UNIVERSITY OF LONDON | INTERNATIONAL PROGRAMMES

Managerial economics

D.J. Reyniers and M. Selvaggi

MN3028

2016

Undergraduate study in
Economics, Management, Finance and the Social Sciences

This is an extract from a subject guide for an undergraduate course offered as part of the University of London International Programmes in Economics, Management, Finance and the Social Sciences. Materials for these programmes are developed by academics at the London School of Economics and Political Science (LSE).

For more information, see: www.londoninternational.ac.uk

THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

This guide was prepared for the University of London International Programmes by:

Professor Diane Reyniers, Professor of Managerial Economics and Strategy, Department of Management, London School of Economics and Political Science

and

Dr Mariano Selvaggi, Independent Academic Adviser; previously of the London School of Economics and Political Science.

This is one of a series of subject guides published by the University. We regret that due to pressure of work the authors are unable to enter into any correspondence relating to, or arising from, the guide. If you have any comments on this subject guide, favourable or unfavourable, please use the form at the back of this guide.

University of London International Programmes
Publications Office
Stewart House
32 Russell Square
London WC1B 5DN
United Kingdom

www.londoninternational.ac.uk

Published by: University of London

Contents

Chapter 1: Introduction

1.1 Road map to the guide

Welcome to the subject guide for **MN3028 Managerial economics**. This course considers a variety of exciting topics that lie at the intersection of management and economics. You will learn how economics can provide a useful framework to think about managerial issues and analyse problems that businesses and societies commonly face in the real world. We hope that this knowledge and critical thinking will serve you well in your future studies and professional career.

This guide will help you to engage actively with the course and develop a robust understanding of the subject matter. It provides a framework for covering the topics in the syllabus and directions to the Essential reading. However, the guide is not a substitute for the careful study of the readings listed at the beginning of each chapter. For each topic in the syllabus it is advisable to start with the relevant chapter of the guide, then do the reading for that particular topic, then come back to the guide and attempt the sample exercises at the end of the chapter. This will help you check your understanding of the topic by applying your knowledge to solve specific problems.

The subject guide is divided into 20 chapters, consisting of an introductory chapter; five main content blocks covering the syllabus (namely, risk and information, game theory and strategic behaviour, demand and supply theory, market structure and competition, and corporate governance); and a concluding chapter. It aims to:

- provide an effective framework for the study of the subject
- introduce you to the relevant subject material
- present the material in a structured/digestible format
- guide you towards appropriate learning resources
- enable you to approach managerial decision problems using economic reasoning
- discuss business practice topics using an analytical approach, relying on equations and numerical insight
- encourage you to take an active approach to learning by reading recommended material; undertaking learning activities; and participating in discussion via the virtual learning environment (VLE).

1.2 Introduction to the subject area

This course is intended as an intermediate economics course for BSc (Management) and BSc (Economics) students. The main objective of the course is to enable you to approach managerial decision problems using economic reasoning. As such, it is less theoretical than a microeconomic principles course and more attention is given to topics that are relevant to managerial decision making. For instance, business practices such as price discrimination, transfer pricing, bundling, resale price maintenance and corporate governance are discussed. Most topics are analysed using equations and numerical examples – that is, an analytical approach is used. The theories that are presented are not practical ‘recipes’; they are meant to give you insight and train your mind to think more like an economist. The emphasis throughout is on the application of knowledge and problem-solving rather than pure memorising and essay writing.

Why should economics and management students study economics? The environment in which modern managers operate is an increasingly complex one. It cannot be navigated without a thorough understanding of how business decisions are and should be taken. Intuition and factual knowledge are not sufficient. Managers need to be able to analyse; to put their observations into perspective; and to organise their thoughts in a rigorous, logical way. The main objective of this course is to enable you to approach managerial decision problems using economic reasoning. At the end of studying this subject, you should have acquired a sufficient level of model-building skills to analyse microeconomic situations of relevance to managers. The emphasis is therefore on 'learning by doing' rather than repeating memorised material. You are strongly advised to develop problem-solving skills using the exercises at the end of each chapter and past examination papers as well as ancillary material available through the VLE to complement and clarify your thinking.

1.3 Syllabus

The course covers basic topics in microeconomics such as supply and demand, consumer theory, labour supply, asymmetry of information, neo-classical view of the firm, production, costs, factor demands, perfect competition, monopoly, monopolistic competition, oligopoly, cartels and tacit collusion. We also analyse some newer material regarding alternative theories of the firm, internal organisation of the firm, market structure, efficiency wages, incentive structures, corporate governance as well as some industrial organisation theories of commonly used pricing practices.

The following topics also form part of the course syllabus:

- individual (one person) decision making under uncertainty, attitudes towards risk and the value of information
- theory of games and strategic decision making, including its applications to oligopoly, collusion among firms, product differentiation, entry deterrence and other market practices
- the effects of asymmetric information in areas such as bargaining, bidding and auctions, situations of moral hazard and adverse selection
- corporate governance in modern organisations.

Some knowledge of constrained maximisation and Lagrangian functions would be helpful for students taking this subject, although this is not a prerequisite.

1.4 Aims of the course

This course has two main objectives, which directly relate to the major themes that will be emphasised throughout this subject guide. The course aims to:

- enable students to approach managerial decision problems using economic reasoning
- present business practice topics using an analytical approach, using equations and numerical insight.

1.5 Learning outcomes for the course

On completion of this course, and having completed the Essential reading and exercises, you should be:

- prepared for Marketing and Strategy courses by being able to analyse and discuss consumer behaviour and markets in general
- able to analyse business practices with respect to pricing and competition
- able to define and apply key concepts in decision analysis and game theory
- able to confidently analyse different market structures and equilibrium outcomes in each of them.

1.6 Overview of learning resources

1.6.1 The subject guide

This subject guide has been written in such a way that you can obtain a good basic understanding of the topics in the syllabus, before going on to read the various other resources to broaden and deepen your knowledge. It has been designed to give you a clear indication of the level of analysis and detail that will be expected of you in the examination. Please note that although the subject guide provides crucial academic guidance, it is different from a traditional academic textbook and it is specifically intended to work **alongside** the other resources detailed below. We strongly encourage you to consult all available resources as you progress with your studies of the subject.

The topics covered in this subject guide are closely aligned with those taught to second year BSc (Management) students at the London School of Economics and Political Science (LSE). Particularly:

- basic microeconomics (i.e. consumer theory, labour supply, neoclassical theory of the firm, factor demands, competitive structure, government intervention, etc.)
- individual (one-person) decision making under uncertainty and the value of information; the theory of games. The latter considers strategic decision making with more than one player and has applications to bargaining, bidding and auctions, oligopoly and collusion, the effects of asymmetric information on market outcomes (one decision-maker has more information than the other) – particularly situations of moral hazard (post-contractual opportunism) and adverse selection (pre-contractual opportunism)
- some interesting topics in labour economics such as efficiency wages, incentive structures and human resource management.

Each chapter of the subject guide contains learning outcomes to help you check your progress and level of understanding of key topics. Each chapter also contains worked-out examples throughout and sample exercises at the end of the chapter. You are strongly encouraged to attempt these exercises as well as the Sample examination paper on the VLE. Simply reading the subject guide and the recommended textbooks without trying to solve exercises on your own will not be sufficient for you to gain benefit from the subject or prepare you for the examination.

1.6.2 Mathematics for managerial economics

The mathematical Appendix 1 in this subject guide gives an indication of the level of mathematics required for the course. Review it! If you are having difficulty with the mathematics in this course, you might find the following textbook useful:

Png, I. *Managerial economics*. (New York: Routledge, 2016) 5th edition [ISBN 9781138810266].

The above textbook does not assume a prior knowledge of economics and offers a less mathematical introduction to managerial economics. However, it is recommended as preliminary reading only and should **not** be used as a substitute for the subject guide. The level of mathematics it uses is much lower than that required for this course and it does not cover the topics in detail. The different models are treated in a very basic way. It can, however, be a useful back-up reference if you don't have the basic knowledge required to understand a topic. In particular, if you use this book, it is recommended that you work through the mathematical appendices that are provided at the end of each chapter. It is not necessary to study the entire book in depth.

1.6.3 Essential reading

This guide is intended for an intermediate-level course on economics for management. It is more self-contained than other subject guides might be. Having said this, we do want to encourage you to read widely from the reading list. Seeing things explained in more than one way should help your understanding. It is therefore vital that you do all the Essential reading specified at the beginning of each chapter, to complement and expand the material discussed in the guide.

In addition to studying the material, it is essential that you practise problem solving. This is important not only to check your own understanding of the topic, but also to prepare effectively for the written examination. Each chapter contains some sample questions and working through these and problems in the recommended texts is excellent preparation for success in your studies. You should also attempt past examination questions from recent years; these are also available online on the VLE.

Throughout this guide, we will recommend appropriate chapters in the following textbook:

Varian, H.R. *Intermediate microeconomics: a modern approach*. (New York: W.W. Norton and Co., 2014) 9th edition [ISBN 9780393123968].

Detailed reading references in this subject guide refer to this specific edition of Varian's textbook. New editions of this textbook may of course have been published by the time you study this course. You can use a more recent edition of any of the books we recommend; use the detailed chapter and section headings and the index to identify relevant readings. Also check the VLE regularly for updated guidance on readings.

1.6.4 Further reading

We list a number of journals throughout the subject guide, which you may find interesting to read. In certain chapters we also list relevant chapters of the following textbooks:

Allen, W.B., K. Weigelt, N.A. Doherty and E. Mansfield *Managerial economics: theory, applications, and cases*. (New York: W.W. Norton and Co., 2013) 8th edition [ISBN 9780393912777].

Besanko, D., R. Braeutigam and K. Rockett *Microeconomics*. (Singapore: John Wiley & Sons, 2015) 5th edition, International Student Version [ISBN 9781118716380].

Besanko, D., D. Dranove, M. Shanley and S. Schaefer *Economics of strategy*. (Singapore: John Wiley & Sons, 2013) 6th edition, International Student Version [ISBN 9781118319185].

Carlton, D.W. and J.M. Perloff *Modern industrial organization*. (Harlow: Pearson Education Ltd, 2015) 4th edition, Global Edition [ISBN 9781292087856].

Tirole, J. *The theory of industrial organisation*. (Cambridge, MA: The MIT Press, 1988) [ISBN 9780262200714].

We strongly encourage you to read those chapters as additional material to support your studies.

Please note that as long as you read the Essential reading you are then free to read around the subject area in any text, paper or online resource. You will need to support your learning by reading as widely as possible and by thinking about how these principles apply in the real world. To help you read extensively, you have free access to the VLE and University of London Online Library (see below).

1.6.5 Online study resources

In addition to the subject guide and the Essential reading, it is crucial that you take advantage of the study resources that are available online for this course, including the virtual learning environment (VLE) and the Online Library. You should use the Online Library and VLE regularly throughout your studies to support the learning material included in this subject guide.

You can access the VLE, the Online Library and your University of London email account via the Student Portal at: http://my.londoninternational.ac.uk

You should have received your login details for the Student Portal with your official offer, which was emailed to the address that you gave on your application form. You have probably already logged in to the Student Portal in order to register! As soon as you registered, you will automatically have been granted access to the VLE, Online Library and your fully functional University of London email account.

If you have forgotten these login details, please click on the 'Forgotten your password' link on the login page.

The VLE

The VLE, which complements this subject guide, has been designed to enhance your learning experience, providing additional support and a sense of community. It forms an important part of your study experience with the University of London and you should access it regularly.

The VLE provides a range of resources for EMFSS courses:

- **Electronic study materials:** All of the printed materials which you receive from the University of London are available to download, to give you flexibility in how and where you study.
- **Discussion forums:** An open space for you to discuss interests and seek support from your peers, working collaboratively to solve problems and discuss subject material. Some forums are moderated by an LSE academic.
- **Videos:** Recorded academic introductions to many subjects; interviews and debates with academics who have designed the courses and teach similar ones at LSE.
- **Recorded lectures:** For a few subjects, where appropriate, various teaching sessions of the course have been recorded and made available online via the VLE.

- **Audiovisual tutorials and solutions:** For some of the first year and larger later courses such as Introduction to Economics, Statistics, Mathematics and Principles of Banking and Accounting, audio-visual tutorials are available to help you work through key concepts and to show the standard expected in exams.
- **Self-testing activities:** Allowing you to test your own understanding of subject material.
- **Study skills:** Expert advice on getting started with your studies, preparing for examinations and developing your digital literacy skills.

Note: Students registered for Laws courses also receive access to the dedicated Laws VLE.

Some of these resources are available for certain courses only, but we are expanding our provision all the time and you should check the VLE regularly for updates.

Making use of the Online Library

The Online Library (http://onlinelibrary.london.ac.uk) contains a huge array of journal articles and other resources to help you read widely and extensively.

To access the majority of resources via the Online Library you will either need to use your University of London Student Portal login details, or you will be required to register and use an Athens login.

The easiest way to locate relevant content and journal articles in the Online Library is to use the **Summon** search engine.

If you are having trouble finding an article listed in a reading list, try removing any punctuation from the title, such as single quotation marks, question marks and colons.

For further advice, please use the online help pages (http://onlinelibrary.london.ac.uk/resources/summon) or contact the Online Library team: onlinelibrary@shl.london.ac.uk

1.7 Examination advice

Important: the information and advice given here are based on the examination structure used at the time this guide was written. Please note that subject guides may be used for several years. Because of this we strongly advise you to always check both the current Programme regulations, for relevant information about the examination, and the VLE where you should be advised of any forthcoming changes. You should also carefully check the rubric/instructions on the paper you actually sit and follow those instructions.

It is worth emphasising again that there is no substitute for practising problem solving **throughout** the year. The best way to internalise the key concepts of managerial economics is to work through the set problems and exercises that are included in this guide and also those that are available on the VLE. It is impossible to acquire a reasonable level of problem-solving skills while you are revising for the examination. You should consciously try to solve the exercises presented at the end of each chapter as well as past examination questions available on the VLE **at the same time** as you progress with your studies. This approach will help you test your true knowledge and understanding of the topic in real time.

The examination lasts for three hours and you may use a calculator. Detailed instructions are given on the examination paper. Read them carefully! The questions within each part carry equal weight and the amount of time you should spend on a question is proportional to the marks it carries. Part A consists of compulsory, relatively short problems of the type that accompanies each of the chapters in the guide. Part B is a mixture of some essay-type questions and longer analytical questions. In part B a wider choice is usually available. Please refer to the VLE for a Sample examination paper to help you prepare. It indicates the level and type of questions you should expect. We would strongly advise you to attempt to complete this paper within the time specified before looking at the solutions provided in the *Examiners' commentaries* on the VLE.

Remember, it is also important to check the VLE for:

- up-to-date information on examination and assessment arrangements for this course
- where available, past examination papers and *Examiners' commentaries* for the course, which give advice on how each question might best be answered.

1.7.1 Some advice and ideas on how to study

(This information was originally written by Tell Fallrath as an Appendix to a previous edition of the guide.)

- Check your understanding of each model in three ways:
    1. Can you explain its main arguments in a few simple words?
    2. Can you solve a basic model analytically?
    3. Can you draw the corresponding graphs?
- Do you understand the models and solutions intuitively?
- What is the overall context of a particular model that you study? How does it relate to what you already know?
- Try to summarise each topic yourself, perhaps by using mind maps. Remind yourself of the examples that you have studied and how they relate to the theory of each chapter.
- Don't memorise, but understand! There is hardly anything that you will need to memorise for this course. Instead make sure you feel comfortable with the exercises.
- Instead of reading the chapter over and over again, practise problem solving!
- Often students complain that there are not enough practice questions. You can create an infinite number of questions yourself by changing given examples slightly (i.e. change a Marginal Cost function from MC = 10 to MC = 20). Check how the results change and that you understand why they change in this particular way. For graphical questions, change an assumption and see how the graph changes.
- Don't blame any deficits in mathematics on the economics course! Throughout the year as soon as they arise, address any difficulties you may have with the mathematics.
- Be curious and get help if you don't understand something. If you are studying at an institution, for example, ask for help from your fellow students, teachers or past-year students.
- Relate the models to your day-to-day experience, which is much more rewarding and fun. Examples give powerful illustrations of how economic theory works. After all, economics is not an abstract science, but models the world around you.

Notes

Chapter 2: Decision analysis

2.1 Introduction

It seems appropriate to start a course on economics for management with decision analysis. Managers make decisions daily regarding selection of suppliers, budgets for research and development, whether to buy a certain component or produce it in-house and so on. Economics is in some sense the science of decision-making. It analyses consumers' decisions on which goods to consume, in what quantities and when, firms' decisions on the allocation of production over several plants, how much to produce and how to select the best technology to produce a given product. However, the bulk of economic analysis considers these decision problems in an environment of certainty. That is, all the necessary information is available to the agents making the decisions. Although this is a justifiable simplification, in reality, of course, most decisions are made in a climate of (sometimes extreme) uncertainty. For example, a firm may know how many employees to hire to produce a given quantity of output but the decision of whether or how many employees to lay off during a recession involves some estimate of the length of the recession. In a similar vein, oil companies take enormous gambles when they decide to develop a new field. The cost of drilling for oil, especially in deep water, can be extremely high (e.g. billions of US$) and the payoffs in terms of the future price of oil and gas as well as the likelihood of striking significant deposits can be very uncertain. Investment decisions would definitely be very much easier if uncertainty could be eliminated. Imagine what would happen if you could forecast interest rates and exchange rates next year with 100 per cent accuracy.

Clearly we need to understand how decisions are made when at least some of the important factors influencing the decision are not known for sure. The field of decision analysis offers a useful framework for studying how these types of decisions are made or should be made. It also provides insight into the cost of uncertainty or, in other words, how much a decision-maker is or should be prepared to pay to reduce or eliminate the uncertainty involved in the decision-making process. To illustrate the concept of value of information, consider the problem of a company bidding for a road maintenance contract. The costs of the project are unknown and the company does not know how low to bid to get the job. An important question to be answered in preparing the bid is whether to gather more information about the nature of the project and/or the competitors' bids. These efforts only pay-off if, as a result, better decisions are taken in the end.

Decision analysis is mainly used for situations in which there is one decision-maker whereas **game theory** deals with problems in which there are several decision-makers, each pursuing their own objectives. You will study game theory in the next chapter. In decision analysis any form of uncertainty can be modelled including that arising from unknown features of competitors' behaviour (as in the bidding example). However, when decision analysis models are used to solve problems with several decision-makers, the competitors are not modelled as rational and strategic agents (i.e. it is not recognised that they are also trying to achieve certain objectives, taking the actions of other players into account). Instead, decision theory takes the view that, as long as probabilities can also be attached to other decision-makers' actions, optimal decisions can be calculated. An obvious objection to this approach is that it is not clear how these probabilities become known to the decision-maker. Game theory avoids this problem as it takes a symmetric, simultaneous view. The reason decision analysis is used in these situations despite these shortcomings is that it is much simpler than game theory. For this reason and because some of the techniques of decision analysis (such as representing sequential decision problems on graphs or decision trees, and solving them backwards) can also be used in game theory, we study decision analysis first.

2.1.1 Aims of the chapter

The aims of this chaper are to consider:

- the concept of expected value of perfect information (**EVPI)** and how it can be used
- why we may want to use **expected utility** rather than expected value maximisation
- the concept of **certainty equivalent** and how it relates to expected value for a risk loving, risk neutral and risk averse decision-maker
- the application of decision analysis in insurance and finance.

2.1.2 Learning outcomes

By the end of this chapter, and having completed the Essential reading and exercises, you should be able to:

- structure simple decision problems in **decision tree** format and derive optimal decisions
- explain attitudes towards risk and the way they may impact on optimal decision-making
- calculate **risk aversion coefficients**
- calculate **EVPI** for risk neutral and non-risk neutral decision-makers
- apply the tools of decision analysis to solve new problems involving uncertainty.

2.1.3 Essential reading

Varian, H.R. *Intermediate microeconomics: a modern approach*. (New York: W.W. Norton and Co., 2014; 9th edition) Chapter 12: Uncertainty.

2.1.4 Further reading

You may also find the following book useful as additional reading:

Allen, W.B., K. Weigelt, N.A. Doherty and E. Mansfield *Managerial economics: theory, applications, and cases*. (New York: W.W. Norton and Co., 2013) 8th edition [ISBN 9780393912777], especially Chapter 14: Risk analysis.

2.1.5 Synopsis of chapter content

This chapter provides a framework for analysing optimal decision-making by a single decision-maker in the context of uncertainty about future outcomes (i.e. risk). It introduces decision trees and the concept of expected utility, examines rational agents' optimal decisions as a function of their risk appetite and presents a way to value perfect information in environments of uncertainty. A number of applications of the concepts to real and financial markets are also considered.

2.2 Decision trees

A **decision tree** is a convenient representation of a decision problem that helps decision-makers visualise their strategic future. It contains all the ingredients of the problem:

- the possible decisions
- the sources of uncertainty
- the payoffs to the decision-maker for each possible combination of probabilistic outcomes and decisions.

Drawing a decision tree forces the decision-maker to think through the structure of the problem faced, thus making the process of determining optimal decisions easier. It illustrates the various options available and provides an easy way to compare pay-offs across strategies. A decision tree consists of two kinds of nodes: decision or action nodes which are drawn as squares; and probability or chance nodes which are drawn as circles. The arcs leading from a decision node represent the choices available to the decision-maker at this point; whereas the arcs leading from a probability node correspond to the set of possible outcomes when some uncertainty is resolved. When the structure of the decision problem is captured in a decision tree, the payoffs are written at the end of the final branches and (conditional) probabilities are written next to each arc leading from a probability node. The algorithm for finding the optimal decisions is not difficult. Starting at the end of the tree, work backwards and label nodes as follows. At a probability node calculate the expected value of the labels of its successor nodes, using the probabilities given on the arcs leading from the node. This expected value becomes the label for the probability node. At a decision node (assuming a maximisation problem), select the maximum value of the labels of successor nodes. This maximum becomes the label for the decision node. The decision which generates this maximum value is the optimal decision at this node. Repeat this procedure, moving further to the left along the decision tree, until you reach the starting node. The label you get at the starting node is the expected pay-off obtained when the optimal decisions are taken. The construction and solution of a decision tree is most easily explained through examples.

Example 2.1

Cussoft Ltd., a firm which supplies customised software, must decide between two mutually exclusive contracts, one for the government and the other for a private firm. It is hard to estimate the costs Cussoft will incur under either contract but, from experience, it estimates that, if it contracts with a private firm, its profit will be £2 million, £0.7 million, or –£0.5 million with probabilities 0.25, 0.41 and 0.34 respectively. If it contracts with the government, its profit will be £4 million or –£2.5 million with respective probabilities 0.45 and 0.55. Which is Cussoft's optimal decision to maximise expected profit?

In this very simple example, Cussoft has a choice of two decisions – to contract with the private firm or to contract with the government. In either case its pay-off is uncertain. The decision tree with the payoffs and probabilities is drawn in Figure 2.1. The label of the top probability node is equal to the expected profit if the contract with the private firm is chosen.

The expected profit is the sum of the amount of money gained (or lost) if each outcome occurs, times the probability of occurrence of the outcome, which in this case is (0.25)(2) + (0.41)(0.7) + (0.34)(–0.5) = 0.617 (£ million). The label of the bottom probability node is equal to the expected profit if the contract with the government is chosen, which is (0.45)(4) + (0.55)(–2.5) = 0.425 (£ million). Working backwards it turns out that the optimal decision at the decision node is to choose the contract with the private firm since it generates the biggest expected profit. Optimal decisions are indicated by thick lines.

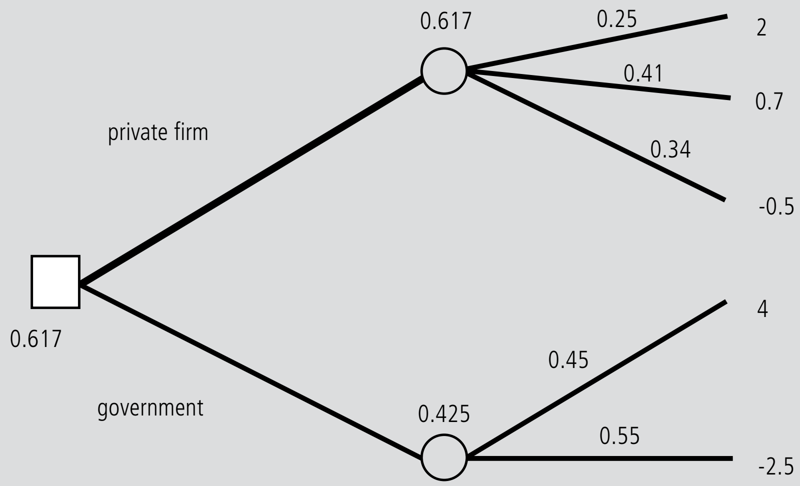

Figure 2.1: Decision tree for Example 2.1

Example 2.2

Suppose the chief executive officer (CEO) of an oil company must decide whether to drill a site and, if so, how deep. It costs £160,000 to drill the first 3,000 feet and there is a 0.4 chance of striking oil. If oil is struck, the profit (net of drilling expenses) is £600,000. If she doesn't strike oil, the executive can drill 2,000 feet deeper at an additional cost of £90,000. Her chance of finding oil between 3,000 and 5,000 feet is 0.2 and her net profit (after all drilling costs) from a strike at this depth is £400,000. What action should the executive take to maximise her expected profit? (Try writing down and solving the decision tree yourself without peeking!)

You should get the following result.

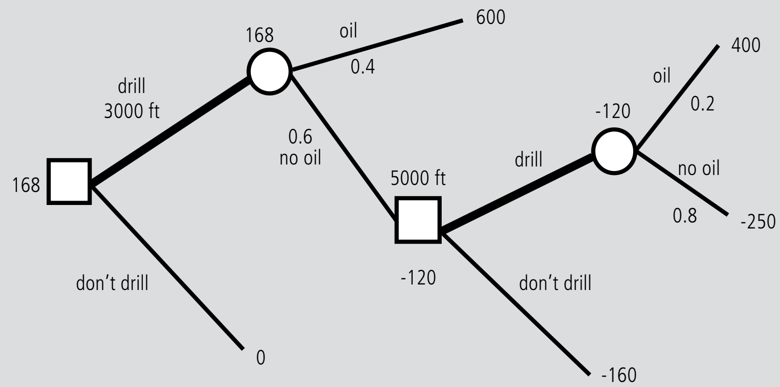

Figure 2.2: Decision tree for Example 2.2

2.3 Attitude towards risk

In the examples considered so far we have used the **expected monetary value (EMV) criterion** (i.e. we assumed that the decision-maker is interested in maximising the expected value of profits or minimising the expected value of costs). In many circumstances this is a reasonable assumption to make, especially if the decision-maker is a large company. To examine how risk may affect managerial behaviour, consider a manager facing the following choice:

1. A certain profit of £1,000,000.
2. A project with a 50:50 chance of a £3,000,000 profit or a £500,000 loss.

The EMV of the project equals:

0.50 (£3,000,000) + 0.50 (£500,000) = £1,250,000

– so to maximise expected profits, the manager should choose the project because it has a bigger EMV. However, it is reasonable to believe that many managers may nonetheless prefer the certainty of £1 million to the risky project. Indeed, this is more likely to be true for relatively small businesses for which losing £500,000 could be disastrous. Whether decision-makers will want to follow the EMV criterion in this situation will ultimately depend on their attitude towards risk. To appreciate that it may not always be appropriate to use EMV, consider the following story, known as the **St. Petersburg paradox**.

I will toss a coin and, if it comes up heads, you will get £2. If it comes up tails, I will toss it again and, if it comes up heads this time, you will get £4; if it comes up tails, I will toss it again and, this time, you will get £8 if it comes up heads, etc. How much would you be willing to pay for this gamble? I predict that you would not want to pay your week's pocket money or salary to play this game. However, if you calculate the EMV you will find:

$$EMV = 2(1/2) + 4(1/4) + 8(1/8) + \ldots + 2^n(1/2^n) + \ldots = 1 + 1 + 1 + \ldots = \infty!$$

Even when faced with potentially large gains, most people do not like to risk a substantial fraction of their financial resources. Although this implies that we cannot always use EMV, it is still possible to give a general analysis of how people make decisions even if they do not like taking risks. As a first step we have to find out the decision-maker's attitude towards risk. A useful concept here is the **certainty equivalent (CE)** of a risky prospect defined as the amount of money which makes the individual indifferent between it and the risky prospect. To clarify this, imagine you are offered a lottery ticket which has a 50–50 chance of winning £0 or £200. So the expected payoff of the lottery is £100. Would you prefer £100 for sure to the lottery ticket? What about £50 for sure? The amount £x so that you are indifferent between *x* and the lottery ticket is your certainty equivalent of the lottery ticket. If your *x* is less than £100, the EMV of the lottery, you are 'risk averse'. In general a decision-maker is **risk averse** if $CE < EMV$, risk **neutral** if $CE = EMV$ and **risk loving** if $CE > EMV$.

Suppose you have an opportunity to invest £1,000 in a business venture which will gross £1,100 or £1,200 with equal probability next year. Alternatively you could deposit the £1,000 in a bank which will give you a riskless return. How large does the interest rate have to be for you to be indifferent between the business venture and the deposit account? (Namely, what is your certainty equivalent? Are you a risk lover?) Note that it is possible to be a risk lover for some lotteries and a risk hater for others.

2.3.1 Expected utility

In the presence of uncertainty, we shall assume that rational agents want to maximise **expected utility**. A utility function represents the level of satisfaction (or benefit or welfare) attached to each possible outcome. It is risk-adjusted, so we can determine a decision-maker's degree of risk appetite summarised in their utility function. The expected utility criterion is convenient because it enables us to still use expected value calculations but with monetary outcomes replaced by utility values. More specifically, the expected utility is the sum of the utility of each possible outcome times the probability of the outcome's occurrence. For example, consider the lottery ticket discussed above and suppose the utility of £200 is 10 and the utility of £0 is 0. In this case the expected utility equals

$0.50\ (10) + 0.50\ (0) = 5.$

It is possible to show that, if a decision-maker satisfies certain relatively plausible axioms, they can be predicted to behave as if they maximise expected utility. Furthermore, since a utility function $U^*(x) = aU(x) + b$, $a > 0$, leads to the same choices as $U(x)$ we can arbitrarily fix the utility of the worst outcome w at $0(U(w) = 0)$ and the utility of the best outcome b at $1(U(b) = 1)$ for a given decision problem. To find the utility corresponding to an outcome x we ask the decision-maker for the value of p, the probability of winning b in a lottery with prizes b and w, which makes x the CE for the lottery. For example, if the worst outcome in a decision problem is £0($U(0) = 0$) and the best outcome is £200($U(200) = 1$), how do we determine $U(40)$? We offer the decision-maker the choice represented in Figure 2.3 and keep varying p until he is indifferent between 40 and the lottery.

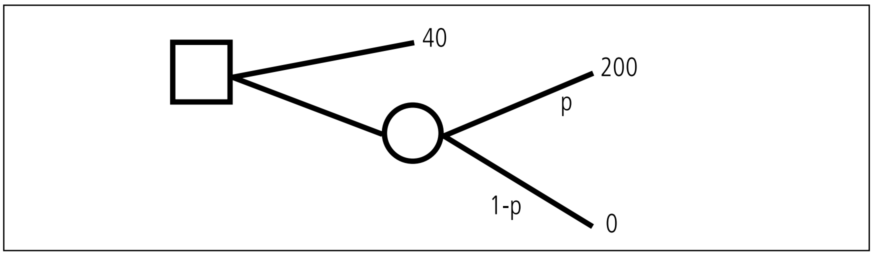

Figure 2.3: Determining utility

When the decision-maker is indifferent, say for p = 0.4, we have:

$$U(40) = U(\text{lottery}) = pU(200) + (1 - p)U(0) = p = 0.4.$$

That is, the utility attached to a gain of £40 is 0.4. Utility values can be obtained in a similar way for the other possible outcomes of the decision problem. Replacing the monetary values by the utility values and proceeding as before will lead to the expected utility maximising decisions.

The definition of risk aversion can be rephrased in terms of the utility function:

- a **risk averse** decision-maker has a concave utility function
- a **risk lover** has a convex utility function
- a **risk neutral** decision-maker has a linear utility function.

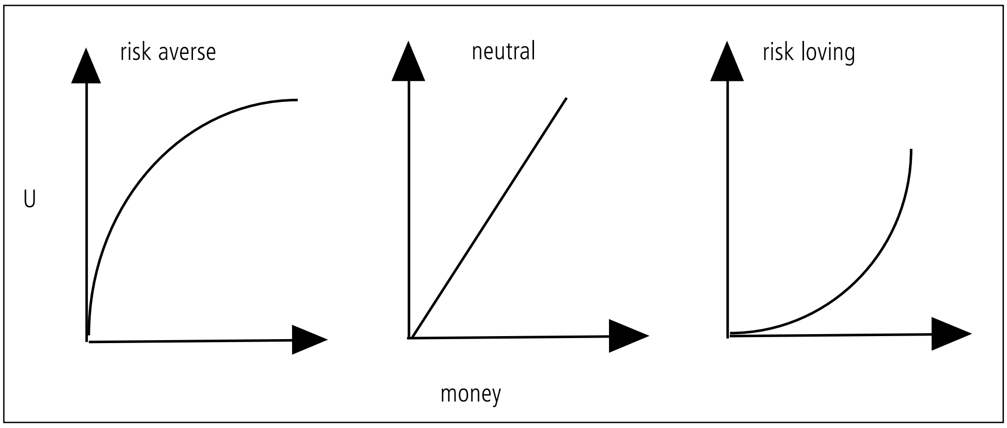

Figure 2.4: Attitudes towards risk

To see why a concave utility function represents risk aversion, consider an individual who will end up with payoffs £100 or £200 which are equally likely so that the expected value of the lottery is £150. For an individual who is risk averse, the CE of this lottery should be less than £150 (you should understand why)! A concave utility function is drawn in Figure 2.5 where you can read off the expected utility of the lottery $EU = 0.5\ U(100) + 0.5\ U(200)$ at point C on the line connecting $U(100)$ and $U(200)$. The CE is the amount of money that delivers the same level of utility as the lottery, i.e. $U(CE) = EU$. For a concave utility function, we always have that $CE<EU$ as can be seen in Figure 2.5 and the individual is therefore risk averse. Following similar steps you should be able to show that $CE>EU$ when the utility function is convex.

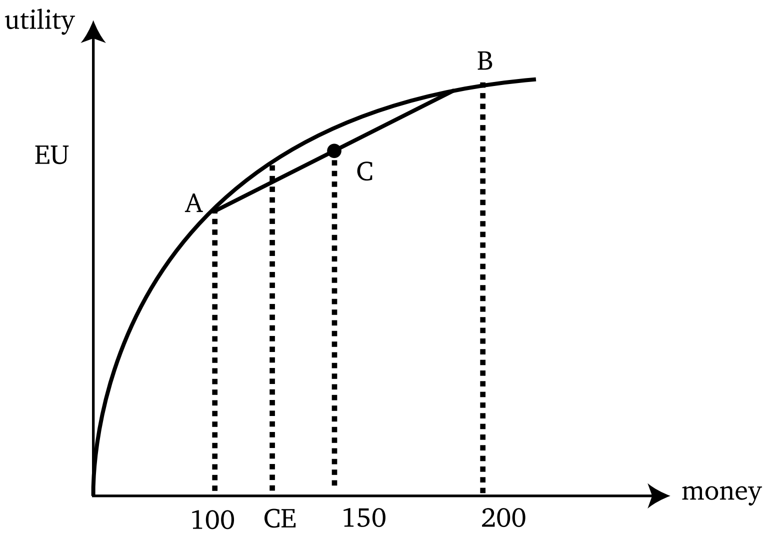

Figure 2.5: Concave utility function and risk aversion

It is possible for a decision-maker to be risk averse over a range of outcomes and risk loving over another range. Indeed, this is how we can explain that the same people who take out home contents insurance buy a national lottery ticket every week. An example of a utility function corresponding to risk loving behaviour for small bets and risk averse behaviour for large bets is drawn in Figure 2.6.

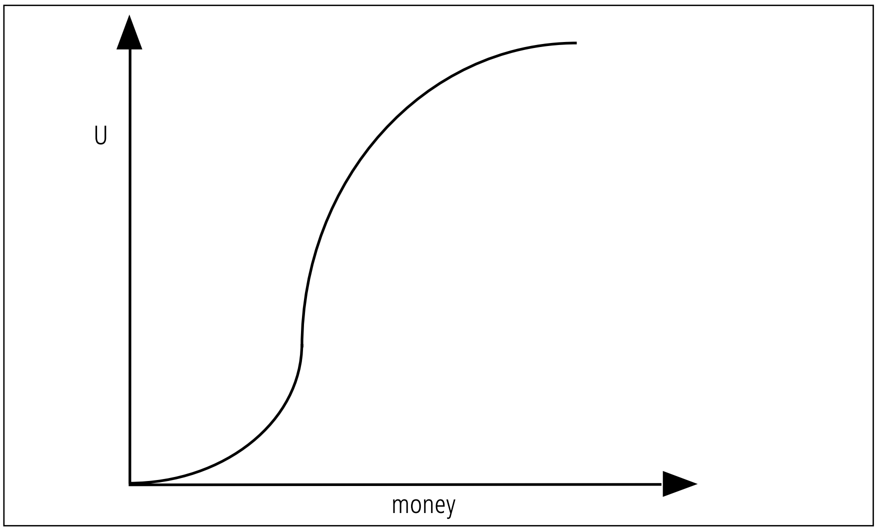

Figure 2.6: Risk loving and risk averse behaviour

2.3.2 Measures of risk aversion

For continuous differentiable functions, there are two measures of risk aversion, the Pratt-Arrow **coefficient of absolute risk aversion** determined as $-U''(x)/U'(x)$ and the Pratt-Arrow **coefficient of relative risk aversion** determined as $-U''(x)x/U'(x)$. If U is concave, then $U''< 0$ and hence both these coefficients are positive for risk averse individuals.[1]

[1] *Note that the coefficient of relative risk aversion is the negative of the elasticity of marginal utility of income and does not depend on the units in which income is measured.*

2.4 Some applications

2.4.1 The demand for insurance

People take out insurance policies because they do not like certain types of risk. Let us see how this fits into the expected utility model. Assume an individual has initial wealth W and will suffer a loss L with probability p. How much would she be willing to pay to insure against this loss? Clearly, the maximum premium R she will pay makes her just indifferent between taking out insurance and not taking out insurance. Without insurance she gets expected utility $EU_0 = pU(W-L) + (1-p)\,U(W)$ and, if she insures at premium R, her utility is $U(W-R)$. Therefore the maximum premium satisfies $U(W-R) = pU(W-L) + (1-p)\,U(W)$. This is illustrated for a risk averse individual in Figure 2.7. The expected utility without insurance is a convex combination of $U(W)$ and $U(W-L)$ and therefore lies on the straight line between $U(W)$ and $U(W-L)$; the exact position is determined by p so that EU_0 can be read off the graph just above $W-pL$ (the EMV of this lottery). This utility level corresponds to a certain prospect $W-R$ which, as can be seen from the figure, has to be less than $W-pL$ (the EMV of this lottery), so that $R > pL$. This shows that, if a risk averse individual is offered actuarially fair insurance (premium R equals expected loss pL), he will insure.

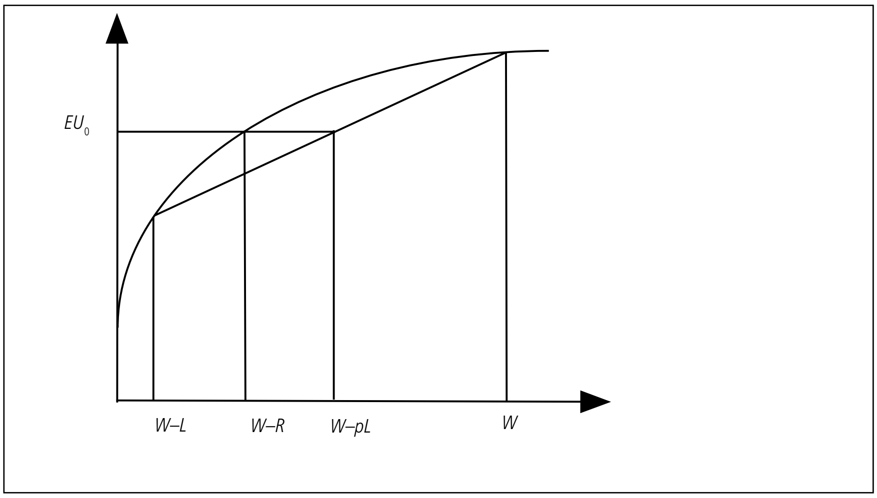

Figure 2.7: Insurance

Example 2.3

Jamie studies at Cambridge University and uses a bicycle to get around. He is worried about having his bike stolen and considers taking out insurance against theft. If the bike gets stolen he would have to replace it which would cost him £200. He finds out that 10 per cent of bicycles in Cambridge are stolen every year. His total savings are £400 and his utility of money function is given by $U(x) = x^{1/2}$. Under what conditions would Jamie take out insurance for a year? What if he has utility of money $U(x) = \ln(x)$?

If he takes out insurance he obtains utility $U(400 - R)$ where R is the premium. Without insurance he gets $(0.1)U(200) + (0.9)U(400)$. Equalising these expected utilities and substituting $U(x) = x^{1/2}$, gives:

$$(0.1)\sqrt{200} + (0.9)\sqrt{400} = \sqrt{400 - R}$$

or

$$R = 23.09$$

which means that insuring is the best decision as long as the premium does not exceed £23.09. Similarly, if $U(x) = \ln(x)$, the maximum premium can be calculated (you should check this!) as £26.79.

2.4.2 The demand for financial assets

Consider the problem of an investor with initial wealth W who wants to decide on her investment plans for the coming year. For simplicity, let us assume that there are only two options: a riskless asset which delivers a gross return of R at the end of the year; and a risky asset which delivers a high return H with probability p and a low return L with probability $1 - p$. It is not difficult to allow for borrowing so that the investor can invest more than W but, to keep things simple, let us restrict the investor's budget to W. The decision problem then consists of finding the optimal amount of money $A(<W)$ to be invested in the risky asset. Given A, the investor gets an expected return of:

$$EU(A) = pU(R(W - A) + HA) + (1 - p)U(R(W - A) + LA) = pU(RW + (H - R)A) + (1 - p)U(RW + (L - R)A)$$

Maximising $EU(A)$ and assuming an interior solution (i.e. $0 < A < W$) leads to the following (first order) condition:

$$EU'(A) = pU'(RW + (H - R)A)(H - R) + (1 - p)U'(RW + (L - R)A)(L - R) = 0$$

Note that, if the risky asset always yields a lower return than the riskless asset ($H, L < R$), there can be no solution to this condition since $U' > 0$. In this scenario the investor would not invest in the risky asset ($A = 0$). Similarly, if the risky asset always yields a higher return than the riskless asset ($H, L > R$) there can be no solution to this condition and under these circumstances the investor would invest all of her wealth in the risky asset ($A = W$). All this makes economic sense. For the other scenarios ($L < R < H$) the first order condition above allows us, for a specific utility function, to calculate the optimal portfolio allocations.

2.4.3 Insurance in financial markets

Consider an asset manager who holds a portfolio of commercial property worth £900 million. They estimate there is a 25 per cent chance the real estate market will worsen next year and the portfolio will be worth only £400 million. They also think there is a 75 per cent chance that property values will remain stable. The expected value of their property portfolio is therefore £775 million.

If the asset manager is risk averse and has utility function $U(W) = W^{1/2}$, where W is the portfolio's worth, the expected utility is given by

$0.25(£400 \text{ million})^{1/2} + 0.75(£900 \text{ million})^{1/2} = 27{,}500$.

The *CE* (that is, the amount of money that would make the manager indifferent between having that amount for certain and holding the risky property portfolio) is calculated as follows:

$U(CE) = CE^{1/2} = 27{,}500$

which means that $CE = £756.25$ million. Suppose a risk neutral insurance company is willing to provide coverage against drops in property values. If the policy covers the manager's entire loss in the portfolio, what premium should the insurer charge for this policy?

The expected payout is £125 million ($0.25 \times £500$m), so the insurance company should charge a premium P of at least this amount to break even on the policy. A higher premium leads to a positive expected profit. The manager's *CE* can in turn be used to calculate the maximum amount she would be willing to pay to hold such a policy. Specifically:

£900 million – $P = CE = £756.25$ million.

Hence the asset manager is indifferent between buying the full coverage insurance policy for £143.75 million and facing the uncertainty in the real estate market because both provide a similar expected utility of 27,500. The difference between the minimum premium that the insurance company is willing to accept (£125 million) and the maximum that the manager is willing to pay (£143.75 million) represents the risk premium, or the amount of money the insured is willing to pay above what is actuarially fair (i.e. the expected value of the loss).

2.5 The expected value of perfect information

In most situations of uncertainty a decision-maker has the possibility of reducing, if not eliminating, the uncertainty regarding some relevant factor by obtaining additional information. Typically this process of finding more information is costly in financial terms (e.g. when a market research agency is contracted to provide details of the potential market for a new product) or in terms of effort (e.g. when you test drive several cars before deciding on a purchase). A crucial question in many decision-making contexts is therefore, 'How much money and/or effort should the decision-maker allocate to reducing the uncertainty?' 'How much should they be willing to pay to buy such information?' In this section, we study a procedure which gives an upper bound (or maximum value) as an answer to this question. This upper bound is called the **expected value of perfect information** (EVPI) and is derived as follows for a risk neutral decision-maker (we will go back to the expected utility model later). The value of perfect information is defined as the increase in expected profit if you can obtain completely accurate information concerning future outcomes.

Assuming you know the outcomes of all probabilistic events in a decision tree, determine the optimal decisions and corresponding pay-off for each possible scenario (combination of outcomes at each probability node). Given that you know the probabilities of each scenario materialising, calculate the expected pay-off under certainty using the optimal pay-off under each scenario and the scenario's probability. This expected pay-off is precisely the pay-off you would expect if you were given exact information about what will happen at each probability node. The difference between this expected pay-off under perfect information and the original optimal pay-off is the EVPI. Since, in reality, it is almost never possible to get perfect information and eliminate the uncertainty completely, the EVPI is an upperbound on how much the decision-maker is willing to pay for **any** (imperfect) information. Generally better decisions are made when there is no uncertainty and therefore the EVPI is positive. However, it is possible that having more information does not change the optimal decisions and in those cases the EVPI is zero. While the concept of EVPI is extremely useful, it is really quite abstract and difficult to grasp. So let us look at an example.

Example 2.2a

The notion of perfect information in Example 2.2 translates into the existence of a perfect seismic test which could tell you with certainty whether there is oil at 3,000 feet, at 5,000 feet or not at all. Assuming such a test exists, how much would the chief executive officer (CEO) be willing to pay to know the test result? The tree in Figure 2.8 represents the (easy) decision problem if all uncertainty is resolved before any decisions are made. There are three scenarios: 'oil at 3,000 feet', 'oil at 5,000 feet' and 'no oil' whose probabilities can be derived from the original tree as 0.40, 0.12, and 0.48 respectively. If the CEO is told which scenario will occur, her decision will be straightforward. Given the optimal decision corresponding to each scenario and the probabilities of the various scenarios, the optimal pay-off with perfect information is £288,000. Recall that the original problem had an EMV of £168,000 and hence EVPI = £120,000. If a perfect seismic test were available, the CEO would be willing to pay up to £120,000 for it.

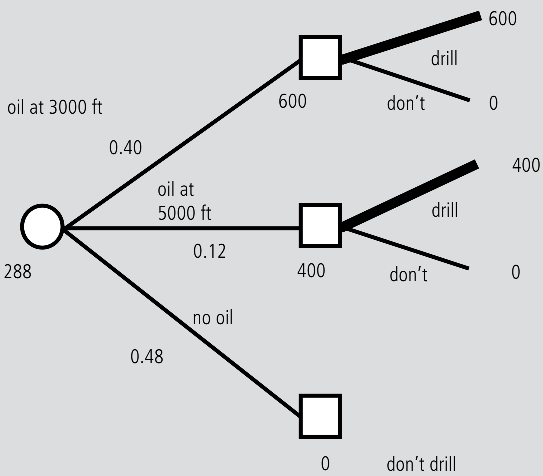

Figure 2.8: Calculating EVPI

If the decision-maker is not risk neutral, a similar method to the one we have just discussed can be used to evaluate the EVPI in utility terms (i.e. we calculate EU under the assumption of perfect information and compare this with the EU in the original problem). However, this does not tell us how much the decision-maker is willing to pay to face the riskless problem rather than the risky one! It is in fact quite tricky to determine the EVPI for an EU maximiser. Consider the simple decision tree in Figure 2.9 where an individual with money utility $U(x) = x^{1/2}$ chooses between a safe and a risky strategy, say investing in a particular stock or not. In either case there are two outcomes – O_1 and O_2 (e.g. the company is targeted for takeover or not) – resulting in the monetary payoffs indicated on the tree. The probabilities of the outcomes are independent of the decision taken. Using the EU criterion, the decision-maker chooses the safe strategy so that EU = 4.

Activity

Explain why the above is the case.

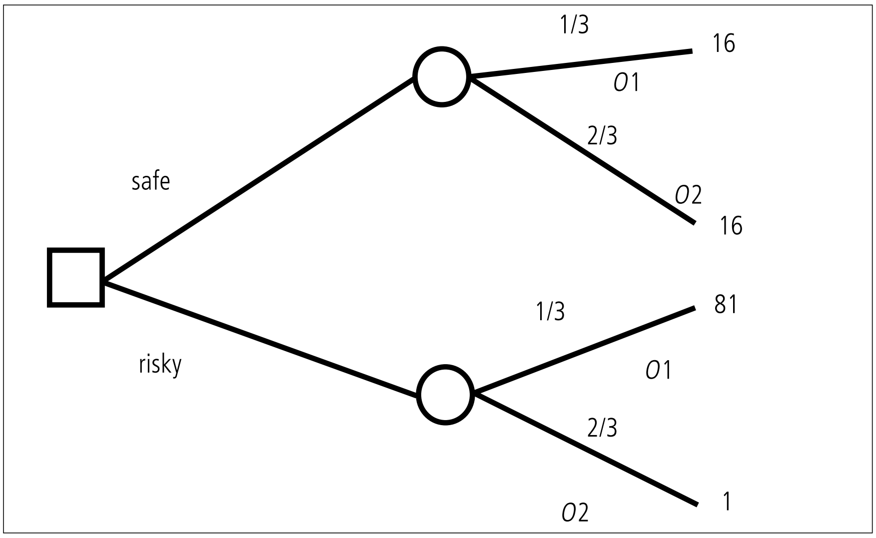

Figure 2.9: EU maximiser chooses safe strategy

With perfect information, however, the decision-maker chooses the 'risky' strategy if O_1 is predicted and the 'safe' strategy when O_2 is predicted, as is indicated in Figure 2.10. This gives her:

$$EU = (1/3)\ (9) + (2/3)\ (4) = 17/3.$$

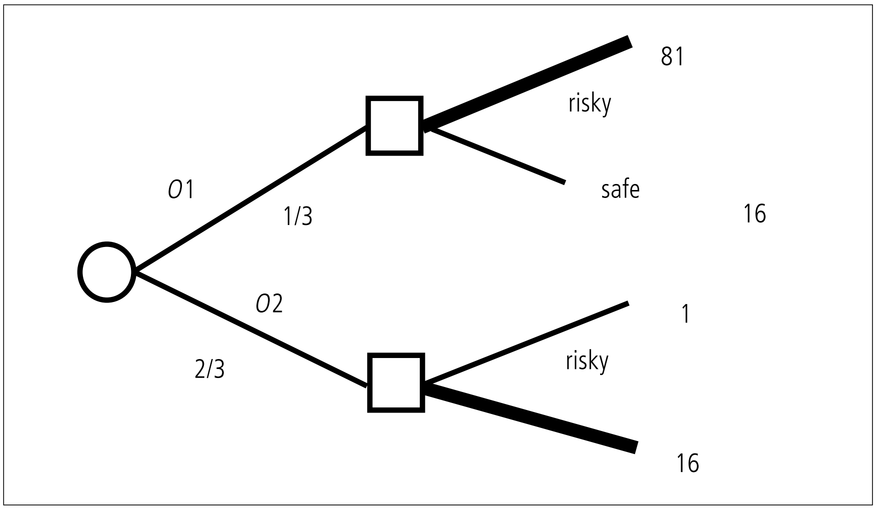

Figure 2.10: Optimal choice under perfect information

How much is the decision-maker willing to pay for this increase in EU from 4 to 17/3? Suppose she pays an amount R for the perfect information. She will be indifferent between getting the information and not getting it if the EU in both cases is equal, or $4 = (1/3)(81 - R)^{1/2} + (2/3)(16 - R)^{1/2}$. Solving this for R (numerically) gives R as approximately 12.5; hence, the EVPI for this problem is about 12.5. Note that this last equation can be solved algebraically.

2.6 Overview of the chapter

This chapter discussed the key concepts of decision analysis, such as: the construction of decision trees and the way they can help identify optimal decisions; the way in which attitudes towards risk may affect decision-making; and common economic measures of risk appetite. A way to value the access to perfect information in a context of uncertainty has also been analysed, as well as applications of decision analysis theory to insurance and financial markets.

2.7 Reminder of learning outcomes

Having completed this chapter, and the Essential reading and exercises, you should be able to:

- structure simple decision problems in **decision tree** format and derive optimal decisions
- explain attitudes towards risk and the way they may impact on optimal decision-making
- calculate **risk aversion coefficients**
- calculate **EVPI** for risk neutral and non-risk neutral decision-makers.
- apply the tools of decision analysis to solve new problems involving uncertainty.

2.8 Test your knowledge and understanding

You should try to answer the following questions to check your understanding of this chapter's learning outcomes.

The VLE also contains additional exercises to test your knowledge of the content of this chapter, together with their respective solutions, as well as past examination papers. We encourage you to use the VLE regularly to support your study of the course material in this subject guide and prepare for the examination.

1. London Underground (LU) is facing a courtcase by legal firm Snook&Co., representing the family of Mr Addams who was killed in the King's Cross fire. LU has estimated the damages it will have to pay if the case goes to court as follows: £1,000,000, £600,000 or £0 with probabilities 0.2, 0.5 and 0.3 respectively. Its legal expenses are estimated at £100,000 in addition to these awards. The alternative to allowing the case to go to court is for LU to enter into out-of-court settlement negotiations. It is uncertain about the amount of money Snook&Co. are prepared to settle for. They may only wish to settle for a high amount (£900,000) or they may be willing to settle for a reasonable amount (£400,000). Each scenario is equally likely. If they are willing to settle for £400,000 they will of course accept an offer of £900,000. On the other hand, if they will only settle for £900,000 they will reject an offer of £400,000. LU, if it decides to enter into negotiations, will offer £400,000 or £900,000 to Snook & Co. who will

either accept (and waive any future right to sue) or reject and take the case to court. The legal cost of pursuing a settlement whether or not one is reached is £50,000. Determine the strategy which minimises LU's expected total cost.

2. Rickie is considering setting up a business in the field of entertainment at children's parties. He estimates that he would earn a gross revenue of £9,000 or £4,000 with a 50–50 chance. His initial wealth is zero. What is the largest value of the cost which would make him start this business:

    a. if his utility of money function is $U(x) = ax + b$ where $a > 0$

    b. if $U(x) = x^{1/2}$; for $x > 0$ and $U(x) = -(-x)^{1/2}$ for $x < 0$

    c. if $U(x) = x^2$, for $x > 0$ and $U(x) = -x^2$ for $x < 0$.

3. Find the coefficient of absolute risk aversion for $U(x) = a - b.exp(-cx)$ and the coefficient of relative risk aversion for $U(x) = a + b.ln(x)$.

4. Find a volunteer (preferably someone who doesn't know expected utility theory) and estimate their utility of money function to predict their choice between the two lotteries in Figure 2.11 below. Payoffs are given in monetary value (£). Check your prediction.

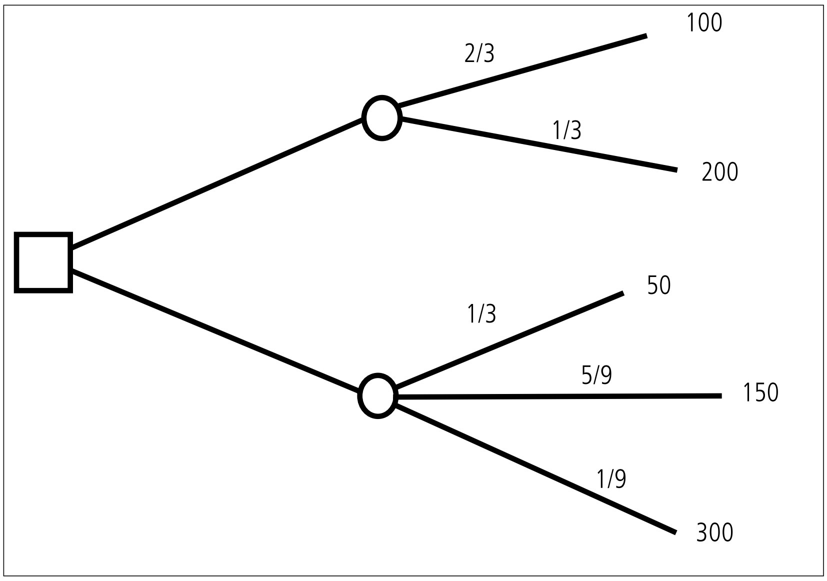

Figure 2.11: Choice between two lotteries

5. An expected utility maximiser spends £10 on a lottery ticket, with a chance of 1 in 1 million of £1 million. He takes out home contents insurance at a premium of £100. His probability of an insurance claim of £1,000 is 1 per cent. Draw his utility of money function.

6. A decision-maker must choose between (1) a sure payment of £200; (2) a gamble with prizes £0, £200, £450 and £1,000 with respective probabilities 0.5, 0.3, 0.1 and 0.1; (3) a gamble with prizes £0, £100, £200 and £520, each with probability 0.25.

    a. Which choice will be made if the decision-maker is risk neutral?

    b. Assume the decision-maker has a CARA (constant absolute risk aversion) utility of money function $U(x) = -a\ exp(-cx) + b$ and her certainty equivalent for a gamble with prizes £1,000 and £0 equally likely is £470. Which choice will be made?[2]

[2] *Hint: show that c has to equal 0.00024 approximately for U(0) = 0 and U(1,000) = 1.*

7. Henrika has utility function $U = M^{1/2}$ for $M \geq 0$ and $U = -(-M)^{1/2}$ for $M < 0$, over money payoffs M.

   a. Given a lottery with outcomes £0 and £36 with respective probabilities 2/3 and 1/3, how much is she willing to pay to replace the lottery with its expected value?

   b. Given the table of money payoffs below, which action maximises her expected utility?

   c. How much would Henrika be willing to pay for perfect information regarding the state of nature?

	S_1	S_2
A_1	4	4
A_2	9	1
	1/3	2/3

Notes

Chapter 3: Game theory

3.1 Introduction

Game theory extends the theory of individual decision-making to situations of strategic interdependence: that is, situations where players (decision-makers) take other players' behaviour into account when making their decisions. The pay-offs resulting from any decision (and possibly random events) are generally dependent on others' actions.

Strategic interaction can involve many players and many strategies, but to simplify our analysis most of the examples and applications discussed in this guide will deal with two-person games with a finite number of strategies. Whereas it is possible to approach many business situations using decision analysis by assigning probabilities to uncertain behaviour of rivals, game theory is a superior methodology for dealing with strategic problems. Rather than assessing probabilities of players' actions, game theory starts out with the assumption that players take actions that promote their self-interest given the information available to them. The introduction of game theoretic reasoning in economics has caused dramatic transformations of several fields of study, most notably – but not exclusively – in industrial organisation.

Game theory can help executives better understand others strategically by providing a robust framework to analyse economic problems and anticipate the behaviour of others (e.g. competitors). This ability, in turn, can increase the payoffs of sound managerial decisions.

A distinction is made between **cooperative game theory** and **non-cooperative game theory**. In cooperative games, coalitions or groups of players are analysed. Players can communicate and make binding agreements. The theory of non-cooperative games assumes that no such agreements are possible. Each player in choosing his or her actions, subject to the rules of the game, is motivated by self-interest. Because of the larger scope for application of non-cooperative games to managerial economics, we will limit our discussion to non-cooperative games.

To model an economic situation as a game involves translating the essential characteristics of the situation into rules of a game. The following must be determined in any game:

- the number of players
- their possible actions at every point in time
- the pay-offs for all possible combinations of moves by the players
- the information structure (what players know when they have to make their decisions).

All this information can be presented in a game tree which is the game theory equivalent of the decision tree. This way of describing the game is called the **extensive form** representation.

It is often convenient to think of players' behaviour in a game in terms of strategies. A strategy tells you what the player will do each time s/he has to make a decision. So, if you know the player's strategy, you can predict his behaviour in all possible scenarios with respect to the other players' behaviour. When you list or describe the strategies available to each player and attach pay-offs to all possible combinations of strategies by the players, the resulting 'summary' of the game is called a **normal form** or **strategic form** representation. A game in normal form will allow us to depict it easily using a pay-off matrix, as will be shown later.

Players in a game can also have complete or incomplete information. In games of **complete information** all players know the rules of the game. In **incomplete information** games, at least one player only has probabilistic information about some elements of the game (e.g. the other players' precise characteristics). An example of the latter category is a game involving an insurer – who only has probabilistic information about the carelessness of an individual who insures his car against theft – and the insured individual who knows how careless he is. A firm is also likely to know more about its own costs than about its competitors' costs. In games of **perfect information** all players know the earlier moves made by themselves and by the other players. In games of **perfect recall**, players remember their own moves and do not forget any information which they obtained in the course of game play. They do not necessarily learn about other players' moves.

Game theory, as decision theory, assumes rational decision-makers. This means that players are assumed to make decisions or choose strategies which will give them the highest possible expected pay-off (or utility). Each player also knows that other players are rational and that they know that the other player knows they are rational and so on. In a strategic situation the question arises whether it could not be in an individual player's interest to convince the other players that one is irrational. (This is a complicated issue which we will consider in the later sections of this chapter. All we want to say for now is that ultimately the creation of an impression of irrationality may be a rational decision.)

Before we start our study of game theory, a 'health warning' may be appropriate. It is not realistic to expect that you will be able to use game theory as a technique for solving real problems. Most realistic situations are too complex to analyse from a game theoretical perspective. Furthermore, game theory does not offer any optimal solutions or solution procedures for most practical problems. However, through a study of game theory, insights can be obtained which would be difficult to obtain in another way and game theoretic modelling helps decision-makers think through all aspects of the strategic problems they are facing. As is true of mathematical models in general, game theory allows you to check intuitive answers for logical consistency.

3.1.1 Aims of the chapter

The aims of this chapter are to consider:

- the concept of information set and why it is not needed in decision analysis
- why it is useful to have both extensive form and normal form representations of a game
- the importance of the prisoners' dilemma as a paradigm for many social interactions
- the concept of dominated strategies and the rationale for eliminating them in the analysis of a game
- the concept of Nash equilibrium (this is absolutely essential!)
- the concepts of backwards induction and subgame-perfect equilibrium
- the concept of non-credible threats and its application to entry deterrence problems.

3.1.2 Learning outcomes

By the end of this chapter, and having completed the Essential reading and exercises, you should be able to:

- represent a simple multi-person decision problem using a **game tree**

- translate from an **extensive form** representation to the **normal form** representation of a game using the **pay-off matrix**
- find **Nash equilibria in pure and mixed strategies**
- explain why in a finitely **repeated prisoners' dilemma** game cheating is a Nash equilibrium
- discuss the **subgame perfect equilibria** and explain the **chainstore paradox**.

3.1.3 Essential reading

Varian, H.R. *Intermediate microeconomics: a modern approach*. (New York: W.W. Norton and Co., 2014; 9th edition) Chapter 29: Game theory and Chapter 30: Game applications.

3.1.4 Further reading

You may also find the following books and articles useful as additional reading:

Allen, W.B., K. Weigelt, N.A. Doherty and E. Mansfield *Managerial economics: theory, applications, and cases*. (New York: W.W. Norton and Co., 2013) 8th edition [ISBN 9780393912777] Chapter 12: Game theory.

Axelrod, R. *The evolution of cooperation*. (New York: Basic Books, 1984) [ISBN 9780465021215].

Kreps, D. and R. Wilson 'Reputation and imperfect information', *Journal of economic theory* (1982) 27, pp.253–79.

Milgrom, P. and J. Roberts 'Predation, reputation and entry deterrence', *Journal of economic theory* (1982) 27, pp.280–312.

Nash, J. 'Noncooperative games', *Annals of Mathematics* 54 1951, pp.289–95.

Selton, R. ' The chain store paradox', *Theory and Decision* (9) 1978, pp.127–59.

Tirole, J. *The theory of industrial organisation*. (Cambridge, MA: The MIT Press, 1988) [ISBN 9780262200714] Chapter 11: Noncooperative game theory: a user's manual.

3.1.5 Synopsis of chapter content

This chapter provides an introduction to non-cooperative game theory by discussing simple games of complete information in normal and extensive form. The concepts of strategic dominance and Nash equilibrium (in pure and mixed strategies) are discussed in the context of insightful applications to managerial economics. We then consider sequential games and discuss the concept of subgame perfect Nash equilibrium as well as some of its applications to industrial organisation.

3.2 Extensive form games

As mentioned above, the extensive form representation of a game is very similar to a decision tree although it can be used for several players rather than just a single decision-maker. The order of play and the possible decisions at each decision point for each player are indicated as well as the information structure, the outcomes or pay-offs and probabilities. As in decision analysis, the pay-offs are not always financial. They may reflect the player's utility of reaching a given outcome. A major difference with decision analysis is that in analysing games and in constructing the game tree, the notion of **information set** is important. When there is only one decision-maker, the decision-maker has perfect knowledge of her own earlier choices. In a game, the players often have to make choices not knowing which decisions have been taken or are taken at the same time by the other players. To indicate that a player does not know her position in the tree exactly, the possible locations are grouped or linked in an information set. Since a player should not be able to deduce from the nature or number of alternative choices available to her where she is in the information set, her set of possible actions has to be identical at every node in the information set. For the same reason, if two nodes are in the same information set, the same player has to make a decision at these nodes. In games of perfect information the players know all the moves made at any stage of the game and therefore all information sets consist of single nodes.

Example 3.1 presents the game tree for a game in which Player 2 can observe the action taken previously by Player 1. However, Example 3.2 presents the game tree for a game in which players take decisions simultaneously and therefore Player 2 does not know the exact decision taken by Player 1. This uncertainty is depicted by Player 2's information set.

Example 3.1

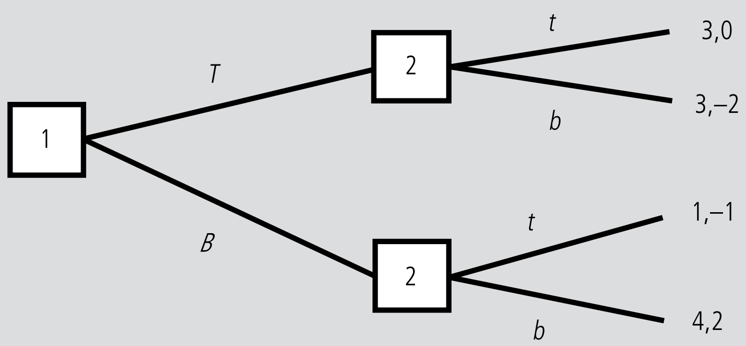

Figure 3.1: Game tree for a dynamic game

In this game tree Player 1 has two strategies: either T (for Top) or B (for Bottom). Player 2 in turn can take actions 't' or 'b' when it's his turn to move.

Player 1 makes the first move; Player 2 observes the choice made by Player 1 (perfect information game) and then chooses from his two alternative actions. The pay-off pairs are listed at the endpoints of the tree. For example, when Player 1 chooses *B* and Player 2 chooses *t*, they receive pay-offs of 1 and –1 respectively. Games of perfect information are easy to analyse. As in decision analysis, we can just start at the end of the tree and work backwards (Kuhn's algorithm). When Player 2 is about to move and he is at the top node, he chooses *t* since this gives him a pay-off of 0 rather than –2 corresponding to *b*. When he is at the bottom node, he gets a pay-off of 2 by choosing *b*. Player 1 knows the game tree and can anticipate these choices of Player 2. He therefore anticipates a pay-off of 3 if he chooses *T* and 4 if he chooses *B*. We can therefore conclude that Player 1 will take action *B* and Player 2 will take action *b*.

Let us use this example to explain what is meant by a strategy. Player 1 has two strategies: *T* and *B*. (Remember that a strategy should state what the player will do in each eventuality.) For Player 2 therefore, each strategy consists of a pair of actions, one to take if he ends up at the top node and one to take if he ends up at the bottom node. Player 2 has four possible strategies, namely:

(*t* if *T*, *t* if *B*)

(*t* if *T*, *b* if *B*)

(*b* if *T*, *t* if *B*)

(*b* if *T*, *b* if *B*)

or {(*t*, *t*), (*t*, *b*), (*b*, *t*),(*b*, *b*)} for short.

Example 3.2

The game tree below is almost the same as in Example 3.1 but here Player 2 does not observe the action taken by Player 1. In other words, it is as if the players have to decide on their actions simultaneously. This can be seen on the game tree by the dashed line linking the two decision nodes of Player 2: Player 2 has an information set consisting of these two nodes. This game (of imperfect information) cannot be solved backwards in the same way as the game of Example 3.1.

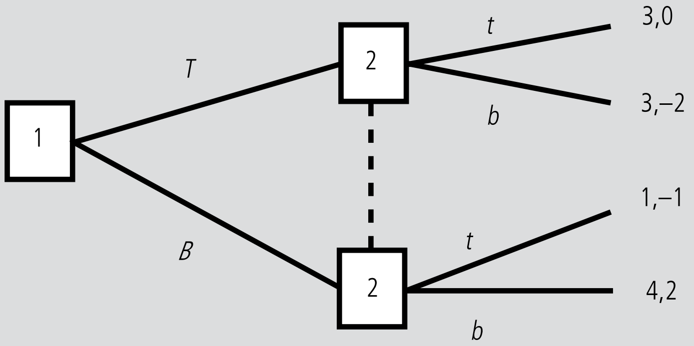

Figure 3.2: Game tree for a simultaneous move game

Note that, although the game trees in the two examples are very similar, Player 2 has different strategy sets in the two games. In the second game his strategy set is just (*t*, *b*) whereas in the first game there are four possible strategies.

3.3 Normal form games

A two-person game in normal form with a finite number of strategies for each player is easy to analyse using a **pay-off matrix**. The pay-off matrix consists of *r* rows and *c* columns where *r* and *c* are the number of strategies for the row and the column players respectively. The matrix elements are pairs of pay-offs (p_r, p_c) resulting from the row player's strategy *r* and the column player's strategy *c*, with the pay-off to the row player listed first. The normal form representations of the games in Examples 3.1 and 3.2 are given below.

		Player 2			
		(t, t)	*(t, b)*	*(b, t)*	*(b, b)*
Player 1	*T*	3,0	3,0	3,–2	3,–2
	B	1,–1	4,2	1,–1	4,2
Normal form for Example 3.1					

		Player 2	
		t	*b*
Player 1	*T*	3,0	3,–2
	B	1,–1	4,2
Normal form for Example 3.2			

Example 3.3

Two competing firms are considering whether to buy television time to advertise their products during the Olympic Games. If only one of them advertises, the other one loses a significant fraction of its sales. The anticipated net revenues for all strategy combinations are given in the table below. We assume that the firms have to make their decisions simultaneously.

		Firm B	
		Advertise	Don't
Firm A	Advertise	10,5	13,2
	Don't	6,7	11,9

If firm *A* decides to advertise, it gets a pay-off of 10 or 13 depending on whether *B* advertises or not. When it doesn't advertise it gets a pay-off of 6 or 11 depending on whether *B* advertises or not. So, irrespective of *B*'s decision, *A* is better off advertising. In other words, 'Don't advertise' is a **dominated strategy** for Firm *A*. Firm *A* has a **dominant strategy**, namely 'Advertise' because this is its optimal choice of strategy no matter what Firm *B* does. Notice that *B* does not have a dominant strategy because its optimal choice depends on what *A* does. Since we are assuming that players behave rationally, dominated strategies can be eliminated. Firm *B* can safely assume that *A* will advertise. Given this fact, *B* now only has to consider the top row of the matrix and hence we can predict that it will also advertise. Note that both *A* and *B* are worse off when both advertise than if they could sign (and enforce) a binding agreement not to advertise.

Example 3.4

In the pay-off matrix below, only one pay-off, the pay-off to the row player, is given for each pair of strategies. This is the convention for zero-sum games where the pay-off to one player is equal to the losses of the other (and therefore the pay-offs to the players sum to zero for all possible strategy combinations). Hence, the entry 10 in the (T, L) position is interpreted as a gain of 10 to the row player and a loss of 10 to the column player. An application of this type of game is where duopolists compete over market share. Then one firm's gain (increase in market share) is by definition the other's loss (decrease in market share). In zero sum games one player (the row player here) tries to maximise his pay-off and the other player (the column player here) tries to minimise the pay-off.

If we consider the row player first, we see that the middle row weakly dominates the bottom row. For a strategy *A* to strictly dominate a strategy *B* we need the pay-offs of *A* to be strictly larger than those of *B* against all of the opponent's strategies. For weak dominance it is sufficient that the pay-offs are at least as large as those of the weakly dominated strategy. In our case, the pay-off of *M* is not always strictly larger than that of *B* (it is the same if Player 2 plays *R*). If we eliminate the (weakly) dominated strategy *B*, we are left with a 2×3 game in which Player 1 has no dominated strategies. If we now consider Player 2, we see that *C* is weakly dominated by *R* (remembering that Player 2's pay-offs are the negative of the values in the table!) and hence we can eliminate the second column. In the resulting 2×2 game, *T* is dominated by *M* and hence we can predict that (M, R) will be played.

		Player 2		
		L	C	R
Player 1	T	10	20	–30
	M	20	10	10
	B	0	–20	10

In general, we may delete dominated strategies from the pay-off matrix and, in the process of deleting one player's dominated strategies, we generate a new pay-off matrix which may contain dominated strategies for the other player which in turn can be deleted and so on. This process is called **successive elimination of dominated strategies**.

Activity

Verify that, in the normal form of Example 3.1 above, this process leads to the outcome we predicted earlier, namely $(B,(t,b))$ but that, in the normal form of Example 3.2 above, there are no dominated strategies.

Sometimes, as in the example above, we will be left with one strategy pair, which would be the predicted outcome of the game but the usual scenario is that only a small fraction of the strategies can be eliminated.

3.4 Nash equilibrium

Game theory can be used to understand and predict behaviour. Much of the theory's predictive power comes from its solution concepts and here we focus on a particular type of solution concept – **Nash equilibrium**.

A Nash equilibrium is a combination of strategies, one for each player, with the property that no player would unilaterally want to change his strategy given that the other players play their Nash equilibrium strategies. So a Nash equilibrium strategy is the best response to the strategies that a player assumes the other players are using.

Pure strategies are the strategies as they are listed in the normal form of a game. The game below has one Nash equilibrium in pure strategies, namely (T, L). This can be seen as follows. If the row player plays his strategy T, the best the column player can do is to play his strategy L which gives him a pay-off of 6 (rather than 2 if he played R). Vice versa, if the column player plays L, the best response of the row player is T. (T, L) is the only pair of strategies which are best responses to each other. It is worth noting that a game may have no Nash equilibrium or more than one equilibrium.

	L	R
T	5, 6	1, 2
B	4, 3	0, 4

In some cases we can find Nash equilibria by successive elimination of **strictly** dominated strategies. However, if weakly dominated strategies are also eliminated we may not find all Nash equilibria of the game. Another danger of successively eliminating **weakly** dominated strategies is that the final normal form (which may contain only one entry) after elimination may depend on the order in which dominated strategies are eliminated.

Example 3.5: 'Battle of the sexes'

The story corresponding to this game is that of a husband and wife who enjoy the pleasure of each other's company but have different tastes in leisure activities. The husband likes to watch football whereas the wife prefers a night out on the town. On a given night the couple have to decide whether they will stay in and watch the football match or go out. The pay-off matrix could look like this.

		Wife	
		In	Out
Husband	In	10, 5	2, 4
	Out	0, 1	4, 8

You should be able to show that this game has two Nash equilibria: (in, in) and (out, out). Only when both players choose the same strategy is it in neither's interest to switch strategies. However, we cannot predict which Nash equilibrium will prevail from the game description alone. The battle of the sexes game is a paradigm for bargaining over common standards.

When electronics manufacturers choose incompatible technologies they are generally worse off than when they can agree on a standard. For example, Japanese, US and European firms were developing their own versions of high definition television whereas they would have received greater pay-offs if they had coordinated. The computer industry, in particular in the area of operating system development, has had its share of battles over standards. This type of game clearly has a first mover advantage and, if firms succeed in making early announcements which commit them to a particular strategy, they will generally do better.

Example 3.6

This example is typical of market entry battles. Suppose two pharmaceutical companies are deciding on developing a drug for combatting Alzheimer's disease or osteoporosis. If they end up developing the same drug, they have to share the market and, since development is very costly, they will make a loss. If they develop different drugs they make monopoly profits which will more than cover the development cost. The pay-off matrix could then look like this:

	A	*O*
A	–2, 2	20, 10
O	10, 20	–1, –1

There are two Nash equilibria in this game: (*A, O*) and (*O, A*). Note that there is a 'first mover advantage' in this game. If Firm 1 can announce that it will develop the drug for Alzheimer's (i.e. strategy *A*) then it can gain 20 if the announcement is believed (and therefore Firm 2 chooses strategy *O*). If Firm 2 could do the same, it would also be better off. Firms in this situation would find it in their interest to give up flexibility strategically by, for example, signing a contract which commits them to the delivery of a certain product. In our scenario a firm could, with a lot of publicity, hire the services of a university research laboratory famous for research on Alzheimer's disease to signal its commitment to the rest of the market.

The type of first mover advantage illustrated in this example is prevalent in the development and marketing of new products with large development costs such as wordprocessing or spreadsheet software packages. The firm which can move fastest can design the most commercially viable product in terms of product attributes and the slower firms will then have to take this product definition as given. Other sources of first mover advantage in a new product introduction context include brand loyalty (first mover retains large market share), lower costs than the second mover due to economies of scale and learning curve effects.

Case study: Airbus versus Boeing

In the early 1990s, Europe's Airbus *Industrie* consortium (Airbus) and Boeing were both capable and committed to developing a large passenger aircraft – a 'superjumbo'. The rationale for pursuing such a project was clear at that time. Airports were getting very crowded and, given the high volume of traffic forecast for the next decades, management predicted that it was going to become increasingly difficult to find take-off and landing slots. Logically an aircraft carrying, say, 700 or 800 passengers rather than 500 was likely to increase efficiency. If the world market has room for only one entrant into the superjumbo segment (predicted sales were about 500) and both firms start development, they will incur severe losses – to bring a superjumbo to market could cost €15bn or more according to press reports. Assume the payoff matrix with strategies 'develop' (D) and 'don't develop' (DD) is similar to the one given below.

A/B	*D*	*DD*
D	–3, –3	10, –1
DD	–1, 10	0, 0

There are two Nash equilibria: one in which Airbus builds the aircraft (and Boeing doesn't) and one in which Boeing builds it (and Airbus doesn't). In this game there is a significant 'first-mover advantage': if we allow Boeing to make a decision before its rival has a chance to make a decision it will develop the aircraft. (In reality the game is, of course, more complicated and there are more strategies available to the players. For example, Boeing could decide to make a larger version of its existing 450-seat 747, which would not be very big but could be developed relatively quickly and at lower cost. Or it could decide to collaborate with Airbus.)

The role of government regulation is not clear-cut here. On the one hand, governments may want to prevent the inefficient outcome of both firms going ahead with development but, on the other hand, the prospect of monopoly (see Chapter 14 of the subject guide) is not attractive either. Particularly if Boeing and Airbus collaborate, the consequences of what would be an effective cartel would be disastrous for struggling airlines. Not only would they have to pay a high price for the super-jumbo but, if they want to buy a smaller aircraft, they would have to turn to Boeing or Airbus who might increase the prices of these smaller, twin-engined jets to promote the super-jumbo.

What happened in reality is that both manufacturers developed their own superjumbos although Boeing, unlike Airbus, did not invest as much. Instead of developing an entirely new aircraft, it developed the 747-8, which, as its name suggests, is an updated version of Boeing's original four-engined 747 jumbo jet. The market for superjumbos turned out to be sluggish and airlines are not placing as many orders as manufacturers (particularly Airbus) initially expected. Part of the reason seems to be flawed strategic business planning. Twin-engined planes improved considerably and became able to fly further and more fuel-efficiently than their four-engined competitors. The main use for the Airbus A380 is on busy routes or into crowded airports like London's Heathrow where landing slots are at a premium, and the only way of flying in more passengers is to use bigger planes. However, at most of the world's airports, landing slots are not necessarily an issue and twin-engined aircrafts are fit-for-purpose. While both companies have order backlogs that will sustain production for a few years, they both face tough choices if new orders do not come through. No passenger airline placed orders for Boeing's 747-8 or Airbus's A380 in 2014.

One way to avoid the problem of both companies starting development would have been for the EU to announce that it would give a large subsidy to Airbus if it developed the aircraft, regardless of whether Boeing also developed it. Then 'develop' may become a dominant strategy for Airbus and one Nash equilibrium would be eliminated. (To see this add 5 to A's pay-off if A chooses strategy D.)

Example 3.7

In the pay-off matrix below there is no Nash equilibrium 'in pure strategies' (i.e. none of the pairs (T, L), (T, R), (B, L) or (B, R) are stable outcomes). Consider, for example, (B, L). If the row player picks strategy B then the best response of the column player is L but, against L, the best response of the row player is T, not B. A similar analysis applies to the other strategy pairs.

	L	R
T	10, 5	2, 10
B	8, 4	4, 2

Is it possible to find an equilibrium concept for games such as the one in Example 3.7? Nash (1951) showed that games in which each player has a finite number of strategies always have an equilibrium. However, players may have to use **mixed strategies** at the equilibrium. A mixed strategy is a rule which attaches a probability to each pure strategy. In other words, a **Nash equilibrium in mixed strategies** is an equilibrium in which each player chooses the optimal frequency with which to play their pure strategies given the equilibrium frequency choices of the others, and no-one has an incentive to deviate.

To see why it makes sense to use mixed strategies, think of the game of poker. The strategies are whether to bluff or not. Clearly, players who always bluff and players who never bluff will do worse than a player who sometimes bluffs. The same applies to tennis players who constantly try to randomise their serve so that opponents cannot easily prepare their return of serve. Players using mixed strategies are less predictable and leaving your opponent guessing may pay off. To see how to find a Nash equilibrium in mixed strategies for a two-player game in which each of the players has two pure strategies, consider the pay-off matrix of Example 3.7 again in 3.7a below.

Example 3.7a

Suppose the row player uses a mixed strategy (x, $1 - x$) (i.e. he plays strategy T with probability x and B with probability $1 - x$) and the column player uses a mixed strategy (y, $1 - y$) (i.e. he plays strategy L with probability y and R with probability $1 - y$). Then the expected pay-offs to the row and column player are respectively:

$$\pi_r = 10xy + 2x(1 - y) + 8(1 - x)y + 4(1 - x)(1 - y)$$

and

$$\pi_c = 5xy + 10x(1 - y) + 4(1 - x)y + 2(1 - x)(1 - y).$$

The row player chooses x so as to maximise her expected pay-off and the column player chooses y so as to maximise his expected pay-off. Given y, the expected pay-off to the row player is increasing in x as long as $y > 1/2$ and decreasing in x for $y < 1/2$ (check this by differentiating π_r with respect to x) and therefore the best response to y is $x = 0$ for $y < 1/2$, $x = 1$ for y>1/2 and any x is optimal against $y = 1/2$. Following a similar derivation for the column player we find that his best response to a given x is to set $y = 0$ for $x > 2/7$, $y = 1$ for $x < 2/7$ and any y is optimal against $x > 2/7$. These best response functions are pictured in bold in Figure 3.3.

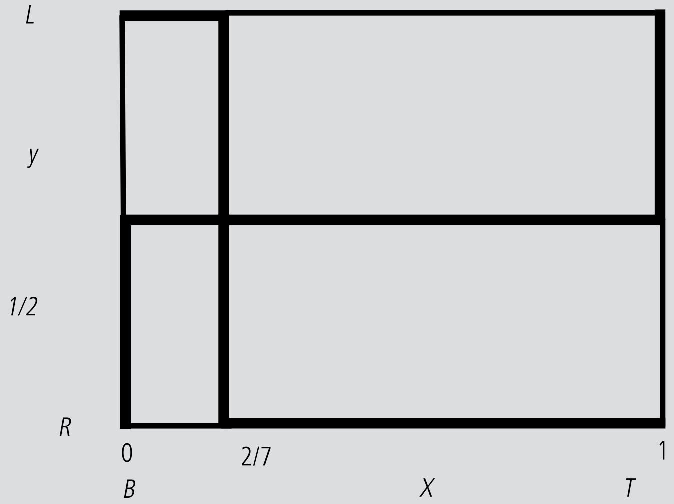

Figure 3.3: Mixed strategy Nash equilibrium

At a Nash equilibrium the strategies have to be best responses to each other. Therefore, the point ($x = 2/7$, $y = 1/2$), where the response functions intersect, forms a Nash equilibrium in mixed strategies. We can then calculate the players' expected pay-offs by substituting x and y in the expressions for π_r and π_c above.

The notion of Nash equilibrium is also very helpful in a negative sense: any combination of strategies which does not form a Nash equilibrium is inherently unstable. However, when there is more than one Nash equilibrium, game theory does not offer a prediction as to which Nash equilibrium is more likely to be played. A topic of substantial research was to try to find justifications for selecting particular types of Nash equilibria – say, the equilibria which are not Pareto dominated – over others. This type of research could help eliminate some Nash equilibria when there are multiple equilibria. Nevertheless game theorists have not succeeded in agreeing on an algorithm which will select **the** equilibrium that is or should be played in reality.

3.5 Prisoners' dilemma

A class of two-person non-cooperative game which has received much attention not only in economics but also in social science more generally, is the class of prisoners' dilemma games. The story the game is meant to model concerns two prisoners who are questioned separately, without the possibility of communicating, about their involvement in a crime. They are offered the following deal. If one prisoner confesses and the other does not, the confessor goes free and the other prisoner serves 10 years; if both confess, they each spend seven years in prison; if neither confesses, they each serve a two-year term. This is summarised in the pay-off matrix below.

	Confess	Don't
Confess	7,7	0,10
Don't	10,0	2,2

This game is very easy to analyse: for both players the strategy 'confess' is a dominant strategy. (Remember that you want to minimise the pay-off here!) There is one Nash equilibrium in pure strategies in which both prisoners serve seven-year sentences. What is interesting about this game is that, if the prisoners could set up a binding agreement, they would agree not to confess and serve only two years. (This type of model is used to explain difficulties encountered in arms control for example, where the strategy 'confess' could be interpreted as 'deploy a new missile'.)

The typical application of the prisoners' dilemma to managerial economics translates the prisoners' plight into the situation of duopolists deciding on their pricing policy. If both set a high price they achieve high profits; if both set a low price they achieve low profits. If one firm sets a low price and its rival sets a high price the discounter captures the whole market and makes very high profits whereas the expensive seller makes a loss. At the Nash equilibrium both firms set low prices and hurt themselves by doing so although in equilibrium, it is in neither firm's interest to deviate.

Of course, in reality firms do not interact only once but they interact in the market over many years and the question arises whether collusive behaviour could be rationalised in a repeated prisoners' dilemma game. When the game is played over several periods rather than as a one-shot game, players might be able to cooperate (set a high price) as long as their rival is willing to cooperate and punish when the rival cheats (deviates from cooperation). This possibility of punishment should give players more of an incentive to cooperate in the long-term. Axelrod (1984) ran a contest in which he asked game theorists to submit a strategy of the repeated version of the prisoners' dilemma. He then paired the given strategies (some of which were very complicated and required significant computer programming) and ran a tournament. The 'tit-for-tat' strategy, which consistently outperformed most of the others, is very simple. This strategy prescribes cooperation in the first round and as long as the other player cooperates but deviation as soon as and as long as the other player deviates from cooperation. In other words, do whatever the other player did in the last round. The tit-for-tat strategy never initiates cheating, and it is forgiving in that it only punishes for one period. If two players use the tit-for-tat strategy, it will be in their own interest to always cooperate.

Let's think about what game theory can contribute to understanding players' behaviour in the repeated prisoners' dilemma. If the game is repeated a finite number of times then collusive behaviour cannot be rationalised. To see this, remember that the only reason to cooperate is to avoid retaliation by your opponent in the future. However, this means that, in the last period, there is no incentive to cooperate. But, if both players are going to cheat in the last period, the next-to-last period can be analysed as if it were the last period and we can expect cheating in that period and so on, so that we end up with the paradox that, even in the repeated prisoners' dilemma game, cheating is the unique equilibrium. (Of course we are assuming, as always, that players are intelligent, can analyse the game and come to this conclusion. If a player is known to be irrational, an optimal response could be to cooperate.)

However, if the prisoner's dilemma is repeated over an infinite horizon or if there is uncertainty about the horizon (i.e. there is a positive probability (< 1) of reaching the horizon), then cooperation can indeed be generated. What is needed is that the strategies are such that the gain from cheating in one period is less than the expected gain from cooperation. For example, both players could use trigger strategies (i.e. cooperate until the other player cheats and then cheat until the horizon is reached). This will be a Nash equilibrium if the gain from cheating for one period is smaller than the expected loss from a switch to both players cheating from then onwards.

3.6 Subgame-perfect equilibrium

So far, with the exception of Section 3.2 'Extensive form games', we have considered games in normal form. In this section we return to the extensive form representation of a game. Consider Example 3.1 and its normal form representation at the beginning of Section 3.3 'Normal form games'. From the pay-off matrix it is clear that there are three Nash equilibria: $(T,(t, t))$, $(B,(t, b))$ and $(B,(b, b))$. Two of these equilibria though, the ones which have Player 2 playing b or t regardless of what Player 1 plays, do not make much sense in this dynamic game. For example, $(B,(b, b))$ implies that Player 2 – if Player 1 plays T – would rather play b and get a pay-off of –2 than t which gives pay-off 0. The reason this strategy is a Nash equilibrium is that Player 1 will not play T.

The notion of subgame-perfect equilibrium was developed as a refinement of Nash equilibrium to weed out this type of unreasonable equilibria. Basically, the requirement for a perfect equilibrium is that the strategies of the players have to form an equilibrium in any subgame. A **subgame** is a game starting at any node (with the exception of nodes which belong to information sets containing two or more nodes) in the game tree such that no node which follows this starting node is in an information set with a node which does not follow the starting node. Put more simply, a subgame is defined as being a segment of a larger game. In the game tree in Figure 3.4, a is the starting node of a subgame but b is not since c, which follows b, is in an information set with d which does not follow b.

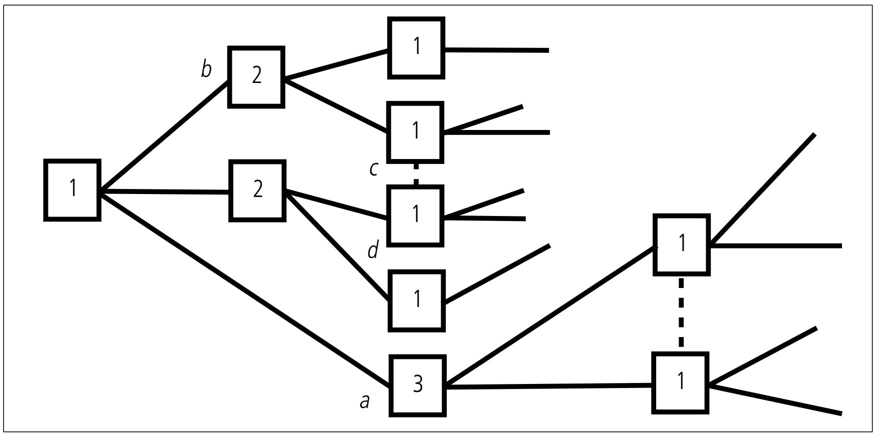

Figure 3.4: *a* starts a subgame, *b* does not

So while $(B,(b, b))$ is a Nash equilibrium in Example 3.1, it is not subgame-perfect since it is not an equilibrium in the (trivial) subgame starting at Player 2's decision node corresponding to Player 1's choice of T. The same applies to Nash equilibrium $(T,(t, t))$.

Example 3.8

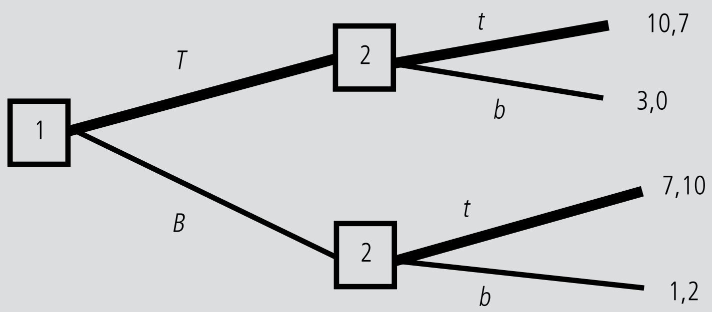

Figure 3.5: Subgame-perfect equilibrium

Consider the game in extensive form depicted in Figure 3.5. Using backwards induction it is easy to see that the subgame-perfect equilibrium is $(T,(t, t))$ as indicated on the game tree. If we analyse this game in the normal form, we find three Nash equilibria (marked with an asterisk in the pay-off table). One of these, namely $(B,(b,t))$ can be interpreted as based on a threat by Player 2 to play b unless Player 1 plays B. Of course if such a threat was credible, Player 1 would play B. However, given the dynamic nature of the game, the threat by Player 2 is not credible since in executing it he would hurt not only his opponent but himself too (he would get a pay-off of 0 rather than 7 which he could get from playing t). By restricting our attention to subgame-perfect equilibria we eliminate Nash equilibria based on non-credible threats.

		Player 2			
		(t,t)	*(t,b)*	*(b,t)*	*(b,b)*
Player 1	*T*	10,7*	10,7*	3,0	3,0
	B	7,10	1,2	7,10*	1,2
Normal form					

Example 3.9 ('entry deterrence')

The industrial organisation literature contains many game theoretic contributions to the analysis of entry deterrence. The simplest scenario is where the industry consists of a monopolist who has to determine his strategy vis-a-vis an entrant. The monopolist can decide to allow entry and share the market with the new entrant, or (threaten to) undercut the new entrant so he cannot make positive profits. The question arises whether the incumbent firm's threat to fight the entrant is credible and deters the entrant from entering. We can analyse this game using the notion of perfect equilibrium. On the game tree in Figure 3.6, E stands for entrant and I for incumbent. The entrant's pay-off is listed first.

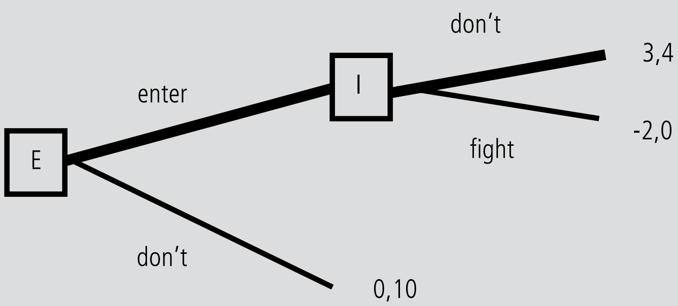

Figure 3.6: Entry deterrence

It is easy to see that, at the subgame-perfect equilibrium, the entrant enters and the incumbent does not fight. It is in the interest of the incumbent firm to accommodate the entrant. There are versions of the entry deterrence game which result in the incumbent fighting entry at equilibrium. In the 'deeper pockets' version, the incumbent has access to more funds since it is likely to be a better risk than the entrant, and can therefore outlast the entrant in a price war. In the incomplete information version, the entrant is not sure about the characteristics or pay-offs of the incumbent (see Example 3.10). In the 'excess capacity' version, incumbent firms make large investments in equipment which affects their pay-off if they decide to fight entry. If firms have a large capacity already in place, their cost of increasing output (to drive the price down) is relatively low and therefore their threat of a price war is credible. To see this in Example 3.9, find the subgame-perfect equilibrium if the incumbent's pay-off of fighting increases to 6 when the entrant enters.

This example could be extended to allow for several potential entrants who move in sequence and can observe whether the incumbent (or incumbents if earlier entry was successful) allows entry or not. It would seem that, for an incumbent faced with repeated entry, it is rational to always undercut entrants in order to build a reputation for toughness and thus deter further entry. Unfortunately, as in the repeated prisoners' dilemma, this intuition fails. Selten (1978) coined the term 'chain store paradox' to capture this phenomenon. The story is about a chain store which has branches in several towns. In each of these towns there is a potential competitor. One after the other of these potential competitors must decide whether to set up business or not. The chain store, if there is entry, decides whether to be cooperative or aggressive. If we consider the last potential entrant, we have the one-shot game discussed above in which the entrant enters and the chain store cooperates. Now consider the next-to-last potential entrant. The chain store knows that being aggressive will not deter the last competitor so the cooperative response is again best. We can go on in this way and conclude that all entrants should enter and the chain store should accommodate them all![1] We should remember that, as for the repeated prisoners' dilemma, the paradox arises because of the finite horizon (finite number of potential entrants). If we assume an infinite horizon, predatory behaviour to establish a reputation for toughness can be an equilibrium strategy.

[1] *For an analysis of different versions of the chain store game in which the paradox is avoided see Kreps and Wilson (1982) and Milgrom and Roberts (1982).*

3.7 Perfect Bayesian equilibrium

In games of imperfect or incomplete information, the subgame-perfect equilibrium concept is not very helpful since there are often no subgames to analyse. Players have information sets containing several nodes.
In these games an appropriate solution concept is perfect Bayesian equilibrium. Consider the imperfect information three-person game in extensive form represented in Figure 3.7. At the time they have to make a move, Players 2 and 3 do not know precisely which moves were made earlier in the game.

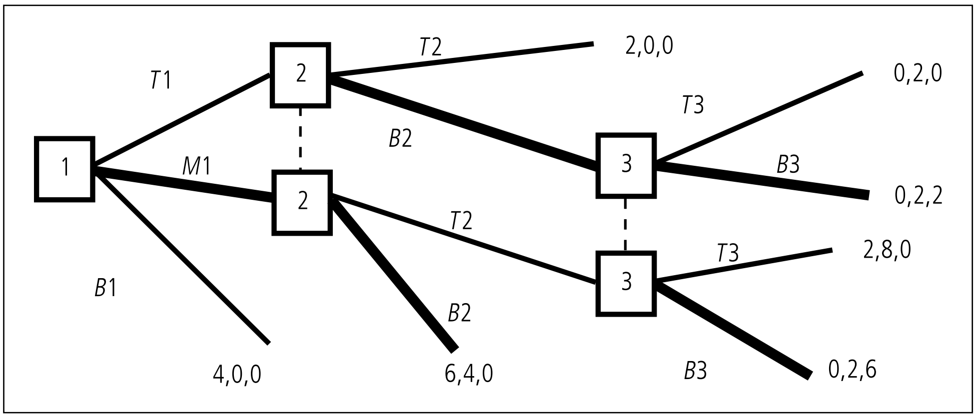

Figure 3.7: Bayesian equilibrium

This game looks very complicated but is in fact easy to analyse if we use backwards induction. When Player 3 gets to move, she has a dominant strategy, *B*3: at each node in her information set, making this choice delivers her the highest pay-off. Hence, whatever probability Player 3 attaches to being at the top or the bottom node in her information set, she should take action *B*3. Similarly, given Player 3's choice, Player 2 chooses *B*2. The choice of *B*2 leads to 2 or 4 depending on whether Player 2 is at the top or bottom node in his information set, whereas *T*2 leads to pay-offs of 0 and 2 respectively. So again, independent of Player 2's probability assessment over the nodes in his information set, he chooses *B*2. Player 1, anticipating the other players' actions, chooses strategy *M*1.

In general for a perfect Bayesian equilibrium we require: (a) that the equilibrium is perfect given players' assessment of the probabilities of being at the various nodes in their information sets; and (b) that these probabilities should be updated using Bayes' rule and according to the equilibrium strategies. In other words, strategies should be optimal given players' beliefs and beliefs should be obtained from strategies. For the game represented in Figure 3.7, these requirements are satisfied if we set the probability of Player 2 being at his bottom node equal to 1 and the probability of Player 3 being at her bottom node equal to any number in [0,1].

Example 3.10

Let us return to the entry game of Example 3.9 and introduce incomplete information by assuming that the incumbent firm could be one of two types – 'crazy' or 'sane' – and that, while it knows its type, the entrant does not. The entrant subjectively estimates the probability that the incumbent is sane as x. This scenario is depicted in the game tree in Figure 3.8 with the pay-offs to the entrant listed first. The important difference with the game of Example 3.9 is that here there is a possibility that the entrant faces a 'crazy' firm which always fights since its pay-off of fighting the entrant (5) is higher than that of not fighting (4). The 'sane' firm, however, always accommodates the entrant. The entrant's decision to enter or not will therefore depend on its belief about the incumbent's type. If the entrant believes that the incumbent firm is 'sane' with probability x then its expected pay-off when it enters is $3x - 2(1 - x) = 5x - 2$. Since the entrant has a pay-off of zero if it doesn't enter, it will decide to enter as long as $5x - 2 > 0$, or $x > 2/5$.

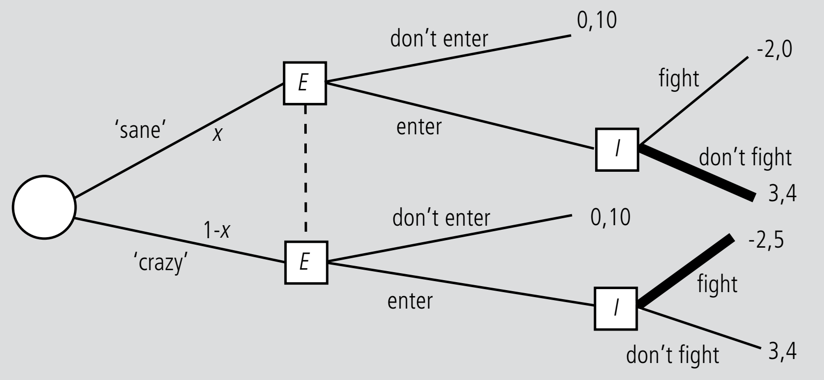

Figure 3.8: Entry deterrence

3.8 Overview of the chapter

In this chapter we considered non-cooperative game theory and introduced the concepts of pay off matrix, dominant strategy and the important solution concepts of Nash equilibrium and subgame-perfect Nash equilibrium. Applications of game theory to simple problems in managerial economics and industrial organisation have also been examined, as well as the concept of perfect Bayesian equilibrium.

3.9 Reminder of learning outcomes

Having completed this chapter, and the Essential reading and exercises, you should be able to:

- represent a simple multi-person decision problem using a **game tree**
- translate from an **extensive form** representation to the **normal form** representation of a game using the **pay-off matrix**
- find **Nash equilibria in pure and mixed strategies**
- explain why in a finitely **repeated prisoners' dilemma** game cheating is a Nash equilibrium
- discuss the **subgame-perfect equilibria** and explain the **chainstore paradox**.

3.10 Test your knowledge and understanding

You should try to answer the following questions to check your understanding of this chapter's learning outcomes.

The VLE also contains additional exercises to test your knowledge of the content of this chapter, together with their respective solutions, as well as past examination papers. We encourage you to use the VLE regularly to support your study of the course material in this subject guide and prepare for the examination.

1. Consider the following matrix game.

	L	R
T	2, 5	1, 4
B	5, −1	3, 1

Are there any dominated strategies? Draw the pay-off region. Find the pure strategy Nash equilibrium and equilibrium pay-offs. Is the Nash equilibrium Pareto efficient? Which strategies would be used if the players could make binding agreements?

2. Find the Nash equilibrium for the following zero sum game. The tabulated pay-offs are the pay-offs to Player 1. Player 2's pay-offs are the negative of Player 1's. How much would you be willing to pay to play this game?

	Player 2			
Player 1	–1	2	–3	0
	0	1	2	3
	–2	–3	4	–1

3. At price 50, quantity demanded is 1,000 annually; at price 60 quantity demanded is 900 annually. There are two firms in the market. Both have constant average costs of 40. Construct a pay-off matrix and find the Nash equilibrium. Assume that, if both firms charge the same price, they divide the market equally but, if one charges a lower price than the other, it captures the whole market. Suppose the two firms agree to collude in the first year and both offer a most favoured customer clause. What is the pay-off matrix for the second year if they colluded the first year?

4. Find the pure and mixed strategy equilibria in the following pay-off tables. How might the –100 pay-off affect the players' actions?

	L	*R*
T	12, 10	4, 4
B	4, 4	9, 6

	L	*R*
T	12, 10	4, 4
B	4, –100	9, 6

5. Students don't enjoy doing homework and teachers don't like grading it. However, it is considered to be in the students' long-term interest that they do their homework. One way to encourage students to do their homework is by continuous assessment (i.e. mark all homework), but this is very costly in terms of the teachers' time and the students do not like it either. Suppose the utility levels of students and teachers are as in the pay-off matrix below.

		teacher	
		check	don't check
student	work	0, -3	0, 0
	no work	-4, 4	1,- 2

a. What is the teacher's optimal strategy? Will the students do any work?

b. Suppose the teacher tells the students at the beginning of the year that all homework will be checked and the students believe her. Will they do the work? Is the teacher likely to stick to this policy?

c. Suppose the teacher could commit to checking the homework part of the time but students will not know exactly when. What is the minimal degree of checking so that students are encouraged to do the work? (Namely, what percentage of homework should be checked?)

6. Consider the two player simultaneous move game below where pay-offs are in £. Find the pure strategy Nash equilibria. How would you play this game if you were the row player? How would you play this game if you were the column player?

	a	*b*	*c*	*d*
u	100, 3	2, 2	2, 1	0, –500
d	0, –500	3, –500	3, 2	1, 10

7. Two neighbours had their house broken into on the same night and from each house an identical rare print went missing. The insurance company with whom both neighbours had taken out home contents insurance assures them of adequate compensation. However, the insurance company does not know the value of the prints and offers the following scheme. Each of the neighbours has to write down the cost of the print, which can be any (integer) value between £10 and £10,000. Denote the value written down by Individual 1 as x_1 and the one written down by Individual 2 as x_2. If $x_1 = x_2$ then the insurance company believes that it is likely that the individuals are telling the truth and so each will be paid x_1. If Individual 1 writes down a larger number than Individual 2, it is assumed that 1 is lying and the lower number is accepted to be the real cost. In this case Individual 1 gets $x_2 - 2$ (he is punished for lying) and Individual 2 gets $x_2 + 2$ (he is rewarded for being honest). What outcome do you expect? What is the Nash equilibrium?

8. Consider the extensive form game below. Find all Nash equilibria. Find the subgame-perfect equilibrium.

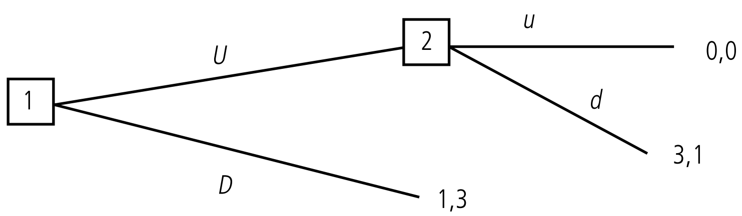

9. Consider the extensive form game below. What are the players' information sets? Write this game in normal form and analyse it using the normal form and the game tree.

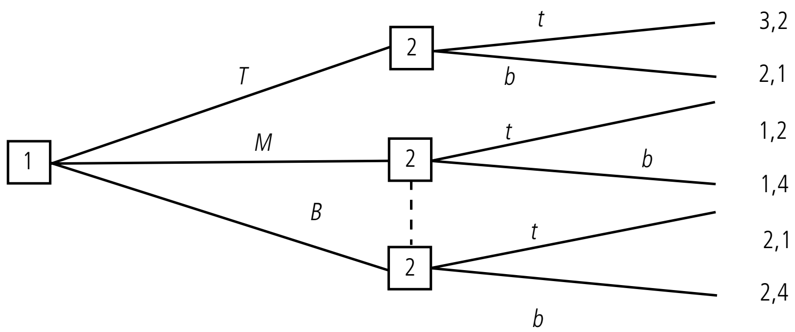

10. A firm may decide to illegally pollute or not. Polluting gives it an extra pay-off of $g > 0$. The Department of the Environment can decide to check for pollution or not. The cost of this inspection is $c > 0$. If the firm has polluted and it is inspected, it has to pay a penalty $p > g$; in this case, the pay-off to the Department of the Environment is $s - c > 0$. If the firm has not polluted and it is checked, no penalty is paid. If the Department of the Environment does not inspect, its pay-off is 0.

a. Suppose that the Department of the Environment can observe whether the firm has polluted or not before it formally decides whether or not to check. Draw the game tree. What is the equilibrium?

b. Now suppose the pollution cannot be observed before checking. Draw the game tree. Is there a pure strategy equilibrium? Compute the mixed strategy equilibrium. How does a change in the penalty affect the equilibrium?

Chapter 4: Bargaining

4.1 Introduction

Consider the following situation. Company B wants to acquire Company S and values it at £1 billion. The current valuation of Company S in the stock market is £700 million, which is the minimum value at which Company S's owners would be willing to sell it. Company B values S more than its current market value because it believes the companies' assets are complementary and by combining them in a single entity, additional value can be created immediately due to strategic synergies. This additional value represents the gains from trade in this example. If the acquisition occurs – that is, if Company B acquires company S – at a price that lies between £700 million and £1 billion, then both companies (i.e. their respective owners) would be better off. The two groups of owners therefore have an interest in completing the deal but, at the same time, they have conflicting interests over the price at which to strike a deal: S would like a high valuation of its assets while B would prefer a low valuation. Any exchange situation, such as the one just described, in which a pair of individuals, groups or organisations can engage in mutually beneficial trade but have conflicting interests over the terms of that trade is a **bargaining situation**.

Generally stated, a bargaining situation is one in which two players have a common interest in cooperating but, at the same time, also have conflicting interests over exactly how to cooperate. The players involved could be individuals, companies or even countries. Many interesting real-life situations (economic, social and political) involve some sort of bargaining over a mutually beneficial outcome. Although many transactions studied in standard microeconomic theory involve markets where at least one side of the market takes the price as given, sale agreements often take place between one individual seller and one buyer. In an industrial context, for instance, this is the case with single sourcing contracts where a firm commits itself to procure all of its supplies of a particular input from a single supplier. In such situations there is no market mechanism to determine the price. Rather than post a price, sellers bargain with buyers over the division of the gains from trade. The number of occasions where negotiation rather than the market mechanism is used is large, and bargaining theory can help us model and analyse outcomes in those environments.

The literature on bargaining has developed dramatically in the last few decades due to advances in non-cooperative game theory. For example, there are multiple applications to political economy, industrial economics, international trade and political science. Bargaining is an interesting topic of study because, as mentioned above, it has both **cooperative** and **conflicting** elements. For example, when a seller has a low reservation price for an object and a buyer has a high reservation price then, clearly, if the two parties can agree to trade, they will both be better off. So both have a strong incentive to reach an agreement. On the other hand, conflict exists regarding the divisions of the gains of trade. Each player would like to reach an agreement that is as favourable to them as possible. The seller will naturally prefer a high price and the buyer will prefer a low price.

Bargaining is any process through which the players **on their own** try to reach an agreement. Typically this process is not frictionless and a basic cost of bargaining comes from the fact that bargaining is time consuming and time is valuable to the players. This creates a sort of urgency to reach an agreement quickly so as to avoid costly delays. Game theory helps us to model bargaining situations carefully and allows us to check our intuition regarding, for example, how the outcome of the bargaining will depend on the parties' bargaining power and strategies, level of impatience, attitudes towards risk and so on. Questions in which economists are interested include:

- conditions under which bargaining will lead to an efficient outcome – the bargaining outcome is Pareto efficient in the sense that parties do reach an agreement without any costly delay;
- the distribution properties of the bargaining outcome – how the gains from trade are divided between the parties;
- what are good bargaining strategies and what determines a player's **bargaining power**?

Bargaining problems arise whenever payoffs have to be shared among several players. When firms succeed in running a cartel, for example, agreeing on how to divide cartel profits is a major problem. Managers are interested in bargaining models for their predictions in the context of management (e.g. for labour (union) negotiations and strikes). However, most game theoretic models of bargaining are either very simplistic (and that is certainly true for the ones we discuss in this chapter) or extremely complex and unrealistic in their assumptions about players' ability to reason and calculate.

In the hope that this does not discourage you too much, let us proceed.

4.1.1 Aims of the chapter

The aims of this chapter are to consider:

- the most prominent axiomatic and game-theoretic **bargaining solutions** used in managerial economics, as well as some of their applications
- the role of **friction** and patience in dynamic bargaining
- the **cooperative** and **conflicting** aspects of bargaining situations
- the nature of the inefficiencies associated with bargaining in the presence of incomplete information.

4.1.2 Learning outcomes

By the end of this chapter, and having completed the Essential reading and exercises, you should be able to:

- apply the **Nash bargaining solution** to economic problems
- analyse **alternating-offers bargaining games** with finite or infinite numbers of rounds
- describe the impact of incomplete information on bargaining outcomes.

4.1.3 Essential reading

Varian, H.R. *Intermediate microeconomics: a modern approach.* (New York: W.W. Norton and Co., 2014); 9th edition, Chapter 30: Game Applications, Section 30.7: Bargaining.

4.1.4 Further reading

Binmore, K. *Playing for real: a text on game theory.* (Oxford: Oxford University Press, 2007) [ISBN 9780195300574] Chapter 17: Cutting a deal.

Muthoo, A. *Bargaining theory with applications.* (Cambridge: Cambridge University Press, 1999) [ISBN 9780521576475] Chapter 2: The Nash bargaining solution; and Chapter 3: The Rubinstein model.

Nash, J.F. Jr. 'The bargaining problem', *Econometrica* 18 1950, pp.155–62.

Rubenstein, A. 'Perfect equilibrium in a bargaining model', *Econometrica* 50 1982, pp.97–109.

4.1.5 Synopsis of chapter content

This chapter examines bargaining models in which players on their own try to reach an agreement over the partition of certain gains from trade (i.e. the 'size of the cake' to be divided). It introduces the Nash bargaining solution to a bargaining problem and then considers Rubinstein's basic alternating-offers model. The chapter discusses some key features of its (unique) subgame-perfect equilibrium as well as the concepts of friction and bargaining power within this model. Finally, the chapter analyses briefly the possible effects of players' incomplete information in the bargaining process and illustrates how inefficiencies may arise in these (more realistic) situations.

4.2 The Nash bargaining solution

A bargaining solution may be interpreted as a recipe that determines a unique outcome for a bargaining situation. The Nash bargaining solution (NBS) is defined by a fairly simple formula. Another attractiveness of the NBS is that it is applicable to a large class of bargaining problems.[1] The NBS is rooted in cooperative game theory. It is therefore axiomatic in that it is focused on predicting outcomes to bargaining problems based on some reasonable properties or axioms. While it does not describe the process by which players arrive at the outcome in terms of individual strategies, it can be shown that it nevertheless possesses some strategic foundations: several plausible (game-theoretic) models of bargaining vindicate their use. Indeed, one such strategic bargaining model that has been widely used in the economics literature will be studied in Section 4.3.

[1] *NBS and the concept of Nash equilibrium studied in Chapter 3 are unrelated, but John Nash created both concepts.*

4.2.1 Bargaining over how to partition a surplus

Here we characterise the NBS of a specific bargaining situation in which two players, A and B, bargain over the partition of a cake (or '**surplus**') of fixed size. This surplus is often also called the **gains from trade** and represents the net value the two players generate when they reach an agreement. For instance, if a buyer and a seller bargain over a mutually acceptable price to exchange an item (as in the example considered in 4.1 Introduction), the gains from trade are given by the difference between the buyer's valuation of the item and the seller's reservation price. The gains from trade can be calculated in a similar fashion depending on the particular problem being analysed.

$U_A(x)$ and $U_B(y)$ are each player's utility from obtaining their respective share of the surplus. For simplicity we assume both utility functions are strictly increasing and concave. Importantly, if the players fail to reach agreement then they obtain their respective **disagreement payoffs** d_A and d_B. These payoffs can be interpreted as a fallback position when bargaining fails. To ensure the existence of a mutually beneficial agreement, let us assume that $U_A(x)>d_A$ and $U_B(y)>d_B$ exist. In other words, there exists a partition of the surplus such that both players are better off than with their disagreement payoffs.

The NBS of the bargaining situation just described is the unique pair of utilities that solves the following maximisation problem:

$$\text{Max } (U_A - d_A)\ (U_B - d_B)$$

where the maximum $(U_A - d_A)\ (U_B - d_B)$, which is the product of the players' gains over their disagreement payoff, is referred to as the Nash product. Under fairly general mathematical conditions, it can be shown that the above problem has a unique NBS which is the solution to the following equation:

$$(U_A(x^N)-d_A)/U'_A(x^N) = (U_B(y^N)-d_B)/U'_B(y^N)$$

where U' denotes the derivative of U, $y^N=G-x^N$ is B's share in the NBS and G is the total surplus.

Example 4.1

Suppose $U_A = x$ and $U_B = y$ while the gains from trade are G. Since in this particular case $U' = 1$, applying the above equation implies that the NBS is characterised by the two payoffs:

$x^N=(G+d_A-d_B)/2$,

$y^N=(G+d_B-d_A)/2$.

This particular bargaining outcome is also known as the split-the-difference rule because these two equations can be rewritten in the following way:

$x^N=d_A+(G-d_A-d_B)/2$,

$y^N=d_B+(G-d_A-d_B)/2$.

The interpretation is as follows. The players agree first of all to give each of them a share d_i of the cake (which gives them a utility equal to the utility they obtain from not reaching agreement), and then they split equally the remaining cake $G-d_A-d_B$. Notice, for example, that player A's share x^N is strictly increasing in d_A and strictly decreasing in d_B.

Example 4.2

This example considers how the degree of risk aversion may affect the players' share of the cake in the NBS. Specifically suppose $U_A = x^\beta$ and $U_B = y$, where $0<\beta<1$. Hence A is risk averse while B is risk neutral. If we also assume that $d_A = d_B = 0$, then the NBS is characterised by the following pair of payoffs:

$x^N = \beta G/(1+\beta)$,

$y^N = G/(1+\beta)$.

Notice that as β decreases, player A's degree of risk aversion increases and her share of the cake shrinks. Player B in turn captures a bigger share of the cake. So a player's share of the cake generally decreases as she becomes relatively more risk averse. At the opposite end of the spectrum, as the value of β approaches 1, the above payoffs converge to a split-the-difference rule as both players share the gains from trade equally.

4.2.2 An application to crime control

In this section we consider a simple application of the NBS concept to crime control in the presence of potential bribery.

A criminal decides whether or not to steal an amount of money $M>0$. If he steals the money, then with probability $\sigma<1$ he is caught by the police who are potentially corruptible. Concretely if the criminal pays a bribe B, the policeman will not report the crime to the authorities and the criminal will walk away free. Let's call the criminal and policeman, respectively, C and P.

Assume the players' utility of money is $U(x)=x$ and the disagreement point $(d_C,d_P)=(\Omega,0)$, where Ω represents the penalty faced by C when he's caught **and** reported to the authorities (the penalty could be monetary or time spent in prison). This bargaining situation is therefore a special case of Example 4.1 above and its NBS is given by:

$U_C=(M + \Omega)/2,$

$U_P=(M - \Omega)/2.$

P's payoff represents the bribe received. So notice that the potential penalty Ω influences both players' payoffs in the NBS even though it is never paid to the authorities.

Under which conditions would crime be successfully deterred? Notice that C's expected utility from stealing is given by:

$EU_C= \sigma (M+\Omega)/2+(1-\sigma)M.$

Since C's utility from not stealing the money is zero, the crime will not be committed so long as:

$\sigma (M+\Omega)/2+(1-\sigma)M<0$; or after rearranging

$M (2 - \sigma)+\sigma\Omega<0.$

Since $\sigma<1$, it can be seen from the above condition that crime is always committed if the penalty matches the crime (i.e. $\Omega = -M$). In effect penalties are not high enough and can be evaded through bribery.

Activity

Show that when the penalty Ω and detection probability σ are sufficiently large, crime can be prevented by the effect they have on C's expected utility and disagreement point at the bargaining stage.

4.3 The alternating-offers bargaining game

We now move on to analyse the more modern non-cooperative approach to modelling the bargaining process. The paradigm is **Rubinstein's model of bargaining** (Rubinstein 1982). In this extensive form game, the players take turns to make offers to each other until agreement is secured. This model has much intuitive appeal, since making offers and counteroffers lies at the heart of many real-life negotiations. One intuitive insight of this model is that some sort of 'friction' is essential to motivate players to reach an agreement. Otherwise there is no meaningful cost of perpetual disagreement. This friction in the bargaining process is typically modelled as cost of delay (or cost of haggling), and the players' respective degree of impatience may in turn influence their bargaining power. Rubinstein's model has been extremely influential and has been adapted and extended over the last few decades to examine a wide variety of new economic applications.

4.3.1 Rubinstein's model

In this model two players have to decide on how to divide an amount of money between them. They alternate in making suggestions about this division and the game ends as soon as an offer is accepted. An offer is a proposal of a partition of the money. If a player accepts the offer, then agreement is struck and the players divide the money according to the accepted offer. If no offer is accepted, the game goes on forever. However, as time goes by, the amount of money available shrinks. This cost of delay represents the 'friction' in the bargaining process alluded to above. It turns out that, at the unique subgame perfect equilibrium of this game, the players agree immediately.

To see how this conclusion is arrived at, let's look at a simpler two-period version of this game. Suppose two players, Bert and Ernie, have a total of £100 to divide between them. Bert makes an offer first. He might decide, for example, to keep £70 to himself and offer £30 to Ernie. Ernie will then either agree to this division or make a counteroffer. If he agrees, he will get the £30; if he refuses the offer and makes a counteroffer, the £100 'cake' will shrink to £100δ. The discount factor δ ($0 < \delta < 1$) can be thought of as a measure of the players' **impatience** or more generally the **cost of a delay** in reaching an agreement (it may represent the cost of a strike in the industrial relations context, for example, where potential output and profit is lost while parties disagree). If the game ends here (i.e. by Bert accepting or rejecting Ernie's counteroffer) and we assume that if no agreement has been reached both players get zero, it is not hard to see what will happen. (You may want to draw a game tree at this point.) Assuming Bert and Ernie are like the usual self-interested, non-altruistic players, then if Ernie decides to make a counteroffer, he will offer Bert one penny, which Bert will certainly accept. The idea is that Ernie can offer some arbitrarily small amount to Bert that would make him better off, and we approximate such an 'arbitrarily small' amount by zero. In this case Ernie ends up with about £100δ. If Bert wants to avoid this outcome, he will have to offer Ernie this same payoff at the start to make him indifferent between accepting his offer in period one or making a counteroffer in the next period. So Bert keeps £100 $(1 – \delta)$ for himself. So at the subgame perfect equilibrium of this two-period game, Bert will offer £100δ to Ernie at the start and Ernie will accept.

Now consider the same game but with Bert willing to make a second offer (i.e. the sequence of moves is Bert, Ernie, Bert). By the time Bert makes his second move, the cake will have shrunk two times to £100δ^2. Now Bert will have the 'last mover advantage' and can make a take-it-or-leave-it offer to Ernie, which he will accept. Ernie will anticipate this and offer Bert £100δ^2 when he has the opportunity, so that he keeps £100δ – £100δ^2 = £100δ $(1 – \delta)$ for himself (remember that in the second period the size of the cake is £100δ). However, Bert can improve on this by offering Ernie £100δ $(1 – \delta)$ in the first move and keeping £100 $(1 – \delta + \delta^2)$ for himself. You should be convinced by now that the subgame perfect equilibrium strategy for each player tells him to make an offer which leaves his opponent very close to being indifferent between accepting the offer and continuing the game. As a consequence, **the first offer is always accepted** and the outcome is efficient.

What happens if we don't give the players a deadline? Suppose they can keep making offers and counteroffers for an infinite number of periods but, as before, the cake shrinks in each period. Since this is an infinite horizon game, you cannot use the backwards induction method.

This artefact also eliminates the 'last mover advantage' and we can look for a symmetric equilibrium (players using the same strategy). Also, the equilibrium strategy must be stationary (i.e. it should give the same prescription in each period because, in the infinite version of the alternating-offers game, when an offer has been rejected, the game is exactly as it was before the offer was made except for the shrinking). So, let us assume that Bert offers Ernie £100x and keeps £100$(1 - x)$ to himself. Ernie will consider accepting £100x or making an offer of £100 δx to Bert and keeping £100 $\delta(1 - x)$. Since he should be indifferent between these two options we find that $\delta(1 - x) = x$ or $x = \delta/(1 + \delta)$. At the (unique) subgame perfect equilibrium Bert gets £100$/(1 + \delta)$ and Ernie gets £100$\delta/(1 + \delta)$. In the infinite version of this game, it is an advantage to be able to make the first move. Again, as in the finite version, the first offer is always accepted and the outcome is **Pareto efficient**. Notice that in contrast to the NBS, in the Rubinstein model, efficiency is a prediction, not an assumption, of the model.

As the time between offers and counteroffers shortens, the discount factor δ approaches 1 and the asymmetry, caused by who moves first, disappears. In this case the equilibrium outcome converges to the split-the-difference rule and each player obtains one-half of the 'cake'. It is not very difficult to extend this analysis to allow for different discount factors for the two bargaining parties. The conclusion of this modified game is that the relatively more patient player will get a larger slice of the cake. Hence in this model a player's bargaining power is inversely related to her degree of impatience. Why does being relatively more patient confer greater bargaining power? Because a player's 'cost' of rejecting an offer is having to wait until next period to be able to make a counteroffer. This intrinsic cost of delay decays as the player becomes more patient.

Another important message of the alternating-offers model is that frictionless bargaining processes are indeterminate. If players do not really care about the time at which agreement is struck, then there is nothing to prevent them from negotiating for as long as they wish. There is no urgency to reach an agreement.

4.3.2 Experimental work

The alternating-offers bargaining game has an elegant and simple 'solution' as well as a stark prediction about the duration of the bargaining (one offer only). However, the game theoretic predictions are not always replicated when tested in experiments with real subjects. Some years ago, a group of students at LSE participated in a series of bargaining experiments. Subjects were paired to an unknown bargaining partner and could only communicate via a formal computer program with their partner who was sitting in a different room. Experiments of this type generally show that 'real' players have a tendency to propose and accept what they consider a fair offer while rejecting what they consider a mean offer even if this rejection means they will be in a worse position. If you were offered one penny in the last round of the game in this section, would you accept?

Disgust and an appetite for fairness are certainly not the only reasons why, in reality, subjects may not behave in accordance with the theory. Some offers may be fully rational but at the same time violate certain social norms of behaviour, which in turn induces retaliation by the other player. Indeed, these social norms may be reflected in significant cross-cultural differences in behaviour that have been discovered by economists when conducting experiments on bargaining games.

4.4 Incomplete information bargaining

While the alternating-offers bargaining game is economically appealing, it nevertheless fails to explain many real life bargaining situations in which disagreement is common and costly negotiations take place over several weeks or months. In the Rubinstein game, agreement is predicted to be immediate. It turns out that to generate delayed agreement you have to assume that the players do not know all the information there is to know about their opponent (i.e. players have some level of private information). In particular, players may have private information about their reservation price (i.e. their maximum willingness to pay) when bargaining over the sale of an item.

Incomplete information bargaining models are classified according to which player(s) has private information and which player(s) makes offers. For example, in a frequently studied bargaining game called the 'Tunisian Bazaar', the value to the seller of the good being sold is common knowledge. The buyer's reservation price, however, is known to the buyer but not the seller, who has certain beliefs about the buyer's reservation price (modelled as a probability distribution). In each period, the seller sets a price and the buyer can either purchase at that price or reject the price offer.

Models of incomplete information bargaining can be extremely complex and I will not discuss them in general or even give an overview. Instead, I will show you in a simple example that inefficiencies can occur. This means that, if it were possible to get both players to reveal their valuations truthfully, they could both be made better off. It is precisely because players are hiding their valuations in order to get an advantage that there are costly delays before an agreement is reached. Consistent with the theory, experiments with subjects have found that when there is complete information about the players' payoffs, bargaining typically leads to quick efficient agreements, but when there is private information, bargaining takes longer and often Pareto efficient outcomes are not reached.

4.4.1 A simple example

Consider the following incomplete information bargaining game. There is a population of buyers and sellers of a good. Half of the seller population has a reservation price of 1 and the other half has a reservation price of 3. One-third of the buyer population has valuation of 2 and two-thirds have valuation of 4. Each player knows his own reservation price or valuation but the others do not.

Two players, a seller and a buyer, are drawn randomly from each respective population and try to come to an agreement about the price at which a good will be sold. We will refer to a player as being of type i if he has reservation price/valuation i. The seller moves first and offers to sell for a price of 2 or 4. The buyer always accepts an offer of 2 but may reject a price of 4. If the buyer rejects, the seller can offer a price of 2 or 4 and the buyer has another chance to accept or reject. If there is delay any pay-offs in the second period are discounted using a discount factor d_s for the seller and d_b for the buyer.

Table 1 contains the possible strategies for each type of player and the pay-offs corresponding to each possible strategy pair. The first pay-off listed is the buyer's. Note that a seller of type 3 will never set the price equal to 2 and that a buyer of type 2 is never willing to buy at price 4. This restricts the number of strategies we have to consider.

	Seller type 1		**Seller type 3**	
	Always ask 2	Ask 4, then 2	Ask 4, then 4	Always ask 4
Buyer type 2 always rejects 4	(0, 2)	$(0, 2d_s)$	(0, 0)	(0, 0)
Buyer type 4 reject 4 once	(2, 2)	$(2d_b, 2d_s)$	$(0, 4d_s)$	$(0, 4d_s)$
Do not reject	(2, 2)	(0, 4)	(0, 4)	(0, 4)

For a buyer of type 4, rejecting a price of 4 in the first period weakly dominates accepting price 4 immediately. If we eliminate the last row in the table then, for a seller of type 1, asking a price of 2 dominates asking 4 first and then asking 2 so that we can eliminate the second column in the table. Given that the seller of type 1 thinks that the buyer he faces is equally likely to be of type 2 as of type 4, asking 2 gives him an expected pay-off of 2 whereas asking 4 twice gives him an expected pay-off of $0(1/3) + (4d_s)(2/3)$. Thus the seller of type 1 asks 2 if $2 > 4d_s(2/3)$ or $d_s < 3/4$. It follows that, if $d_s > 3/4$, there is no trade, at the equilibrium, between a seller of type 1 and a buyer of type 2 which is clearly inefficient. There is also an inefficient delay of the agreement between a seller of type 3 and a buyer of type 4.

4.5 Overview of the chapter

The chapter examined bargaining models in which two players try to reach an agreement over the partition of certain gains from trade. In particular it considered the (axiomatic) Nash bargaining solution and the dynamic Rubinstein's alternating-offers game. It discussed some key features of the subgame perfect equilibrium of the Rubinstein model and the concepts of friction and bargaining power. The chapter also considered briefly the effects of incomplete information on the bargaining process and how inefficiencies may arise in this context.

4.6 Reminder of learning outcomes

Having completed this chapter, and the Essential reading and exercises, you should be able to:

- apply the **Nash bargaining solution** to economic problems
- analyse **alternating-offers bargaining games** with finite or infinite numbers of rounds
- describe the impact of incomplete information on bargaining outcomes.

4.7 Test your knowledge and understanding

You should try to answer the following questions to check your understanding of this chapter's learning outcomes.

The VLE also contains additional exercises to test your knowledge of the content of this chapter, together with their respective solutions, as well as past examination papers. We encourage you to use the VLE regularly to support your study of the course material in this subject guide and prepare for the examination.

1. Consider the alternating-offers bargaining game over £100, with Bert making the first offer and Ernie making a counteroffer if he wants to. Suppose Bert has an outside option of £50, that is, at any point during the game, Bert can stop bargaining and get £50 (discounted if he gets it after the first period). If Bert takes his outside option, Ernie gets zero. How does this affect the equilibrium strategies and payoffs?
2. Consider the following bargaining game. Player 1 offers the division of a cake of size 1 to player 2. Player 2 can either accept or reject. If he accepts, the proposed division is implemented. If player 2 rejects, then player 1 can once more propose a division which player 2 can accept or reject. If player 2 rejects this second offer, he can propose a division to player 1 and player 1 can accept or reject. If player 1 rejects, player 2 can make another offer. If player 1 rejects this second offer, the game ends and both players get zero. Both players discount their payoffs after each rejection with discount factor δ $(0 < \delta < 1)$. What is the subgame perfect equilibrium of this game?