of the bets of S_1^e win.

Equivalently, $\mathcal{I}$ should be prepared to pay some small amount to win on the bets of S_1^u and lose (and hence pay) on those of S_1^e.

But S_1^e is exhaustive: exactly one event is certain to happen, so $\mathcal{I}$ will certainly pay out if they do this and there's a chance that the third colour of ball will be drawn in which case $\mathcal{I}$ doesn't win anything. Their position is irrational as they value a system of bets in which they might win 1 unit or may win nothing more highly than one in which they certainly win 1 unit. Behaviourally, they commit to paying an amount $c > 0$ for a bet in which they certainly lose 1 unit and possibly win 1 unit; in the best case the two bets cancel and they have payed c for nothing; in the worst they pay c and then lose 1 unit.

Case 2: $\mathbb{P}(A) + \mathbb{P}(A^c) > 1$

Using a similar argument as previously, we know that the number of balls required n, for large enough n there exists $k > 0$ such that the number of balls required are

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

Consider an urn with $m^\star(A)$ green balls and $m^\star(A^c) - k$ blue (and hence $n = m^\star(A) + m^\star(A^c) - k$ balls in total) and the following systems of bets:

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

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

The stated probabilities mean, $\mathcal{I}$ will pay £c to win on S_2^e and lose on S_2^u (as they value the right hand bet in S_2^e more highly than that in S_2^u). However, everything cancels: $\mathcal{I}$ will always win 1 unit and lose 1 unit so there's no logic in paying an additional amount to do so.

In summary: a rational individual won't pay for a bet which certainly returns £0. If $\mathbb{P}(A) + \mathbb{P}(A^c) \neq 1$ then there is a contradiction: so $\mathbb{P}(A) + \mathbb{P}(A^c) = 1$.

□

Theorem 3.2. *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 (outline). Again, there are two situations which must be considered. The proof proceeds in the same way as that of theorem 3.1.

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

The probabilities are such that given an urn containing a large enough number of balls, n, the number of balls which must be a particular colour for a bet that the drawn ball being of that colour to be comparable to a bet on the event of interest are such that there exists $k > 0$ such that:

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

where $n^\star(A)$ is the number of the n balls which must be of a given colour in order for drawing a ball of that colour to be as likely as the occurrence of event A and so on.

Let

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

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

$$S_3^{e'} = [b(A \cup B)]$$

As $\mathcal{I}$ thinks the right hand bet in S_3^u is more valuable than that in S_3^e (and the left hand bets are of equal value) they will pay a small amount to win if S_3^u happens but lose if S_3^e does. They value S_3^u and $S_3^{e'}$ equally and so will (at neutral cost) simultaneously place a bet in which they win if $S_3^{e'}$ occurs but lose if S_3^u does. However, the bets on S_3^u cancel one another and the remaining two systems of bets are logically equivalent.

Hence they will pay to win and lose on equivalent events! A rational individual wouldn't do so and hence we have a contradiction.

Similar reasoning holds when $\mathbb{P}(A) + \mathbb{P}(B) > \mathbb{P}(A \cup B)$. □

It may help to see a concrete example of how things go wrong:

Example 3.2 (Football betting). – Football team C is to play A. A friend says:

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

- This is annoying as it is clearly illogical and violates the coherence requirements that a rational being should follow. However, there is a course of action available which will have one of two possible outcomes: your friend allows you to place bets with them such that they will certainly lose or they accept that the probabilities which they have asserted are essentially meaningless as they aren't prepared to support them with action.
- Consider an urn containing 7 balls; 6 are green and the "sure-thing" system of bets:

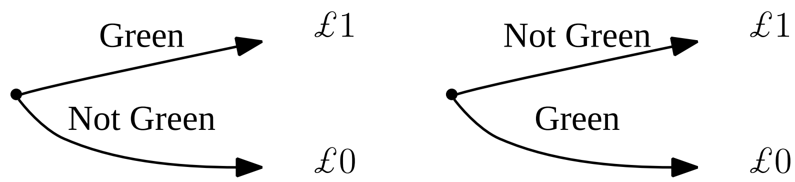

- Your friend's probabilities mean that the two urn bets are inferior to $b(C)$ and $b(A)$, respectively. Your friend should pay $£c$ to win on $[b(A), b(C)]$ but lose on the urn system.
- If they agree, you could simply take $£c$ as a small profit. Alternatively, you could argue that logically, $b(C)$ and $b(A)$ are not exhaustive (there may be a draw). So your friend should pay a little to switch back to a system in which they cannot lose. Iterate until your point has been made.
- If your friend refuses you can quite reasonably argue that their "probabilities" are meaningless.

◁

There are certain other things which must be proved in order for assertion 3.1 to have been proved. These are dealt with on exercise sheet 2.

The Cox-Jaynes Axioms. It can also be shown that any quantification of belief must obey certain rules under some rather weak assumptions. The following does not require that we identify degrees of belief with probabilities from the outset and by convention we refer instead to "degrees of plausibility".

If we want the following (so-called Cox-Jaynes axioms) 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.

In general this can make it difficult to reach conclusions upon which everyone can agree. However, in the context of decisions analysis for individuals it is convenient: it means that we can deal with the actual beliefs of the individual making the decision rather than some collective objective assessment with which they may not agree.

– Elicited probabilities should be coherent.

The decision analyst must ensure this.

Whatever an individual believes, if we represent their strengths of belief using real numbers then we should do so in a way consistent with the axioms of probability. It is likely that if asked to give such a set of numbers, most clients will not provide a decision analyst with a coherent set of values and so it is important for the analyst to use techniques which ensure that they obtain a coherent representation of their client's beliefs.

– Temporal coherence is not assumed or assured.

You are permitted to change your mind.

That is, although the set of beliefs must at any instant of time correspond to a coherent collection of probabilities, it is possible for them to change from one time instant to another. The client's strength of belief in the proposition that it will rain tomorrow afternoon may not be the same as the strength of their belief in the proposition tomorrow morning. This is re-assuring, but how *should* we update our beliefs?

4. Conditions

So far we have dealt with simple probabilities and elicitation. One extremely important consideration when conducting inference generally and decision analysis in particular is how our beliefs should be updated in light of new information. This chapter is concerned with relating the idea of conditional probability to the subjective interpretation of probability and showing that Bayes' rule and some related results are all that is required (at least in principle) to allow us to update our beliefs in a logical and coherent fashion.

4.1 Conditional Probability

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. That is, we can view this definition as the statement that knowing that B has occurred tells us that whatever the outcome of the experiment is it must lie in the set B but providing no additional information. That is, the *relative* probability of any event contained within B is unaltered by the knowledge that B has occurred. Thinking of a Venn diagram for the events, A, B and $A \cap B$ may make this clearer.

Example 4.1 (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 4.2 (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. We need to know that these objects have a meaningful interpretation when we are dealing with subjective probabilities.

The standard construction for relating conditional probabilities to the behaviourally-justified subjective probabilities of de Finetti is the *called-off bet*. Consider a type of bet, $b(A|B)$ which pays

– £1 if A happens *and* B happens,

– nothing if B happens but A does not (any stake, i.e. payment made in order to make the bet is lost),

– nothing and is called off (any stake is returned) if B does not happen.

We can view this bet graphically:

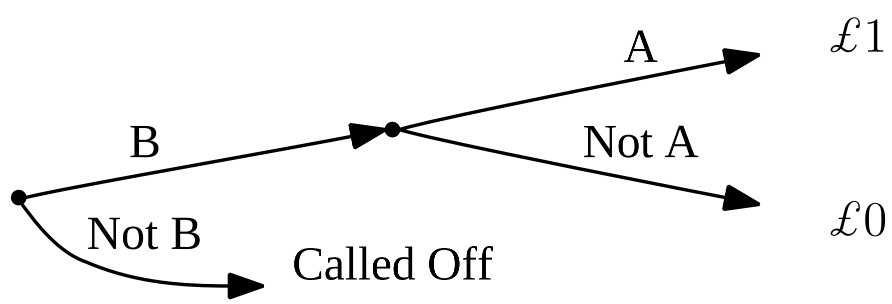

How would a rational being value such a bet? In order to answer this question we can relate this bet to a simple bet on A made only once we've observed that B occurs. Logically, these two bets must have the same value.

Theorem 4.1 (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 the number of balls of each colour such that $\mathcal{I}$ is values a bet on a red ball being drawn to be of the same value as one on $A \cap B$ and so on.

Logically, a bet on B or B^c is of the same value as one on (red or blue) or on (green or yellow), respectively.

Consider a second bet: B occurs. What are the probabilities $\mathcal{I}$ attaches to A and A^c conditional upon this? Given an urn containing 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 (considering all possible outcomes when one ball is drawn from each urn), we may deduce that:

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

Hence

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

The proof was covered in more detail in the lecture, but the important elements are those covered here.

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 4.1 (Independence of Events). *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. Notice that this definition of independence applies to random *events* although it is closely related to the concept of independence of random *variables* (see section 4.3).

4.2 Useful Probability Formulæ

A number of simple results from probability are indispensable when performing inference in general and for decision making in particular. Some of the most useful of these are summarised here.

The first result of interest simply tells us that if exactly one of a collection of events must happen then any other event can be decomposed as the union of its intersection with each of those events. This can be extremely useful as it allows the calculation of the probability of complicated events as the sum of the probabilities of mutually exclusive, simpler events. Although the result itself may seem rather trivial it finds widespread application.

Theorem 4.2 (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 (from the same probability space). 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.

Combining this result with Bayes' law provides an expansion which is particularly useful for elicitation and the calculation of conditional probabilities more generally:

Theorem 4.3 (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 $\mathbb{P}(A|B_i)$. □

This is directly useful when we have access to conditional probabilities, $\mathbb{P}(A|B_i)$ given the individual elements of the partition and the marginal probabilities $\mathbb{P}(B_i)$ but don't know $\mathbb{P}(A)$. It may be less obvious that it can also be extremely useful for accurate elicitation:

Example 4.3 (Buying a house). Consider advising a client on whether or not they should by 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:

$$\mathbb{P}(A) = \sum_{i=1}^{3} \mathbb{P}(A|E_i)\mathbb{P}(E_i)$$

where

$$E_1 = \text{[Inflation is low.]}$$
$$E_2 = \text{[Inflation is high; salary rises]}$$
$$E_3 = \text{[Inflation is high; salary doesn't rise]}$$

◁

This is just a simple example of a general principle. It is typically easier to accurate express opinions about simpler events in which various aspects of uncertainty are separated out and then to combine our estimates than it would be to obtain an accurate assessment of the overall probability.

The core of Bayesian analysis is the following elementary result:

Theorem 4.4 (Bayes' Law). *If A and B are events of positive probability, then:*

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

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

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

In order to evaluate the denominator of this fraction, it's often useful to note that $\mathbb{P}(A|B) = \frac{\mathbb{P}(A)\mathbb{P}(B|A)}{\mathbb{P}(A)\mathbb{P}(B|A)+\mathbb{P}(A^c)\mathbb{P}(B|A^c)}$ or to otherwise employ the partition formula to evaluate this probability.

This allows us to update our beliefs.

Example 4.4 (Disease Screening). Consider screening a rare disease.

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

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 $\mathbb{P}(B|A^c)$ to be useful.

Although a positive test provides some evidence to support the hypothesis that the subject is diseased, because the probability that they had the disease before this evidence was observed was so low the increase in probability associated with the new evidence is not large enough to make us confident that they have the disease. This has implications throughout the world of diagnostics and elsewhere. ◁

– 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.

The example illustrates that conditional probabilities $P(A|B)$ and $P(B|A)$ not only need not be (and generally aren't) the same, but they can be very different. One issue with elicitation is that if asked for $P(A|B)$ most people would actually give an assessment of their view of $P(B|A)$, although one would hope that experienced mathematicians would not fall into this trap. This is a simple consequence of the way in which people thing of the problem in a forward direction.

There is a direct implication of this effect: when performing elicitation in examples like this it is much more likely that a good result will be obtained if $P(B|A)$ and $P(A)$ are elicited and Bayes' law used to determine $P(A|B)$ than it is if $P(A|B)$ is elicited directly.

4.3 Random Variables and Expectations

In previous probability courses it is likely that the fundamental objects you have encountered are random variables. These can be extremely useful and we will wish to make use of them in the process of our decision making. However, it is convenient to think of the events of our probability space as being fundamental with random variables arising from those. This is the approach which is taken in more advanced probability and has a number of advantages — including the fact that it is straightforward to define complicated collections of related random variables.

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 in $\mathcal{X} = \{x_1, \dots, x_n\}$. 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 $\mathbb{P}$. This is *one way* of constructing a probability space whose sole function is to provide a definite interpretation of a particular random variable. In general, it is often more useful to take an already given probability space and to associate with it one or more random variables of interest. For example, one could associate with a probability space which models the rolling of two fair dice a random variable which describes the sum of the values shown of each die and another corresponding to the product of those values.

As illustrated in the next example, perhaps the best way to think of random variables is as functions which assign particular values of elementary events in the sample space. Provided that the associated algebra is rich enough it is then possible for us to assign probabilities to events corresponding to the random variable taking particular values. This formalism can be easily extended to deal with collections of several random variables and random variables which take continuous values (you could think of the angle at which a spinner comes to rest, for example) although the second of these things involves certain technicalities which we won't be dealing with in this course. Indeed, for the purposes of this course it is sufficient to consider only random variables which take their values in a discrete set of possible outcomes.

Example 4.5 (Roulette). Consider spinning a roulette wheel with $n(r) = n(b) = 18$ red/black spots and $n(g) = 1$ green one, which we can model as a discretised spinner.

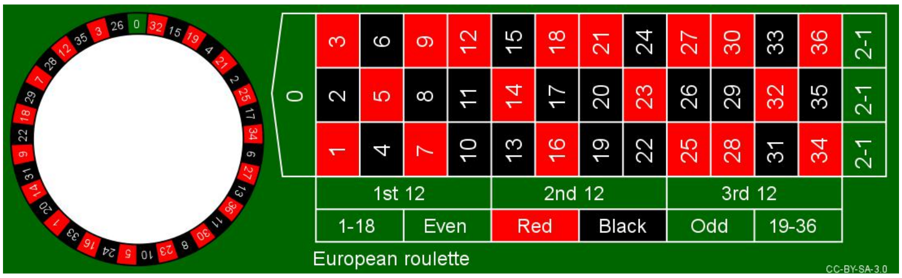

Set X to -1 if the ball stops in a red region, +1 for a black one and +39 for a green. This corresponds to the casino's profit on a spin of the wheel in which customers have staked 20 on red and 19 on black.

Under a suitable assumption of symmetry, the probability mass function

$$\begin{aligned}\mathbb{P}[X = -1] &= n(r)/n \\ \mathbb{P}[X = +1] &= n(b)/n \\ \mathbb{P}[X = +39] &= n(g)/n\end{aligned}$$

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

We will not worry too much about the specific probability spaces with which our random variables are associated in this course, but the mapping from a probability space to the real line interpretation is extremely general and you will encounter it in much more detail in any subsequent courses you take on probability.

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

Definition 4.2. *Random variables, X and Y, are independent if for all possible x_i, y_j:*

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

You can think of this in terms of events associated with random variables. A random variable taking a particular value is an event and the definition above essentially says that providing all events which can be defined in terms of one random variable are independent of all events which can be defined in terms of a second random variable then it is reasonable to consider the random variables themselves as being independent.

[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 4.3. *The expectation of a random variable, X, is:*

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

where the sum is taken over all possible values.

When introduced this quantity was known as *mathematical expectation* and in English language versions of much of the earlier literature the symbol M was used to denote it. This is not a small point: mathematical expectation has a precise interpretation as the centre of probability mass and it does, in this sense, characterise the typical behaviour of the random variable. However, it need not have many of the properties that you would expect a "typical" value to have. For example, the expected value of a random variable, X, need not even be a value which it is possible for X to take.

Example 4.6 (Number of Heads). Let X be the number of heads observed when a fair coin is tossed.

$$\mathbb{P}(X = 0) = \mathbb{P}(X = 1) = \frac{1}{2}$$

Hence

$$\mathbb{E}[X] = \sum_{i=0}^{1} i\mathbb{P}(X = i) = 0 \times \frac{1}{2} + 1 \times \frac{1}{2} = \frac{1}{2}$$

which is a perfectly reasonable average value, but we wouldn't *expect* to see $\frac{1}{2}$ heads if we tossed the coin. ◁

Useful Properties of Expectations. One of the reasons that expectations are so widely used is that they have a number of useful properties.

– Expectation is linear: if we have two random variables X and Y and three real-valued constants, a, b and c, then:

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

This statement can be proved by considering the definition of the expectation. It compactly encodes a number of statements:

$$\mathbb{E}[aX] = a\mathbb{E}[X]$$
$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$
$$\mathbb{E}[c] = c$$

It is also possible to employ these techniques iteratively to show, for example, that given n random variables, $X_1, \ldots, X_n$:

$$\begin{aligned}
\mathbb{E}\left[\sum_{i=1}^{n} X_i\right] =& \mathbb{E}\left[X_1 + \sum_{i=2}^{n} X_i\right] \\
=& \mathbb{E}\left[X_1\right] + \mathbb{E}\left[\sum_{i=2}^{n} X_i\right] \\
=& \mathbb{E}\left[X_1\right] + \mathbb{E}\left[X_2\right] + \mathbb{E}\left[\sum_{i=3}^{n} X_i\right] \\
=& \sum_{i=1}^{n} \mathbb{E}\left[X_i\right]
\end{aligned}$$

and we will make use of these results later.

– 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 of this is that: a function of a random variable is itself a random variable. If X takes values $x_i \in \Omega$ with probabilities $\mathbb{P}[X = x_i]$ then $f(X)$ takes values $f(x_i)$ in $f(\Omega)$:

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

If we simply replace all of the values which X is allowed to take by mapping them through f then we obtain a new random variable which behaves in exactly the same way as the object obtained by realising the random variable X and passing the realised value through that function.

One can also think of $\mathbb{E}\left[f(X)\right]$ as being the centre of probability mass of X after X is transformed by the function f.

– These properties may naturally be combined, so that expectations of linear combinations of functions of random variables may be expected as linear combinations of the expectations of those functions, e.g.:

$$\begin{aligned}
\mathbb{E}\left[(X - \mathbb{E}\left[X\right])^2\right] =& \mathbb{E}\left[X^2 - 2X\mathbb{E}\left[X\right] + \mathbb{E}\left[X\right]^2\right] \\
=& \mathbb{E}\left[X^2\right] - 2\mathbb{E}\left[X\right]\mathbb{E}\left[X\right] + \mathbb{E}\left[X\right]^2 = \mathbb{E}\left[X^2\right] - \mathbb{E}\left[X\right]^2
\end{aligned}$$

notice that when evaluating the outer expectation here, the inner expectation which is just a number may be treated just like any other constant and so it doesn't introduce any nonlinearity here.

– There is a limit to the operations which can be justified. For general nonlinear $f(x)$, $E(f(X)) \neq f(E(X))$. You can verify this easily by considering a random variable which takes values ± 1 with equal probability and allowing $f(x) = x^2$.

Example 4.7 (Die Rolling). Consider rolling a six-sided die:

$\Omega = \{1, 2, 3, 4, 5, 6\}$ Let X be the number rolled. This is equivalent to letting X be the identity mapping, so that $X(\omega) = \omega$ for any $\omega \in \Omega$.

Under a symmetry assumption:

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

Hence, the expectation, and that of $f(X) = X^2$, is:

$$\mathbb{E}[X] = \sum_{x \in \Omega} x\mathbb{P}[X = x] = \sum_{x=1}^{6} x\mathbb{P}[X = x] = 21 \times 1/6 = 7/2$$

$$\mathbb{E}\left[X^2\right] = \sum_{x \in \Omega} x^2\mathbb{P}[X = x] = \sum_{x=1}^{6} x^2\mathbb{P}[X = x] = 91 \times 1/6 = 15\frac{1}{6}$$

◁

Example 4.8 (A Roulette Wheel Again). We could consider the expectation of some random variables associated with a Roulette wheel.

– Recall the roulette random variable introduced earlier.

$$\begin{aligned}\mathbb{E}[X] =& \sum_{x_i} x_i \times \mathbb{P}[X = x_i] \\ =& -1 \times \mathbb{P}[X = -1] + 1 \times \mathbb{P}[X = 1] + 39 \times \mathbb{P}[X = 39] \\ =& -1 \times n(r)/n + 1 \times n(b)/n + 39 \times n(g)/n \\ =& (n(b) - n(r) + 39 \times n(g))/n = 39/37 = 1 + 2/37.\end{aligned}$$

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

$$\begin{aligned}\mathbb{E}\left[X^2\right] =& \mathbb{E}[f(X)] \\ =& 1^2 \times \mathbb{P}[X = -1] + 1^2 \times \mathbb{P}[X = 1] + 39^2 \times \mathbb{P}[X = 39] \\ =& (n(r) + \times n(b) + 1521 \times n(g))/n \\ =& (18 + 18 + 1521)/37 = 42.08\end{aligned}$$

◁

5. Making Decisions

5.1 Decision Problems

From the outset our intention has been to make decisions in the presence of uncertainty. Having looked in some detail at the way in which we can deal with uncertainty — specifying beliefs, updating those beliefs in the presence of new information and interpreting statements about those beliefs — we are now in a position to start answering the question: what should we advise our client to do.

We will begin by formulating a simple generic decision problem and the consider some *decision rules*: ways of choosing an action in light of the uncertain information which is available.

Decision Ingredients. The basic components of a decision analysis are:

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

Our control over the problem is limited to selecting an element of D. This is a rather limited form of control as the particular outcome, $x \in \mathcal{X}$ which occurs will interact with our action and the relationship between the act and outcome may be complicated.

Definition 5.1 (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.*

Notice that L is a function of two variables which takes both a *decision* and an *outcome* as arguments and returns the (real-valued) loss resulting from that particular combination. The loss depends not only on the final outcome but on what decision was made: in many cases there will be some sort of cost associated with each potential decision and balancing this cost with the intrinsic values of the different outcomes is part of the problem. The relation between decision, outcome and resulting loss may be an intricate one.

An algorithm for choosing a particular action $d \in D$ is a *decision rule*.

It's convenient to begin with a simple example to provide a definite instance of each of these ingredients:

Example 5.1 (Insurance). Should we buy insurance? You must decide whether to pay c to insure your possessions of value v against theft for the next year:

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

That is the decision space contains two elements, each corresponding to one possible action.

In a simple,stylised analysis, 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}

these are the elements of the outcome space, $\mathcal{X}$.

Considering the cost of buying insurance and the losses associated with uninsured thefts, we arrive at the following tabulation of our loss function.

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

◁

We will return to this example and some simple variants later to answer the question posed (and to illustrate how a rational decision maker might arrive at a conclusion in this and similar problems).

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. For example, that in the insurance example our choice to buy or not buy insurance will not influence whether or not we then suffer a burglary. This allows us to avoid the added complications of *moral hazard* familiar to the insurance industry: imagine, if you will, purchasing life insurance on the life of an individual whom you do not much like. The insurer would have assumed that this does not alter the probability of that individual dying. In some instances that may not be the case (imagine that the amount for which they are insured is large and that they are an individual whom you severely dislike). In order to avoid this problem, there are significant limitations upon who can take out life insurance on any particular individual!

It is convenient to work with a probability space $\Omega = \mathcal{X}$ and the algebra generated by the collection of single elements of $\mathcal{X}$. It then suffices to specify a probability mass function for the elements of $\mathcal{X}$. One way in which we might choose to address uncertainty is to work with expectations so that we make decisions which in some sense will be good ones on average.

Example 5.2 (Insurance Continued). Returning to our insurance example. Not being intimate with the workings of the local criminal fraternity, we need some way to assign probabilities to the events which might occur. It may be possible to elicit these from a suitably knowledgeable individual but this might not be something that we wish to do. Assuming that burglaries are largely independent will allow us to obtain a crude estimate, albeit one based upon an assumption which is violated to a significant extent and neglects lots of other information of which we could make use.

There are 25 million occupied homes in the UK (2001 Census). Approximately 280,000 domestic burglaries are carried out each year; approximately 1.07 million acts of "theft from the house" were carried out (2007/08 Crime Report).

We might naïvely assess our pmf using the observed frequencies of occurrence and an assumption that all houses are equally likely to be burgled and furthermore that houses are burgled independently:

$$p(x_1) = \frac{25 - 1.07 - 0.28}{25} = 0.946$$

$$p(x_2) = \frac{1.07}{25} = 0.043$$

$$p(x_3) = \frac{0.28}{25} = 0.011$$

◁

The EMV Decision Rule. We wish to experience as small a loss as we can, but we have incomplete information available to us: we know only what decisions are available to us, the probability distribution associated with the possible outcomes and the loss of particular (decision, outcome)-pairs. Unfortunately, without knowing the outcome in advance, which we don't, we can't choose the decision which minimises our loss in any particular decision. What we need is some way to eliminate this explicit dependence on outcome from our decision-making process.

If we calculate the expected loss for each decision, by taking its expectation over the possible outcomes, 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*. One justification: If you make a lot of decisions in this way the you might expect an averaging effect... but a more fundamental reason will be given later.

Example 5.3 (Still insurance). Back to the insurance question. We had formulated a decision problem.

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 summarising our beliefs about the likelihood of burglary or theft over the period of the insurance:

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

If we consider the expected loss under our pmf for each of our possible actions, we obtain expected losses:

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

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

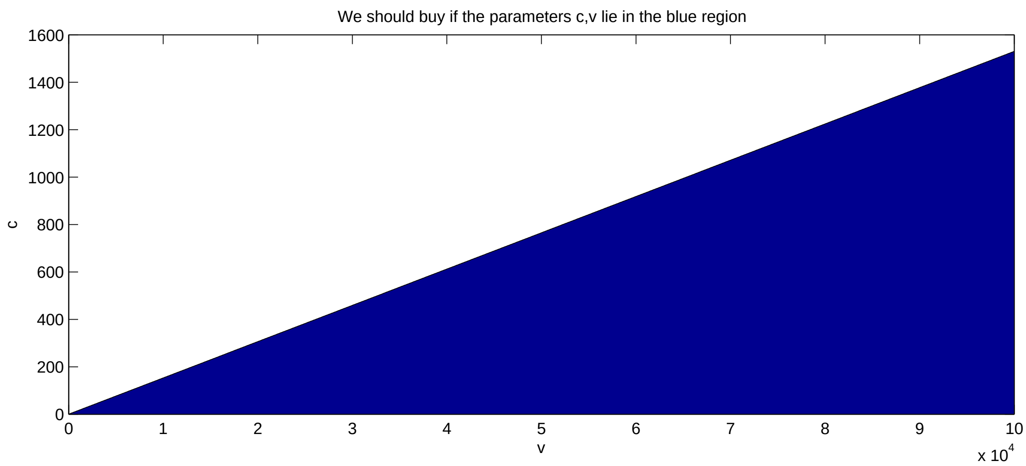

◁

Optimistic EMV. All this talk of loss is a little pessimistic and isn't always the most natural way to formulate a problem. There may be cases in which a profit is expected under some (or even all) of the possible decisions. In this case it makes more sense to consider profits than losses.

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

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

This is, of course, simply a semantic change. Although some signs have been changed and the objects being used given different names the underlying procedure is identical. Whether loss or reward is the more natural way to describe the outcomes of a particular problem depends upon the particular characteristics of the problem (and perhaps the people to whom you're planning to show the decision analysis to). Choosing the less natural form will, at least mathematically, just lead to an analysis containing an unnecessarily large number of negative numbers; the outcome and calculations are essentially unchanged.

5.2 Decision Trees

Although we have a formal rule for making decisions, in large problems (with perhaps dozens of decisions and hundreds of different random outcomes) it becomes rather difficult to keep track of everything that needs to be calculated and to make the right calculations. In order to deal with problems such as this, it is necessary to make use of the right notation.

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 possible set of decisions. It must allow us to calculate the EMV decision for a problem in a systematic way and it should be sufficiently flexible that we are able to adapt it to other decision rules which we may wish to employ.

Graphical Representation: Decision Trees. The *decision tree* is a useful graphical representation is a way to describe a complete decision problem in a single diagram. We shall see that it also provides a simple mechanism for determining EMV decisions for a problem and for visualising and interpreting why that decision is optimal in an EMV sense and where other decisions are inferior to it.

There is a simple procedure for constructing a decision tree for any problem.

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 (that is, placing decisions and random events in the order in which they must be made/revealed in the real decision problem):
   a) Indicate decisions with a □.
   b) Draw forks from decision *nodes* labelled with the particular decisions which can be made.
   c) Indicate sets of random outcomes with a ◯.

d) Draw edges from random event *nodes* labelled with their (conditional upon everything to the left of that point in the tree) probabilities.
e) Continue iteratively until all decisions and random variables are shown.
f) At the right hand end of each path indicate the *loss/reward* associated with the sequence of decisions/random outcomes which must occur to arrive at the end of that particular path. Of course, it's important to stick the *loss* or *reward* in any particular problem: which one is used doesn't matter, provided that you are consistent.

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

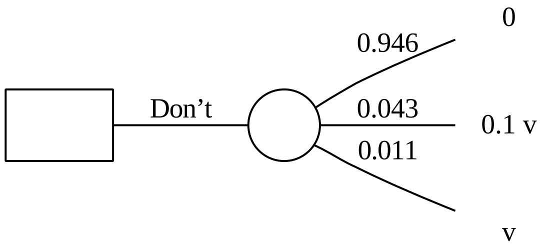

Constructing the fork for the other possible decision and combining them produces a complete decision tree. We can then attach probabilities and expected values to the tree (a general approach for doing this will be described below):

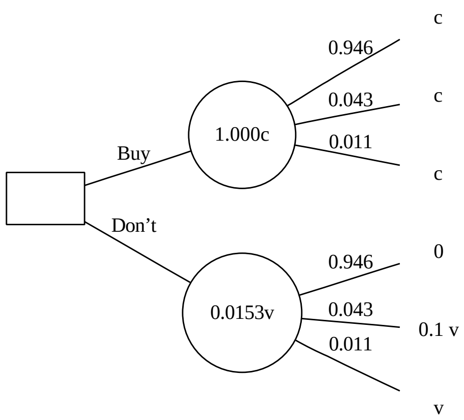

We've worked backwards from the RHS filling in the expected losses associated with each decision.

But we didn't need to make things that complicated: there is only one outcome if we buy insurance (we only need to include forks from chance nodes if they influence the loss and in some instances it may be possible to merge a number of random variables to produce a single chance node with forks for all relevant differing outcomes and their associated probabilities):

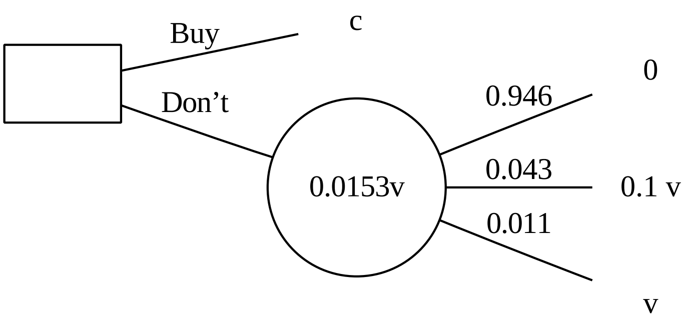

In more complex examples, we should label the random events (say N for no robbery, T for small theft and B for burglary...

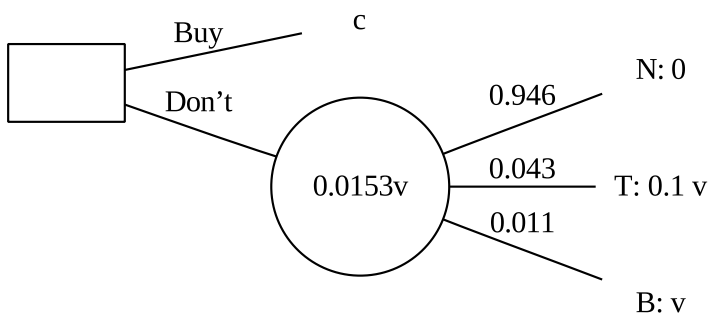

Calculation and Decision Trees. In the insurance example, we calculated the expected loss of each of the two possible decisions and it would be easy to decide how to act on the basis of this calculation. In more complicated problems we would need to be systematic about these calculations in order to keep track of everything that's going on and to minimise the number of redundant calculations. Fortunately, there is a convenient algorithmic approach to making EMV decisions using a decisions tree and, as a side effect, we gain a significant amount of information about the structure of the decision problem along the way.

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

1. Starting at the right hand end of the graph, trace each path back to ◯ nodes:
   a) Fill in the rightmost ◯ nodes with the (conditional on all earlier events – i.e. ones to the left) expected losses (the probabilities and losses are indicated at the edges and ends of the edges).
   b) For each decision node which now has values at the end of each branch, find the branch with the smallest loss (or largest reward).
   c) Eliminate all of the others: we wouldn't choose to make those decisions so they can be eliminated from consideration.
   d) This produces a reduced decision tree.
   e) If there are several levels of decision to be made it may be necessary to iterate, filling in the next level of chance nodes now that we've decided how we would act if we arrived at points immediately to their right.
2. When left with one path, this is the EMV decision.

Do Not Laugh at Notations. 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. It is also a compact and efficient notation which is easy to interpret and which can be used to justify any action guidance which it produces to a client who may not be particularly mathematically inclined.

Good mathematicians are aware of the power of good notation and appropriate representations:

> "We could, of course, use any notation we want; do not laugh at notations; invent them, they are powerful. In fact, mathematics is, to a large extent, invention of better notations."

Richard P. Feynman

"By relieving the brain of all unnecessary work, a good notation sets it free to concentrate on more advanced problems"

Alfred North Whitehead

5.3 Decision Trees — Example

Looking at a slightly more involved decision problem starts to show why decision trees are useful and the extent to which they automate the decision making process.

Consider this case as explained to you by a client:

- 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 but will provide no reward.

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 or drilling in both but there would be no fundamental difficulty in including these options).

Your Knowledge. We begin by eliciting the following information from the oil company. After some time we are left with the following, apparently accurate representation:

- 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?. From this collection of events, we need to encode things in a formal way that will allow us to attack the decision problem.

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 we have may be written as:

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

Some Calculation is Needed. We know $\mathbb{P}(a|A)$ and related quantities. This is the probability that an investigation will indicate oil is present if there is; we really need to know the probability that an expert will indicate that there is oil present in these fields if they are asked and, further, the probability that oil is present in a field given that investigation indicates that there is.

The first of these things may be calculated by the partition theorem:

$$\begin{aligned}\mathbb{P}(a) =& \mathbb{P}(a|A)\mathbb{P}(A) + \mathbb{P}(a|A^c)\mathbb{P}(A^c) \\ =& 0.8 \times 0.4 + 0.2 \times 0.6 \\ =& 0.32 + 0.12 = 0.44\end{aligned}$$

$$\begin{aligned}\mathbb{P}(b) =& \mathbb{P}(b|B)\mathbb{P}(B) + \mathbb{P}(b|B^c)\mathbb{P}(B^c) \\ =& 0.8 \times 0.2 + 0.2 \times 0.8 \\ =& 0.16 + 0.16 = 0.32\end{aligned}$$

and the second by Bayes' rule:

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

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

Once again, we must be careful: investigation actually provides weaker evidence than may at first appear to be the case. The probability that experts recommend drilling in either field is greater than the probability that oil will be found there; in the case of the second field, even if we are advised to drill by an expert the probability of finding oil is still only 0.5.

This is a reasonably detailed decision problem and it probably isn't immediately obvious what the most sensible strategy would be. The process of constructing and then solving a decision tree breaks the problem down into manageable steps which can be completed in a systematic way. We begin by constructing the tree without probabilities:

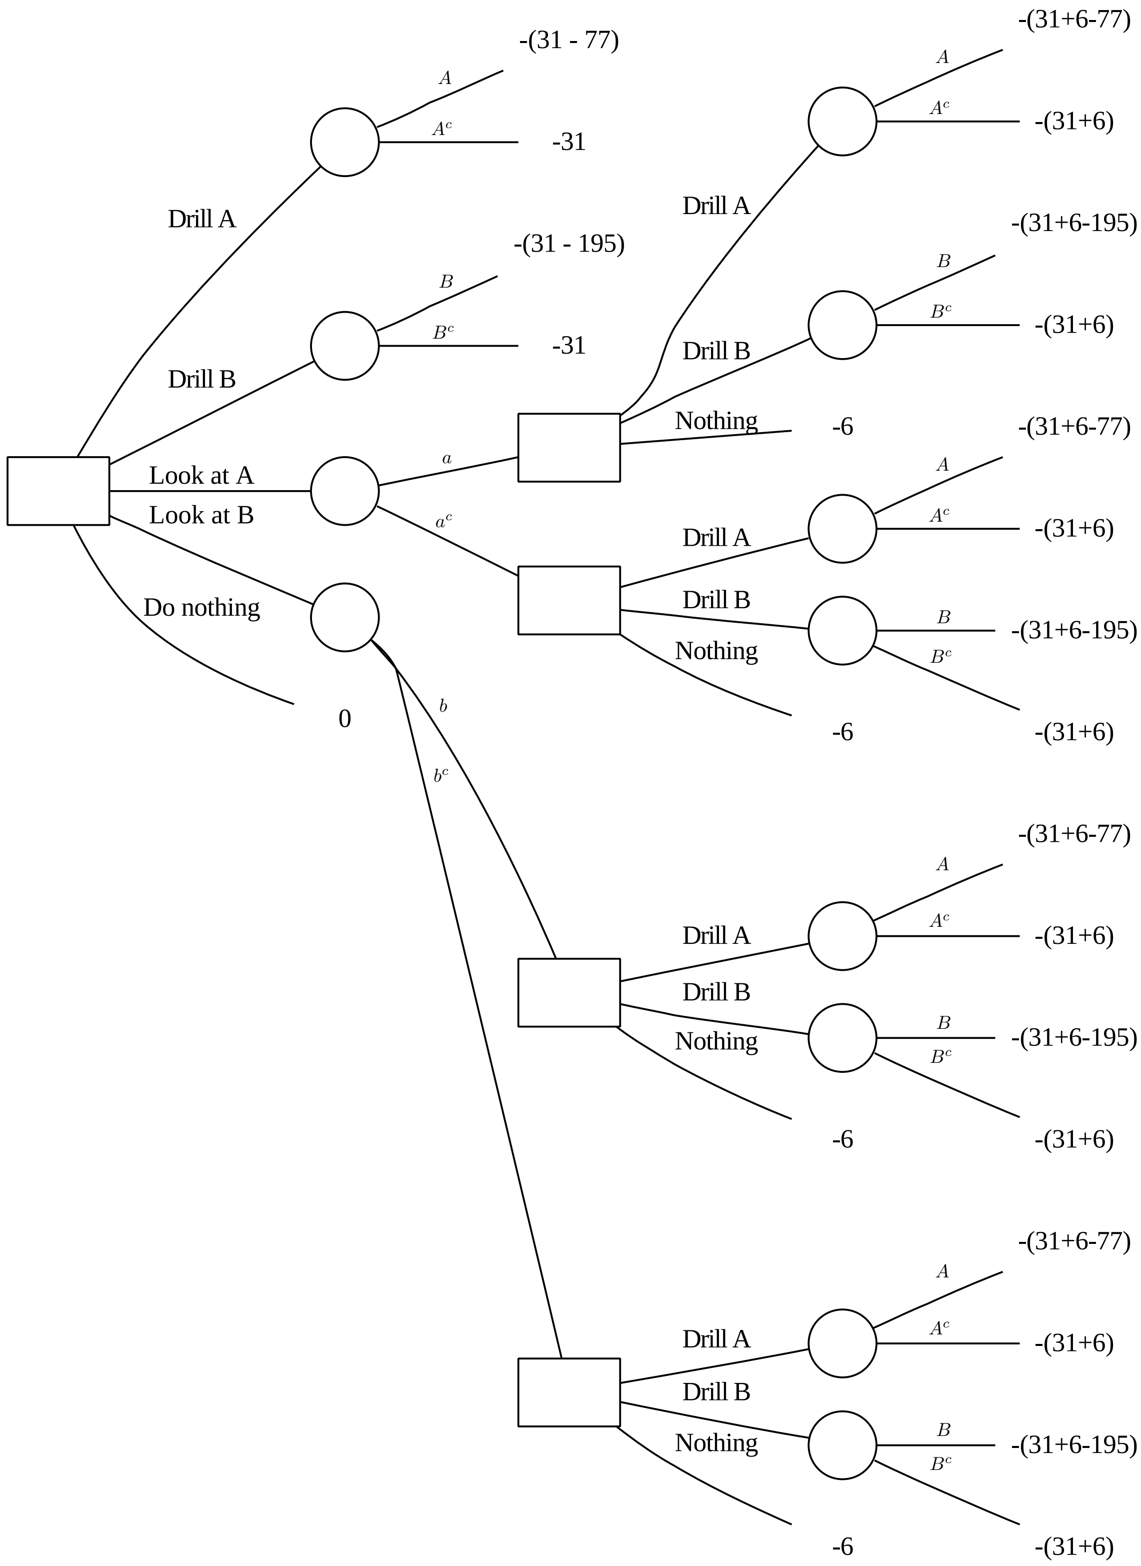
-(31 - 77)
A
A^c
-31
Drill A
-(31 - 195)
B
B^c
-31
Drill B
Look at A
Look at B
Do nothing
0
a
a^c
b
b^c
-(31+6-77)
A
A^c
-(31+6)
Drill A
-(31+6-195)
B
B^c
-(31+6)
Drill B
Nothing
-6
-(31+6-77)
A
A^c
-(31+6)
Drill A
Drill B
B
-(31+6-195)
B^c
Nothing
-6
-(31+6)
-(31+6-77)
A
A^c
-(31+6)
Drill A
Drill B
B
-(31+6-195)
B^c
Nothing
-6
-(31+6)
-(31+6-77)
A
A^c
-(31+6)
Drill A
Drill B
B
-(31+6-195)
B^c
Nothing
-6
-(31+6)

Then work out what each probability should be:

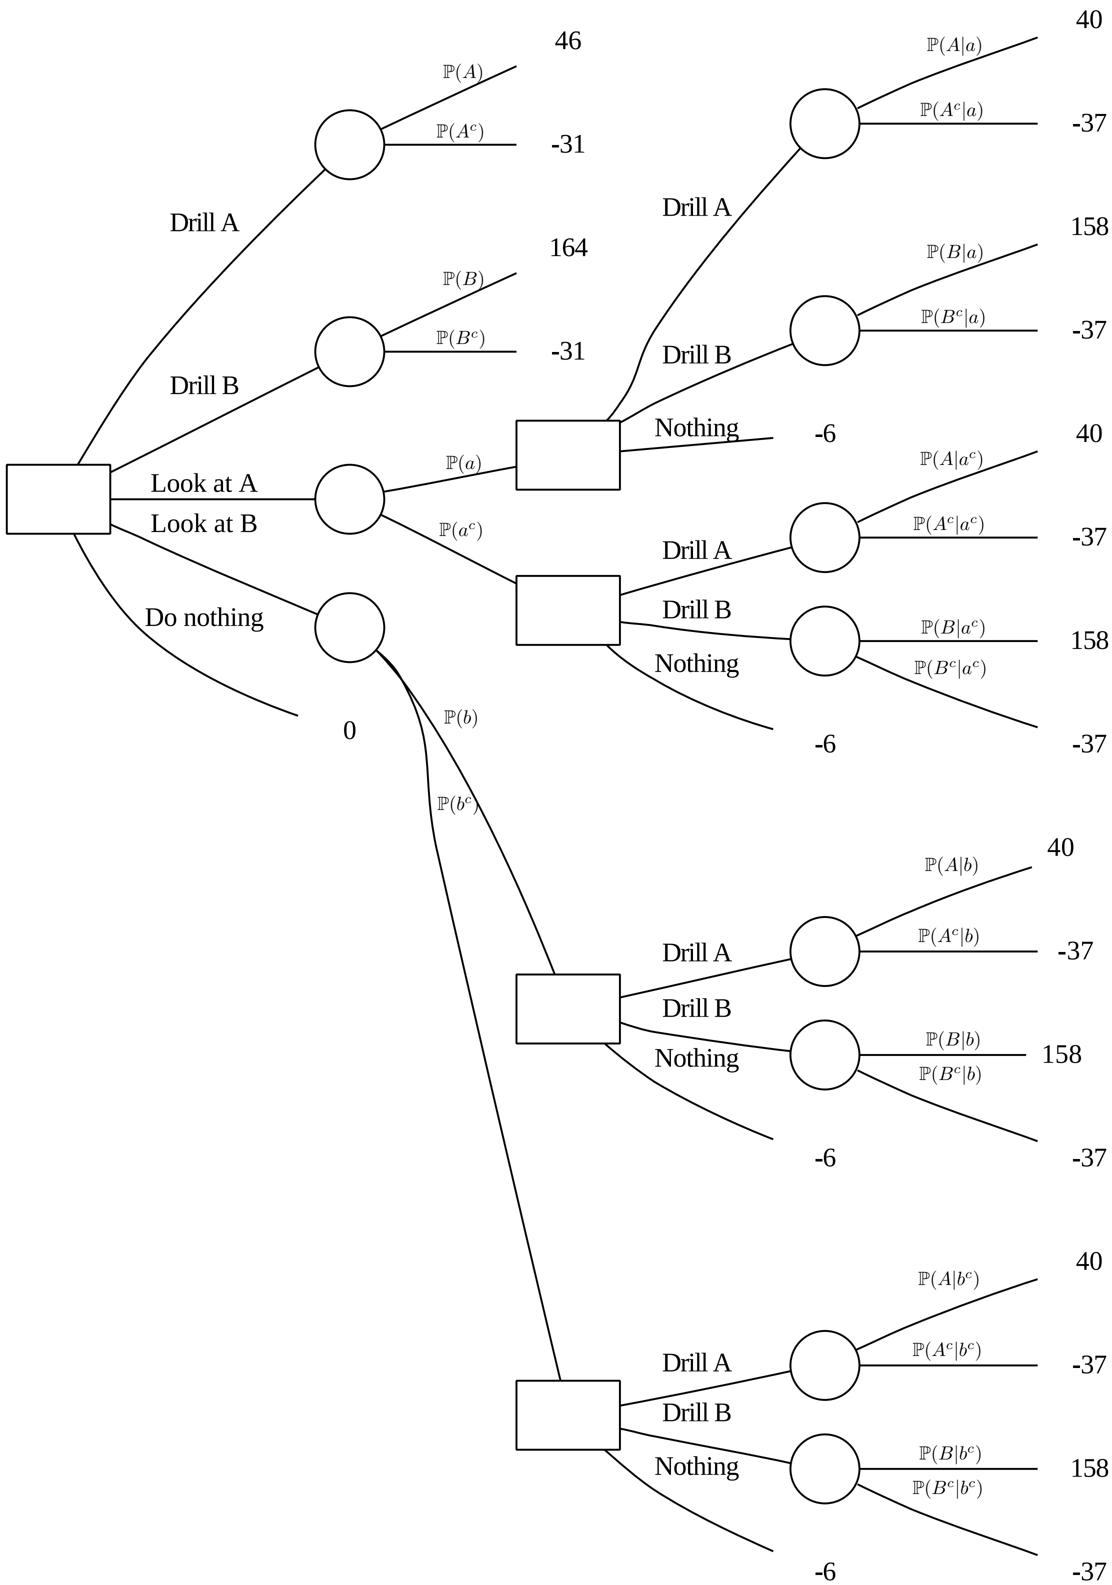

Then work out what each probability should be numerically:

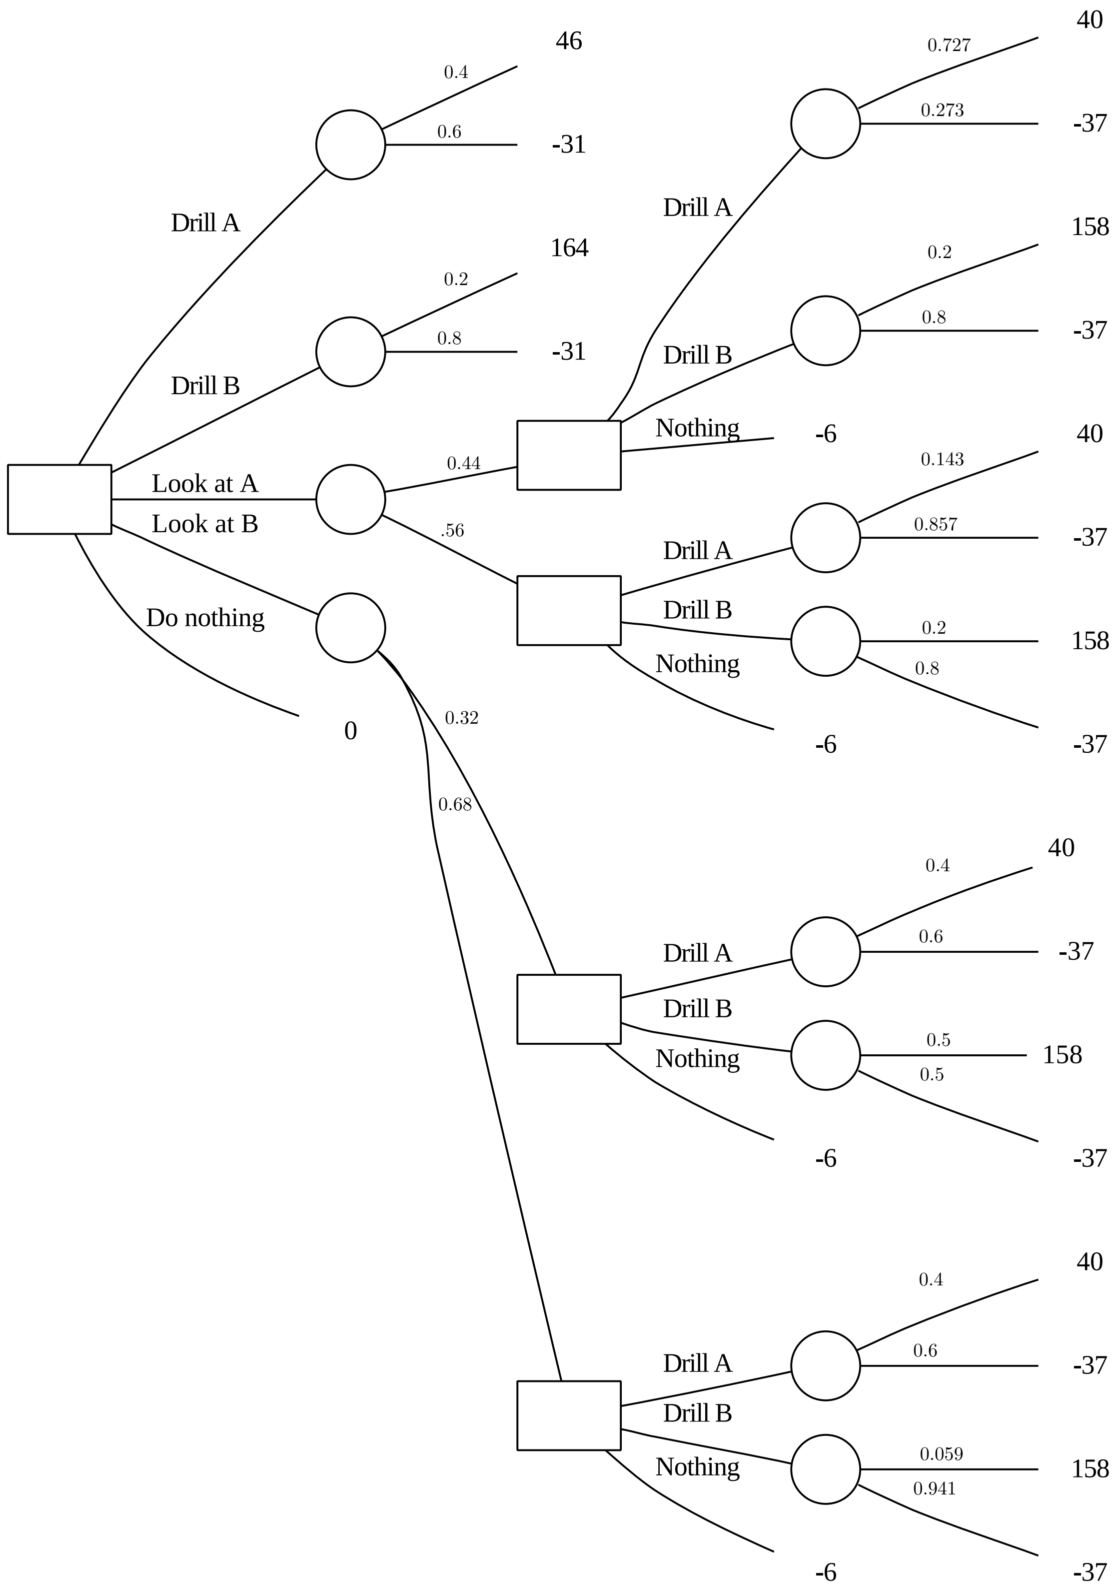

Then starting at the RHS calculate expectations and make optimal decisions to determine the solution.

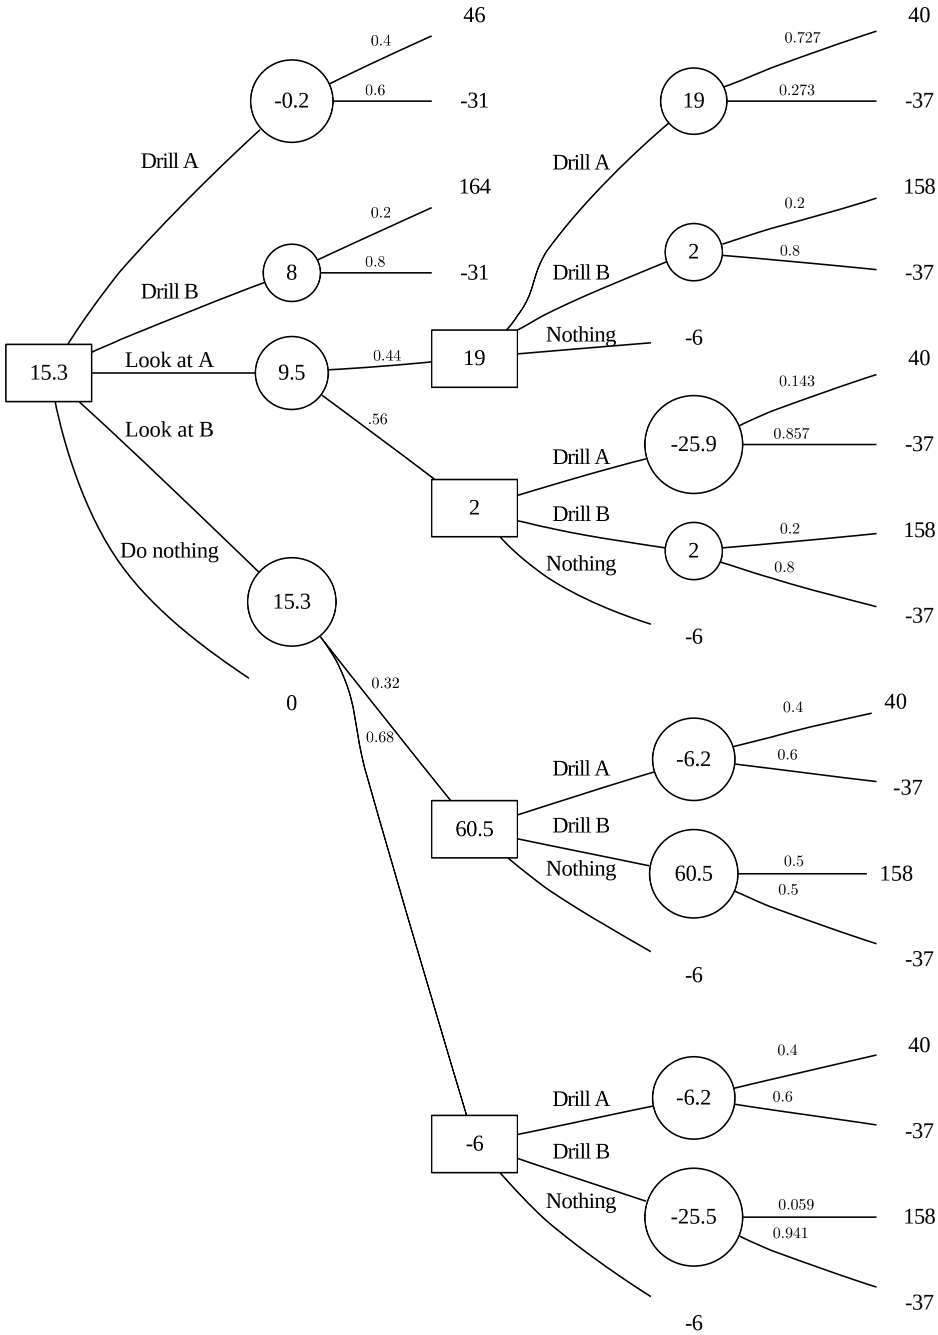

In this case we should investigate B; if it is suggested that there is oil there then we should drill there, otherwise we should do nothing. If you'd like to see some more examples then have a look at `http://people.brunel.ac.uk/~mastjjb/jeb/or/decmore.html` which provides numerous worked examples from old exam papers used at Imperial College.

Perfect and Imperfect Information.

Definition 5.2 (Expected Value of Perfect Information (EVPI)). *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.*

Definition 5.3 (Expected Value of Imperfect Information (EVII)). *The difference in the expected value of a decision problem in which decisions are made with access to an imperfect source of information and one in which no additional knowledge is available.*

The result of carrying out a preliminary trial is an example of *imperfect information.* Investigating B is part of our EMV strategy, so the information we obtain is clearly valuable: it worth more than the £6M cost of performing the trial.

For the sake of comparison, forget for the moment about the possibility of doing a preliminary trial. The EMV strategy is then to drill at B with an expected reward of £8M.

Suppose that we can only do a trial at A, but that the trial does not cost us anything. Then our EMV strategy is to look at A, and then drill at either A or B. The expected reward is £15.5 M ($£9.5M + £6M$ as we are ignoring the cost of the trial). The EVII associated with the trial at A is the increase in the expected reward: £15.5M - £8M = £7.5M.

Alternatively, suppose that we do a trial at B and that the trial does not cost us anything. Then our EMV strategy is to look at B, and then drill at either A or B. The expected reward is £21.3 M ($£15.3.5M$ + $£6M$). The EVII associated with the trial at B is the increase in the expected reward: £21.3M - £8M = £13.3M.

Thus looking at either site is worth more than the £6M cost. In the example, we were limited to only looking at one site so the EMV strategy found the better value-for-money source of imperfect information for us. Note that the value of imperfect information is not generally additive. Suppose we could look (for free) at both A and B, but that we are still limited to drill in at most one location. The increase in expected reward would be less than £7.5M + £13.3M. This is because if we find evidence of oil at both A and B, we cannot take advantage of both pieces of information.

A natural question to ask is how much more would we be willing to pay for *perfect information*, that is to find out for certain which of the four cases $\{A \cap B, A \cap B^c, A^c \cap B, A^c \cap B^c\}$ we are in. In the table below our EMV strategy is shown in bold:

$R(d, x)$	$A \cap B$	$A \cap B^c$	$A^c \cap B$	$A^c \cap B^c$
Drill A	46	**46**	-31	-31
Drill B	**164**	-31	**164**	-31
Do Nothing	0	0	0	**0**
$\mathbb{P}$	0.08	0.32	0.12	0.48

The expected reward for the EMV strategy is $(0.08 + 0.12) \times £164M + 0.32 \times £46M + 0.48 \times £0M = £47.52M$. The EVPI is the increase in expected reward compared to having no extra information: £47.52M-£8M=£39.52M. It might seem odd that there is still an expectation involved when we are considering the problem with perfect information—the EVPI measure the value of the information *before* we receive the information, so the situation is still uncertain.

The large difference between the EVPI and the EVIIs suggests that it might be worth putting effort into improving the oil detection procedure.

6. Preferences

6.1 The Trouble With Money

We will consider the following example throughout much of this chapter as we investigate whether EMV decisions are necessarily always appropriate and look at some possible alternatives.

Example 6.1 (The Farmer's Trilemma). A farmer must decide which crop to plant; whichever crop they choose to plant, the resulting 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? In the previous chapter we considered the EMV decision rule and we could apply that here. However, it may not always be appropriate as this considers only the expected value and so in some sense it fails to take into account the level of risk associated with the decision. People vary in their attitudes to risk: most people are risk-averse and would prefer to avoid it whilst others are risk-seeking and actively enjoy risk. What other approaches could we employ to advise our farmer? ◁

Maximin Decisions. One farmer believes that the weather will do whatever makes things worst, whatever decision he makes. He's either pessimistic or paranoid although as Jan Phillipe Reemstma once wrote, "Sometimes paranoia is just what psychologists refer to as a well-adjusted perceptual attitude." He logically (given his beliefs) seeks to 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 and guarantees himself a reward of at least 2; he is happy with this knowing that things could have been worse if he had opted to plant either of the other crops.

This is known as a *maximin* decision: it maximises the minimum reward. In other settings, you may see this referred to as a *minimax* approach (as it minimises the maximum loss).

Maximax Decisions. A second farmer believes that the weather will do whatever makes things best, whatever decision he makes. He's either optimistic or feeling lucky. He logically (given his beliefs) wishes to 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. If things go well he will receive a reward of 11; he is happy with this knowing that he wouldn't have the opportunity to do so well if he had chosen to plant either of the other crops.

This is known as a *maximax* decision: it maximises the maximum reward.

The Hazards of Extremism. Maximin and maximax solutions may sometimes be acceptable they provide an encoding of a specific approach to decision making which may, in simple scenarios, correspond to exactly what is wanted.

However, one problem which rapidly arises in realistic problems is that 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 upon which you base your decision. But, in decision problems, you work with an idealisation in which you haven't really considered *every* possible outcome. This seems rather inconsistent. At the very least is seems that it will be necessary to worry about precisely which outcomes are considered possible and as soon as we start excluding extremely unlikely severe events we weaken the robustness of the approach.

By way of an example, consider our pair of farmers again. If the pessimistic farmer heard there was a tiny chance of a hurricane devastating his crops then he might choose to base all of his decisions upon maximising the amount of money he would make in the event of a hurricane. If he then heard that there may also be a volcanic eruption from an as yet undiscovered volcano in the vicinity of his farm then this may further influence his decision. The strategy seems to become less and less sensible as further options are included. Similarly, the optimistic farmer might hear than there's a tiny chance of a really fantastic year and decide that it's time to start growing lemons in the north of Scotland.

As such, we'd clearly like a way in which we could deal with these two different attitudes to risk in a way which will allow us to keep our farmers happy without taking quite so extreme a position and without attaching undue significance to extremely unlikely events: in reality there is a limit to how extreme an event we can realistically be prepared for without dramatically impairing our ability to function in more normal circumstances.

Paradoxes in St. Petersburg. There is a rather well-known illustration of the general problem with dealing exclusively with expected monetary value. How much is the following bet worth?

Example 6.2 (St. Petersburg Lottery). You are allowed, for a price, to enter the following lottery:

- The prize is initially £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.

What price would be a fair one? ◁

If we consider the expected value approach things start to look rather strange. The probability of obtaining $n-1$ heads before the first tail is 2^{-n} (it's a geometric random variable). If there are $n-1$ heads then the reward is doubled, from one, $n-1$ times, and the prize is $1 \times 2^{n-1} = 2^{n-1}$. Consequently, the expected value of the decision to play this game is, allowing $R(d,n)$ to denote the reward obtained if decision d is made and the first tail is seen at toss number n and $p(n)$ denotes the probability that the first tail is seen at toss number n:

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

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. Equivalently, a rational being employing an EMV strategy should be prepared to pay *any* finite amount in order to be allowed to play the game once. Would you rather play this game or have £1, 000, 000?

A surprising amount of effort has gone into the development of explanations of this apparent paradox and in particular into finding ways to prevent this problem from arising in other situations. We will investigate the most widespread such approach in the next section.

6.2 Utility

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. This may initially seem like a reasonable proposition. The value of receiving £2 with a probability of 0.5 would seem to be roughly the same as being given £1 in many circumstances. However, would you rather have £10^8 with certainty or a probability of 10^{-9} of having £10^{17}? I would certainly prefer the first, although it seems likely that this is one decision which I will never be called upon to make.

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: it can be rather more complicated than the product of the two: indeed, in general, it will be a nonlinear function which will dramatically influence our decision-making process.

In order to formalise things it will be useful to introduce some notation. Let A, B and C be simple random rewards: particular rewards (say r_A, r_B and r_C) with some probability (say p_A, p_B and p_C) or nothing otherwise — each corresponding to a particular simple bet of the sort which were used when attaching a behavioural interpretation to subjective probabilities.

We wish to discuss which of various outcomes are preferable to a client, and it would be useful to have a compact representation of preference statements:

- Write $A \succ B$ (or $B \prec A$) if A is strictly preferred to B.
- Write $A \sim B$ if A and B are equally preferable.
- Write $A \succeq B$ (or $B \preceq A$) if A is at least as good as B.

It will also be useful to have some notation for describing a compound bet which with some probability yields one outcome but otherwise yields another. For some $t \in (0, 1)$, let $tA + (1 - t)B$ denote random rewards A occurring with probability t and B with probability $1 - t$. For the sake of definiteness, $tA + (1 - t)B$, means with probability t the reward associated with A is received with whatever probability that reward would occur in bet A; whilst with probability $1 - t$ whatever reward is associated with B is received with the appropriate probability. Let p_A and p_B denote the probability of receiving a reward under A and B and let r_A and r_B denote the associated rewards. In this setting, $tA + (1 - t)B$ provides a reward of r_A with probability tp_A, reward r_B with probability $(1 - t)p_B$ and nothing otherwise.

Using this notation we can make a number of statements about the properties which an individual's preferences should have if they are to be considered rational.

Definition 6.1 (Rational Preferences). *If a collection of preferences obey the following four axioms:*

1. *Completeness: For any A, B one of the following holds:*

$$A \succ B \qquad\qquad A \sim B \qquad\qquad A \prec B$$

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.

It should be noted that there are various criticisms of these requirements. They may initially seem rather mild, but there are circumstances in which one might question whether they are realistic. In particular, the completeness axiom as the first of these statements is generally known is rather controversial: are we really able to state a preference between any two possible outcomes? That said, these requirements are a rather sensible starting point and for the remainder of the present course we will assume that we are happy to abide by them.

If the axioms of rational preference are satisfied, then a number of useful things prove to be true. In particular, 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. More specifically, any decision rule compatible with the preferences stated is equivalent to maximising the expectation of a suitable utility function. That is, anyone making decisions in line with a set of rational preferences as defined previously makes decisions as if they seek to maximise the expected value of some utility function.

It's outside the scope of this course to prove this but it will become apparent that it is reasonable from the next pages. Some of the recommended references include further details if you are interested in this. Notice, that the EMV decision rule is a particular case of the approach of maximising expected utility in which the utility function is linear and it is this which really justifies the use of EMV decision rules in such settings as it is possible to justify a linear utility.

Eliciting Utilities. If preferences are to be represented by utilities, we must be able to determine utility functions. When we considered the behavioural interpretation of probabilities, we assumed that the rewards, M, being considered were small. The reason for this was that it makes it reasonable to assume that a rational person would behave according to EMV-type principles. When large amounts are at stake it can dramatically influence the value people attach to things: for example, many people are prepared to buy lottery tickets which give them a very small chance of winning a large amount of money; how many of them would be prepared to accept the cost of a lottery ticket in exchange for guaranteeing to pay out millions of pounds if it won?

We can use an individual's perceived values of different bets as a method for eliciting their utility function. If t is the (monetary) value of the best possible outcome for the class of problems which are being considered, and s is that of the worst possible outcome, then we could consider bets which pay one or other of these amounts with various probabilities as shown:

How much is the following bet worth to you? In other words, what m would you accept not to benefit from the bet shown (if the alternative is benefitting from this bet at no expense)? This is a function of α, and so we could write $m = f(\alpha)$.

This bet:

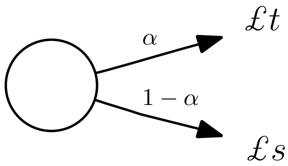

has certain monetary equivalent (CME) value $m = f(\alpha)$. That is, our individual considers receiving m with certainty to have the same value as receiving t with probability α and s with probability $1-\alpha$. There are three qualitatively different possibilities:

- $m < \alpha t + (1-\alpha)s$: the CME of the bet is less than its expected value. Someone who takes this position is referred to as *risk averse*: they value certain rewards more highly than fair but risky bets with the same expected reward.
- $m = \alpha t + (1-\alpha)s$: the CME of the bet is equal to its expected value. This is a *risk neutral* view: an individual with such values does not take the risk associated with a bet, only its expected value.
- $m > \alpha t + (1-\alpha)s$: the CME of the bet is more than its expected value. Such a view is referred to ask *risk seeking.*

Most, but not all, individuals are risk-averse when dealing with sums of money which are significant to them. Typically, when dealing with small enough sums the assumption of risk-neutrality is a sensible one. However, we will see some examples in which a risk-seeking position can be justified. All of these positions are sensible in the right circumstances and it is convenient to have a single formalism in which al of them can be readily encoded.

A utility function can be defined so that the utility of m is $U(m) = f^{-1}(m)$. This will be consistent with the preferences between certain amounts of money and bets in which the two most extreme events occur and any other preferences which are rationally implied by these. In fact, this class is rich enough that we need only determine $f(\alpha)$ for $\alpha \in [0,1]$ to adequately specify a utility function for a particular individual confronted with a problem whose possible outcomes lie in $[s,t]$.

If we define utility in this way, we obtain a method for mapping the value of an outcome, which must lie between s and t, to a number between zero and one. Under the axioms of rational preference, this utility function is sufficiently powerful that it will allow us to compare any set of possible outcomes and to ascribe a utility value to each of them in such a way that any outcome with a higher utility than another is strictly preferred to that other outcome and any pair of outcomes with common utility are equally valuable.

A Family of Utilities. It is often convenient to have available a parameterised family of utility functions which can capture a variety of qualitatively-different collections of preferences.

Figure 6.1 shows a number of representatives of one such family of utility functions:

$$U_\alpha(x) = x^\alpha \qquad \alpha > 0$$

This collection has a single parameter, α, which determines what sort of behaviour the utility function will lead to. If $\alpha = 1$ it corresponds to linear utility and anyone with such a utility function should behave in a *risk-neutral* manner, caring only about expected value and acting in accordance with the principles of EMV decision making. The $\alpha < 1$ utility functions are *risk-averse* and individuals would prefer to accept a smaller amount (than the EMV) with certainty than to indulge in a risky bet. In the unusual $\alpha > 1$

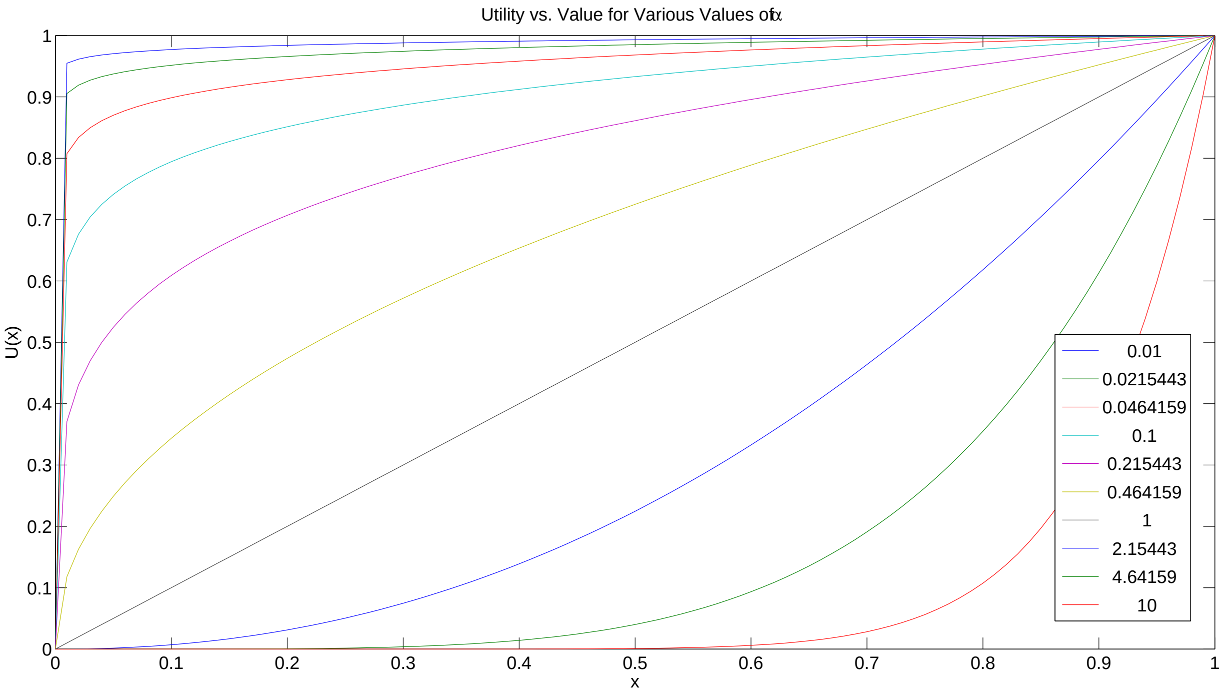

Fig. 6.1. A family of utility functions, $U_\alpha(x) = x^\alpha$.

cases, the utility function describes *risk-seeking* behaviour in which the value attached to risky bets is actually more than the EMV.

Example 6.3 (The Utility of Insurance). Remember the insurance example...

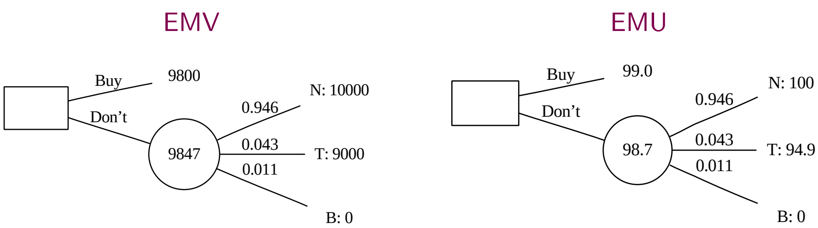

The first figure shows the EMV position: you'd prefer not to buy insurance with this reasoning. For the same reason, the insurer would prefer you to buy it.

The second shows the EMU position with $U(x) = \sqrt{x}$ — you prefer to insure. This utility function increase more slowly than linearly and so it encodes a risk-averse collection of preferences.

EMV makes sense for the insurer; EMU for you. Both parties acting rationally and using the same beliefs arrive at different conclusions because of differing preferences. The large company has a largely neutral attitude to risk on the scale of a single insurance pay-out whilst you, as most individuals do, are risk-averse when considering possible losses on a scale corresponding to all of your accumulated worldly goods.

This may seem like a rather convoluted route by which we have arrived at an argument that could be considered rather obvious, but is it not impressive that using a rather simple mathematical framework we have arrived at a precise and formal encoding of a rational argument which actually does describe the way in which be act? ◁

Example 6.4 (The Value of Money). We don't (always) really think that a chance p of winning an amount r has a value of pr...

Consider a lottery which pays a reward £X where X is a random number distributed uniformly over $[0, 4]$. This is an example in which the reward is a *continuous* random variable and it is included here just to show that no particular difficulties arise when we depart from the discrete cases considered elsewhere through the course. We won't be making use of continuous random variables elsewhere in this course.

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:

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

The fair price, x_f is such that

$$U_\alpha(x_f) = \mathbb{E}[U_\alpha(X)]$$

i.e., the fair price is the solution of the equation:

$$\begin{aligned} U_\alpha(x_f) = x_f^\alpha &= \frac{4^\alpha}{\alpha + 1} \\ x_f &= \frac{4}{(\alpha + 1)^{1/\alpha}} \end{aligned}$$

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. ◁

Some Aspects of Utility. The family of utility functions described previously were each within one of three categories: risk averse, risk neutral or risk seeking. In general, a risk averse utility function is one which increases more slowly than linearly (and so it has negative second derivative); a risk neutral utility is linear and a risk-seeking one increases more rapidly than linearly (with positive second derivative). Of course, more complicated utility functions may not have a second derivative with the same sign throughout the range of values of interest (or, indeed, even have a second derivative which is defined everywhere in this range).

A common way of measuring the degree of risk aversion is to take, noting that a utility function must logically be increasing and hence have positive first derivative, the function:

$$-\frac{\mathrm{d}^2 U}{\mathrm{d}x^2} \Big/ \frac{\mathrm{d}U}{\mathrm{d}x}$$

is often used as a measure of the degree of risk aversion. Notice that by taking the ratio of derivatives in this way the absolute scaling of the utility function has been eliminated. In general, of course, this value will be a function of the particular value of utility under consideration.

When we discussed the elicitation of utilities by elicitation of the values of particular bets, it was noted that the utility function would only take values between zero and one. In fact, in order to entirely avoid the St. Petersburg paradox and related problems it is necessary that any utility function should be bounded (at least when considering the utility of the possible outcomes involved in a particular decision problem). Using functions such as the U_α described above is common practice and is perfectly acceptable provided that the (monetary) value of any possible outcome of the decision problem being considered is bounded, but care is needed when applying them to problems with unbounded possible rewards such as the St. Petersburg lottery.

It is only the boundedness of utility functions which is required, not that they should lie in a particular interval. This is because we are interested in maximising the expectation of a utility and expectation is a linear operator. As such, any linear transformation of a utility function will lead to precisely the same decisions and conclusions: indeed, such a transformation leads to a utility function which encodes exactly the same preferences.

Beyond Simple Utility Functions. Of course, we have been looking at relatively simple problems here and other complications can and do arrive in practice. One of the most immediate problems is that in many decision-making problems we have several different factors to consider. For example, when considering the treatment of patients a hospital must balance the cost of various treatment regimes with the resulting length and quality of life of the patients involved. It is not usually straightforward to map all of the conflicting requirements onto a simple linear scale and so another level of reasoning is required. This leads in a number of directions, one of which is multi-attribute utility theory which you may encounter in later courses.

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. . .

7. Games

Many real-world decision-making problems involve other people whose actions may have a complex interaction with our own. In order to deal with problems of this sort we need a more general approach than the one which we have considered so far. This brings us naturally the the idea of games.

7.1 What is a Game?

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

– Several *agents* or *players* make one 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.

Definition 7.1 (Simple 2-Player Game). *A simple two player game has the following structure:*

– *Player 1 chooses a move for a set* $D = \{d_1, \ldots, d_n\}$.
– *Player 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*, respectively, 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	$\cdots$	δ_m
d_1	$(R(d_1, \delta_1), S(d_1, \delta_1))$	$\cdots$	$(R(d_1, \delta_m), S(d_1, \delta_m))$
$\vdots$		$\ddots$	$\vdots$
d_n	$(R(d_n, \delta_1), S(d_n, \delta_1))$	$\cdots$	$(R(d_n, \delta_m), S(d_n, \delta_m))$

Notice that each entry in this “matrix” is actually a two element vector whose first element corresponds to the first player’s reward and the second element to that of the second player. It’s sometimes useful to consider a single player's payoff as a function of the possible decisions. Doing this leads us to a standard matrix of payoffs for each player.

Player 1 and player 2 have these individual payoff matrices:

Player 1	δ_1	$\dots$	δ_m
d_1	$R(d_1, \delta_1)$	$\dots$	$R(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$R(d_n, \delta_1)$	$\dots$	$R(d_n, \delta_m)$

Player 2	δ_1	$\dots$	δ_m
d_1	$S(d_1, \delta_1)$	$\dots$	$S(d_1, \delta_m)$
$\vdots$			$\vdots$
d_n	$S(d_n, \delta_1)$	$\dots$	$S(d_n, \delta_m)$

Some Examples of Two-Player Games

Here are some examples of widely-studied games which we will revisit later. Although apparently rather simple each of these games provides a concrete example of a game with a particular property of interest.

A particularly simple game with which you may already be familiar is often played by children and the extremely bored:

Example 7.1 (Rock-Paper-Scissors). Each player picks from the same set of decisions:

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

Rock beats **S**cissors; **S**cissors beat **P**aper and **P**aper beats **R**ock. One possible payoff matrix (in which winning a game has a payoff one 1, losing of -1 and for which a draw receives 0) 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)

Notice that this payoff matrix tells us that each player values winning with any move equally. It also suggests that winning once and losing when playing a second time would have the same value as drawing both games. ◁

Another example which is rather more widely studied and which you are likely to have encountered before is the Prisoner's Dilemma:

Example 7.2 (The Prisoner's Dilemma). Again, each player picks from the same set of decisions, they must choose whether to stay silent once arrested or to betray their partner in crime:

$$D = \Delta = \{\text{Stay } \textbf{S}\text{ilent}, \textbf{B}\text{etray Partner}\}$$

If they both stay silent they will receive a short sentence (1 year); if they both betray one another they will get a long sentence (4 years); if only one betrays the other the traitor will be released and the other will get a long sentence (5 years, say).

One possible payoff matrix for this sort of problem is:

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

Notice that each player wishes to maximise this everywhere non-positive payoff. It might be less cumbersome to replace every element of this matrix with its negative to produce a matrix filled with losses; in which case, each player wishes to minimise this payoff! Just as in the case of decisions we can work equally well with rewards or losses provided that we are consistent and don't confuse the two.

This is an interesting game as it is clear how the players should act if they can cooperate. However, in the absence of explicit cooperation each player must act in accordance with their own best interests (in game theory it is assumed that they payoff encodes all of a players preferences between outcomes and so sympathy for ones fellow man is, in this context, not a valid reason to alter ones behaviour). ◁

A third example along similar lines also encodes a problem of conflicting objectives, albeit of a rather less extreme sort:

Example 7.3 (Love Story). A boy and a girl must go to either of:

$$D = \Delta = \{\textbf{F}\text{ootball}, \textbf{O}\text{pera}\}$$

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)

This encodes the fact that both players regard being in the same place as the best possible outcome, regardless of which place that is. The boy is reasonably happy if he goes to the football alone; and the girl if she goes to the opera alone. Both players are unhappy if the boy goes to the opera and the girl to the football alone.

In some sense, this example allows each player to maximise their own worst case enjoyment by participating in their own favourite activity although if they both take such a course of action it guarantees that they will never meet. We will see that the less extreme nature of this payoff matrix means that things work a little better than they do in the case of the prisoner's dilemma. ◁

Each of these games exhibits a number of special features, a few of these are particularly significant. 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 $D = \Delta$ in these cases. However, although $D = \Delta$ in all three of these examples, it need not be the case and many interesting problems provide different players with different possible moves.

Uncertainty in Games

As neither player knows what action the other will take, there is uncertainty.

Thankfully, the subjective 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. Notice that they need not believe that their opponent is acting in a classically random way in order for these probability distributions to be meaningful within a subjective Bayesian setting. We will, however, see below that sometimes acting in a "genuinely random" manner may have advantages.

Just as in a decision problem, we can think about expected rewards and this may be a useful thing to do. For player 1, the expected reward of move d_i is:

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

notice that the expectation is taken over the actions of the second player in this case (i.e. δ is the random variable). Whilst, for player 2, we have

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

with the expectation over the actions of the first player (d is the random variable).

Expected rewards in this context tell us how the player would perform on average in a large number of repetitions of the game against a player whose strategy really did lead them to play according to the probability distribution used to assign weight to each of their possible moves. Again, it is quite possible to justify the use of expectation as a criterion on grounds of linearity alone without reference to ensembles of repeated experiments.

At this stage a number of interesting questions are beginning to arise, including these:

– When can a player act without considering what the opponent will do? That is, when is player 1's optimal 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?

7.2 Separability and Domination

Separable Games. If there is a way of writing the reward functions as one component due to the first player's actions plus a second component due only to the second players actions then this means that in some sense there is no interaction between the choices made between the two players in terms of the reward functions.

A game is *separable* if there exist functions $r_1 : D \to \mathbb{R}$, $s_1 : D \to \mathbb{R}$, $r_2 : \Delta \to \mathbb{R}$ and $s_2 : \Delta \to \mathbb{R}$ such that

$$\begin{aligned}R(d,\delta) \equiv& r_1(d) + r_2(\delta)\\ S(d,\delta) \equiv& s_1(d) + s_2(\delta)\end{aligned}$$

If such a decomposition exists, then the effect of one player's choice on either reward does not depend upon the other player's choice. The contribution of the two players' actions to the reward functions can be separated:

$$\begin{aligned}\bar{R}(d_i) =& r_1(d_i) + \sum_{j=1}^{m} p(\delta_j) r_2(\delta_j)\\ \bar{S}(\delta_j) =& \sum_{i=1}^{n} q(d_i) s_1(d_i) + s_2(\delta_j)\end{aligned}$$

Strategy in Separable Games. We will shortly see that this property makes it extremely easy to see how each player should act in games whose reward functions have this separability property.

If a game is separable and the two players aim to maximise their expected rewards, then they can exploit that separability. Looking at the expression for their expected rewards its clear that the first player can influence only the first term in the expansion of $\bar{R}$ and the second player can influence only the second term in the expansion of $\bar{S}$ (and they are interested only in their own reward functions). As such, the first player may as well make the decision which maximises the first term in their (expected) reward function and the second player should act in order to maximise the second term in their (expected) reward function.

The conclusions which we draw are simple:

- Player 1's strategy should depend only upon r_1 as the decision they make doesn't alter the additive contribution to the reward from r_2.
- Player 2's strategy should depend only upon s_2 as the decision they make doesn't alter the additive contribution to 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:

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

In many problems these sets will contain a single element; if they contain more than one element then the players should regard any action within these sets as being equivalent and equally good. In either case each player's strategy should be, simply, to maximise the component of the reward which they deterministically control.

The Prisoner's Dilemma and Separability. It can be shown that the prisoner's dilemma is a separable game:

- Let $r_1(S) = 0$ and $r_1(B) = 1$.
- Let $r_2(S) = -1$ and $r_2(B) = -5$.
- The game is symmetric, so $s_2 = r_1$ and $s_1 = r_2$.
- Now, $R(d, \delta) = r_1(d) + r_2(\delta)$.
- And $D^\star = \{B\}$ as $r_1(B) > r_1(S)$.
- Similarly for the second player, whose payoff function has an equivalent decomposition, $\Delta^\star = \{B\}$ as $s_2(B) > s_2(S)$.

This is the so-called paradox of the prisoner's dilemma: both players acting rationally and independently leads to what would seem to many to be the worst possible solution!

Note that the important thing here is that functions r_1 and r_2 of the correct form exist and so we can decompose the reward as a sum. The precise values of these functions are not as important as there is a degree of freedom in their specification. The idea of separability is that we may write player one's reward function, R, as a sum of a component due to their actions r_1 and another component due to their opponent's actions r_2 (and the same thing for player 2, with different functions).

It is useful to have a mechanistic approach to the calculation of functions r_1 and r_2 from the overall reward. As we only know $R(d, \delta)$ we can't obtain a unique value for $r_1(d)$ or $r_2(\delta)$, because we could add any constant to r_1 as long as we subtract the same constant from r_2 (as the sum of r_1 and r_2 will remain the same if we do this).

So, it's convenient to begin by saying that $r_1(S) = 0$ (although it could be any other constant as we have an under-constrained problem).

Now we've set our constant level, all of the other values of r_1 can be calculated by saying that:

$$R(S, S) = r_1(S) + r_2(S)$$
$$R(B, S) = r_1(B) + r_2(S)$$

subtracting the second of these from the first:

$$R(S, S) - R(B, S) = r_1(S) - r_1(B)$$

which means that

$$r_1(B) = R(S, S) - R(B, S) + r_1(S)$$
$$r_1(B) = (-1) - (0) + 0 = +1$$

And in a case with more actions available to player 1, you could obtain $r_1(d)$ for each of them in turn using the same approach (i.e. comparison of $R(S, S)$ with $R(d, S)$).

Given r_1, it's easy to calculate r_2:

$$r_2(S) = R(S, S) - r_1(S) = (-1) - 0 = -1$$
$$r_2(B) = R(S, B) - r_1(S) = (-5) - 0 = -5$$

So, by comparison and based upon the assumption that the game is separable we've calculated the values of the components of the first player's reward function. The same could be done for the second player's. At this stage, you must check that

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

for all value of d and δ (if this is the case then the game was separable and we have decomposed; if it's not then we either made a mistake or the game is not separable).

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 (we simply rewrite the payoff matrix, replacing each payoff with the utility of that payoff and proceed in exactly the same way).

For a given probability mass function p, player 1 has:

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

Whilst for given q, player 2 has:

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

We want p and q to be consistent with the assumption that the opponent is rational. We wish to 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 and it may not be precisely what you would expect:

– 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

$$\vdots$$

– 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.

So if x is common knowledge then all of the players know x and each player knows that each of the other players know x and each player knows that each of the other players know that they know that the other players all know x and... This is an example of an infinite regress, which can lead to some complication in games which are not finite. The games we shall consider are all simple enough that no problems will arise.

Domination

We have seen that in separable games, some subset of a player's moves are simply better than the others. Whatever their opponent chooses to do they will benefit more from playing one of these preferred moves to one of the other moves that they have available. In fact, we can generalise this concept to other games. It isn't surprising that there are often actions available which are inferior in all circumstances to some other move.

We call a move *dominant* if it is better than any of the other possible moves regardless of how an opponent acts. That is, 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:

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

We can also look at the definition the other way around: a dominant move is one which is at least as good as any of the others under all circumstances; sometimes it will be useful to identify those moves which are never better than at least one of the other available moves (as we'd be inclined to avoid playing moves of this type). 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:

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

The following theorem formalises the idea which is rather obvious from our definition that if there is a unique dominant move then a player should always play that move (if there are several dominant moves then they are exactly as good as one another and it doesn't matter which of them is played).

Theorem 7.1 (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

$$\sum_j R(d_i, \delta_j)p(\delta_j)$$

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

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

whatever may be p. □

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 7.4 (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 an example of what is known as a discriminating solution. In this game, and other games with such a solution, rational players may arrive by simple logical arguments at a strategy which is as good as it can be for either one of them given that their opponent is rational.

Iterated Strict Domination. In fact, we can use the idea of common knowledge to justify carrying out the elimination of dominated strategies iteratively, with each player eliminating any strategies which are dominated conditional upon the knowledge that the other player will not use any strategy which they have already seen dominated. This leads to the following systematic approach to eliminating strategies from consideration deterministically (in the particular case of two player games).

The players take it in turns to eliminate any of their remaining strategies which are strictly dominated. The procedure is repeated until neither player has any remaining strictly-dominated strategies. Slightly more formally:

1. Let $D_0 = D$ and $\Delta_0 = \Delta$. Let t = 1
2. Player 1 checks D_{t-1} to see if it contains one or more strictly dominated moves – given that player 2 must choose from Δ_{t-1}. Let D'_t be the set of such moves. Let $D_t = D_{t-1} \setminus D'_t$.

3. Player 2 updates Δ_{t-1} in the same way noting that player 1 must choose a move from D_t.
4. If $|D_t| = |\Delta_t| = 1$ then the game is solved.
5. If $|D_t| < |D_{t-1}|$ or $|\Delta_t| < |\Delta_{t-1}|$ let $t = t + 1$ and goto 2.
6. Otherwise, we have reduced the game to the simplest form we can by this method.

It's probably sensible to employ this strategy whenever presented with a new game as it will reduce the game to a simpler form if any strategies are dominated. In some cases, it will leave a discriminating solution.

Example 7.5 (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 — we are, of course, free to start with whichever player seems more convenient. 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 and 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". ◁

7.3 Zero-Sum Games

In order to make further progress, it is convenient to consider a particular class of games. The zero-sum games described in this section are relatively easy to analyse and have some particularly useful structure.

Purely Competitive Games. In a *purely competitive* game, one player's reward is improved only at the cost of the other player. This means, that if $R(d', \delta) = R(d, \delta) + x$ then $S(d', \delta) = S(d, \delta) - x$. That is, if player 1 deciding to play d' instead of d would improve their payoff by x then it must reduce player 2's payoff by x.

As this must be true for any d and d', it provides the general relationship:

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

That is, the sum over all players' rewards is the same for all sets of moves which they could choose to make. It doesn't change the domination structure or the ordering of expected rewards if we add a constant to all rewards (just as linear transformations of utility make no difference to expected-utility-based decisions in the context of decision theory). 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:

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

We need specify only one payoff and we can recover the other one as its negative. Payoff matrices may be simplified to specify only one reward[1].

Example 7.6 (Rock-Paper-Scissors is a zero-sum game). We may write the the payoff matrix as:

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

◁

Noting that this is simply the first player's payoff matrix with the second player receiving a payoff equal to its negation. As we shall see later, it can be convenient to use standard matrix notation, with $M = [m_{ij}]$ and $m_{ij} = R(d_i, \delta_j)$. That is, the matrix, M has element m_{ij} as the entry in row i, column j and m_{ij} is equal to the first player's payoff if they choose move d_i and their opponent chooses δ_j.

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 player could exploit that commitment (if they knew what it was, that is). We need a strategy for dealing with such games. Remembering the approaches we considered when dealing with simple decision problems, it would seem as though this might be a case for a maximin approach: one which makes the worst case scenario as good as it can be: at least this will minimize the harm which a cunning opponent could cause us.

In the case of zero-sum games, at least, a maximin strategy doesn't immediately seem too far-fetched as an opponent attempting to maximise their own payoff function in such a setting is at least trying to make things as bad as they can for us. Working on a maximin basis is somehow equivalent to assuming that they are rather successful at doing so.

Maximin Strategies in Zero-Sum Games. If a player adopts a maximin strategy, he is acting as though 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, in a zero-sum game, player 1's expectation of 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 also swap the two players to obtain a maximin strategy for player 2.

If we consider a symmetric game it seems a little odd that player 1 is assuming that they will always do as badly as possible given their choice of move and simultaneously that their opponent will do as well as they can. For example, consider the rock paper scissors game once again:

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

Example 7.7 (Rock-Paper-Scissors and Maximin). Let $M = [m_{ij}]$ denote the payoff matrix for the RPS game. Then,

$$\min_j R(d_i, \delta_j) = \min_j m_{ij} = -1 \text{ for all } i.$$

Thus any move is maximin for player 1, who expects to receive a pay-out 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? Although maximin strategies can be seen as pessimistic ones it seems a little too pessimistic for both players to be confident that they will always lose in a game as equally matched as this one. ◁

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 in the sense that they reach inconsistent conclusions.
– It cannot be common knowledge that *both* players will adopt a maximin strategy!

If either player believes that the other will adopt a maximin strategy in this sense, they shouldn't themselves adopt one. In fact, even if a player really believes their opponent can predict their move then they can use randomization to make their action less predictable and hence improve their prospects.

Mixed Strategies. Actually, random strategies feature prominently in game theory. It is very often more profitable to adopt a strategy which contains an element of randomization in order to protect against the various possible approaches that an opponent could take.

A *mixed strategy* for player 1 is a probability distribution over D. If a player has mixed strategy $\mathbf{x} = (x_1, \dots, x_n)$ 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 (the only difference being that they must consider the moves available to them, Δ, rather than those available to player 1.

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:

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

The final equality, in which M is the first players payoff matrix, illustrates one of the advantages of viewing this payoff matrix as a matrix: if the mixed strategy of each player is represented by a standard column vector then the expected reward can be calculated using simple matrix multiplication.

Maximin Revisited. When we first considered the use of maximin strategies in games we implicitly assumed that each player would adopt a pure strategy whilst assuming that their opponent could predict what that strategy would be. It would seem likely that a player would expect to do better by adopting a randomized mixed strategy if they felt that their opponent was able to predict their strategy. If we consider the maximin approach applied to mixed strategies we will obtain better results than we did previously.

Player 1's maximin *mixed* strategy is the $\underline{x}$ which maximises that players worst-case expected payoff given that such a strategy is adopted:

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

and by playing such a strategy, they obtain an expected payoff of

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

Player 2's maximin *mixed* strategy is the $\underline{y}$ which maximises their own worst-case expected payoff and as we are considering zero-sum games, this is equivalent to the strategy which minimises their opponents best-case expected reward:

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

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 7.2 (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.*

This theorem tells us that if both players adopt maximin strategies then they both have consistent expectations of their rewards. It also provides a value with which player 1 can bound their expected reward below (with respect to all possible strategies which player 2 can adopt) by a judicious choice of mixed strategy and with which player 2 can bound player 1's expected reward above (with respect to all possible strategies which player 1 can adapt) by a good choice of mixed strategy.

This provides us with the concept of a solution for this type of game. However, there are some things which this theorem does not tell us. In particular: it does not tell us what the strategies which achieve this value are — and the strategies $\underline{x}$ and $\underline{y}$ which achieve this value may not be unique.

Before considering how can we find suitable strategies in general it is useful to look at a simple example.

Example 7.8 (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.

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:

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

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 and this is the approach which would be taken when dealing with these problems. If one of the two players has only two possible decisions, however, a simple graphical method can be employed.

Graphical Solution, Part 1: Player 1's approach. In this section we shall consider how to deal with two-player zero-sum games in which player 1 has only two moves available to them. The same approach can, of course, be adopted if player 2 has only two moves available to them, even if the first player has a wider range of options, by exchanging the rôles of the two players.

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 described by treating these expected rewards as a function of x we obtain the following graph:

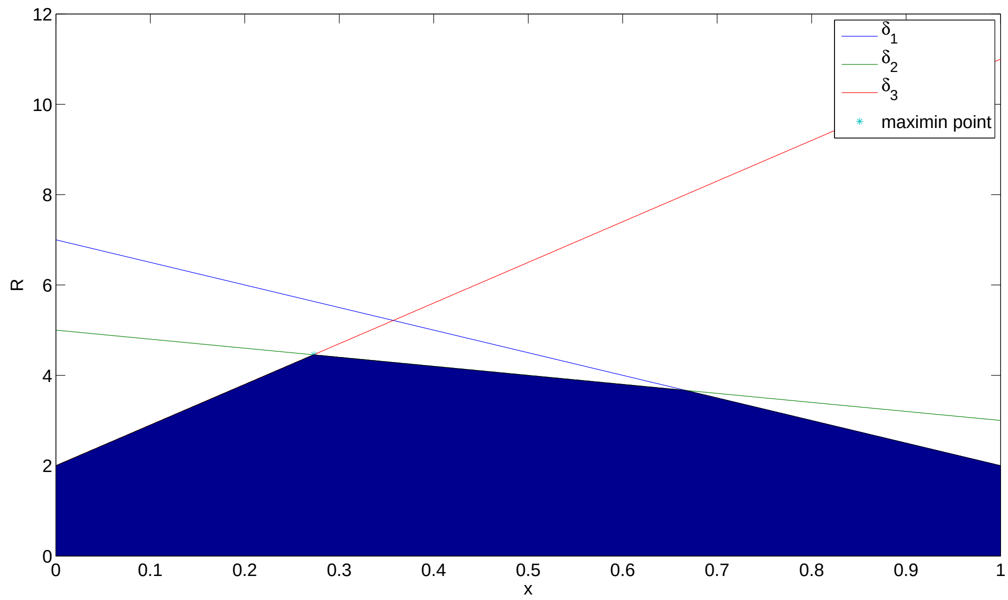

The maximin strategy maximises the player's expected return in the worst case (i.e. when their opponent takes the approach which the player would least like them to). In terms of our graph, this means we choose x to maximise the distance between the lowest of the lines and the ordinate axis (because the other player can control which of these lines is relevant, so making the worst of them provide as good a result as possible is the best that the first player can do). In this particular case, we can see from the graph that this is at the point where the lines associated with δ_2 and δ_3 intersect, at $x^\star$ which solves:

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

Notice that we have identified which lines are relevant by using the graph but have then obtained an accurate value for $x^\star$ algebraically rather than by trying to read the value from the graph.

Player 1's maximin mixed strategy is $(x^\star, 1 - x^\star) = (3/11, 8/11)$. Playing this strategy, his expected return is:

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

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) as the fundamental theorem tells them that they can at best restrict the first player to an expected reward of V_1. This means that they may consider a mixed strategy $(0, y, 1-y)$. By the fundamental theorem, player 2's maximin strategy leads to the same expected payoff for player 1 as his own maximin strategy:

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

And so, they should play $y^\star$ which solves:

$$V_2 = 3y + 11(1-y) = 49/11$$
$$8y = (121-49)/11 = 72/11 \qquad \Rightarrow y^\star = 9/11$$

Leading to a mixed strategy $(0, y^\star, 1-y^\star) = (0, 9/11, 2/11)$.

Of course, this technique applies completely to any game with the structure being considered, the only difference being which of the strategies which need to be considered (depending upon which lines intersect when the first player considers their approach).

Example 7.9 (Spy Game). Consider this stylised game.

A spy has escaped and must choose to flee down a *river* or through a *forest*. Their guard must choose to chase them using a *helicopter*, a pack of *dogs* or a *jeep*.

They agree that the probabilities of escape are as given in this payoff matrix and view the probability of escape as being the quantity which they wish to maximise/minimise the expectation of:

	H	D	J
R	0.1	0.8	0.4
F	0.9	0.1	0.6

Both players wish to adopt maximin strategies. 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:

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

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

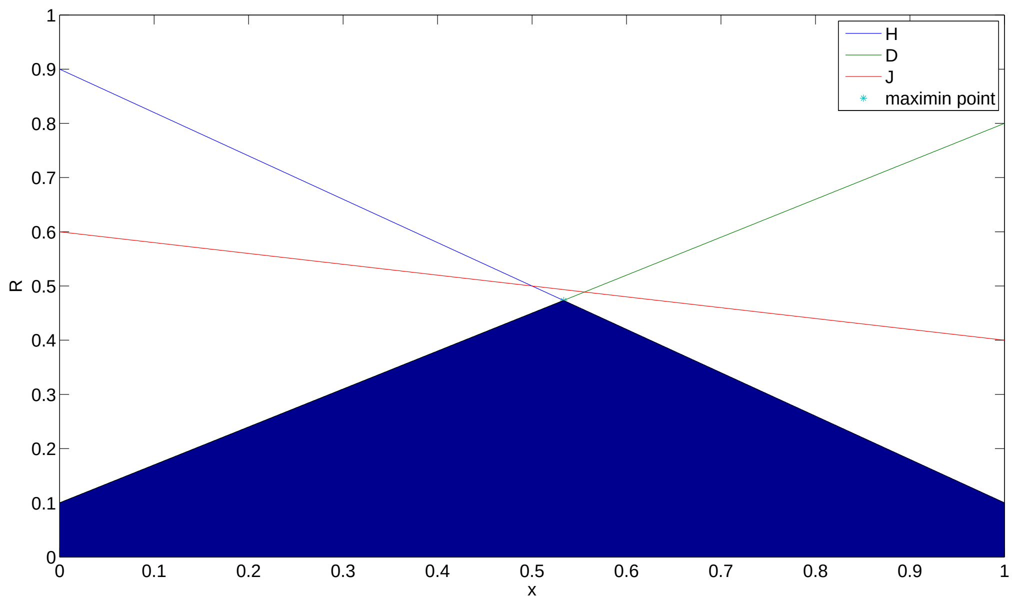

The maximin solution is the intersection of the lines for strategies D and H. This occurs at the solution, $x^\star$ of:

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

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

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

◁

We have been considering games with a rather specific structure: two player zero-sum games (and an even more specific form when dealing with the graphical solution approach described above). Some natural questions spring to mind in relation to the fundamental theorem. Unfortunately, the answers to most of the more interesting ones is "no". In particular: The "fundamental theorem" does not generalise to games of more than two players nor does it generalise to non-zero-sum games.

Although zero sum games are widely studied as they have a structure which is amenable to analysis, in many ways games with an element of co-operation are much more interesting.

7.4 Selected Game-Theoretic Concepts (Not Examinable)

This section briefly summarises some of the things which become important when dealing with more general games. These ideas don't form a core part of the current course and will not be examined. They are included for the benefit of anyone interested in these ideas and it is likely that they are all things you will encounter in any subsequent course on game theory which you choose to take.

We will consider what we need to do in order to generalise what we have thought about so far to the case of games that do not have the zero-sum property in which it is possible for one player to improve their own outcome without their opponent suffering commensurately.

We have seen that 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: neither player is explicitly attempting to minimise their opponent's payoff when that payoff is not the negative of their own. It's not obvious what properties a solution should have in these more general games.

However, it is now generally agreed that the important properties relate to ideas of equilibrium and stability. A stable solution to a game is one which is in some sense resistant to perturbation: no player should be inclined to move unilaterally away from it. These ideas are connected with those of balance: it must somehow be in the interest of all players to remain at a particular point in the solution surface or some subset of the players are likely to try to move away from that point. It is also useful to consider how good players are likely to think particular solutions are: there may be many stable solutions and some may be somewhat better than others.

Pareto Optimality. To begin with, it is useful to consider a popular optimality concept from the field of economics. 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, such as Rock-Paper-Scissors, all sets of pure strategies are Pareto optimal as any improvement in payoff for one player is matched by a reduction in payoff for another.

Nash Equilibrium. A collection of strategies (one per player) in a game is a *Nash equilibrium* if no player can improve their expected reward by unilaterally changing their strategy. This means that no player can do better by changing their own strategy; it would require cooperation with another player — if it is possible to do better at all.

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

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

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

These inequalities simply formalise the above statement.

Nash Equilibria in 2 Player Zero Sum Games. The particular structure of zero-sum games can be connected to the solution of more general games, and to Nash equilibria in particular.

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 $S = -R$). This property does not extend to non zero-sum games.

Nash Equilibria and the Prisoner's Dilemma. As a particular example, we can return to the familiar 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: any change by a single player makes things strictly worse for that player.

(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. In particular, a player choosing to betray their opponent would make things strictly worse for that opponent.

However, the (S, S) strategy set is not stable: it is not an equilibrium as either player can unilateral improve their own reward by betraying their opponent.

Solutions of Games

We are now in a position to formulate the idea of solutions of games in a number of ways. The fact that it was not clear from the outset exactly what properties a "solution" to a game should have is reflected in the existence of a number of different notions of a solution.

Solutions I: The Nash Sense. Two pairs of strategies $(\underline{x}, \underline{y})$ and $(\underline{x}', \underline{y}')$ are interchangeable with respect to some property if $(\underline{x}', \underline{y})$ and $(\underline{x}, \underline{y}')$ have that same property. A game is *Nash solvable* if all equilibrium pairs are interchangeable (with respect to being equilibrium pairs).

That is, if $(\underline{x}_1, \underline{y}_1), \ldots (\underline{x}_n, \underline{y}_n)$ are all the equilibrium pairs of a game then the game is Nash solvable if (x_i, y_j) is an equilibrium pair for any $i, j = 1, \ldots, n$.

All zero-sum games are Nash solvable; not many other games are.

The Nash sense of solution is rather a strong one and it is not applicable to very many games and so it's useful to look at some other formulations of "solutions" to games.

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).

8. Closing Comments

This is a course entitled “Games and Decisions” and its motivation, throughout, has been to answer one question: how should a rational being act; indeed, by what process should such a being make decisions?

Along the way, it has been necessary to consider a number of auxiliary topics which may at the time have seemed somewhat removed from that objective. However, each part of the course was either directly concerned with the making of decisions or was an essential precursor: some of the early material covered the philosophical foundations of inference and decision making, other parts covered some technical machinery required to implement its logical conclusions.

Hopefully, it is now clear how the various parts of the course are connected and why it is necessary to have a firm grasp on the philosophy behind probability and inference if one is to use probabilistic arguments to provide real-world action guidance.

If this course has been a success you will now have an understanding of the concepts which underly Bayesian inference and subjective probability *as well* as how to apply these ideas in the context of decision making.

It is probably apparent that we have barely scratched the surface of these two major topics; game theory is addressed in more detail in EC220/221 “Mathematical Economics” next year, more advanced aspects of decision theory are addressed in ST301/ST413 Bayesian Statistics & Decision Theory.

Much of what has been considered in this module is investigated much more thoroughly in the recommended texts. The book of Körner, in particular, is an extremely entertaining and readable introduction to all that is important in the theory of Decisions and Games.