AM30GT: Game Theory

Yordan P. Raykov

Acknoledgements: David Lowe

Preface

This book hopefully will help students to understand the basics of Game Theory.

The material is standard in that the topics covered are traditional game theory topics but non standard to conventional courses in that it requires active participation and learning and a combination of prior skills in statistics, analysis, calculus, logic, linear algebra, optimisation, and humour.

This book's presentation emphasises motivation and active learning driven by a variety of examples and by extensive and careful exercises.

The clearest example of the active developmental approach are the exercises. A student progresses most while doing exercises, so the ones included here have been selected with care to reflect the fundamental concepts. The problems range from simple checks to reasonably involved proofs. The student needs to take ownership of these exercises, by active investigation and participation and hopefully present me with new ones.

Applications and computers. The point of view taken here, that students should think of game theory as an exercise in *thinking*. Computers, especially when coupled with some of the excellent open source mathematical software (for example, take a look at the SAGE project `http://www.sagemath.org/`) should be regarded merely as brain extensions to help us think, in the same way that a pen and paper are extensions allowing us to visualise and assist our thinking. The use of computational approaches is often key to understanding game theory concepts since sometimes game theory can seem counter-intuitive, and counter common sense.

The license. This book will be freely available but it is prepared only for Aston students. It will be downloadable and may be read without restriction.

Any error reports are welcome. The book will be periodically updated to include your fixes and new chapters; people who contribute in this way are acknowledged in the text's source files.

For people reading this book on their own.

This book's emphasis on motivation, active learning and development make it a good choice for self-study. But it can only be a focus. It is important that the good student reads more widely beyond the material presented here.

Game Theory benefits from its recent popularised status and the resulting emergence of many on-line sources. Amongst some that I recommend you might want to investigate to augment your understanding are:

Ben Polak's open Yale course:
`http://oyc.yale.edu/economics/econ-159`

The Stanford Coursera on-line free course:
http://game-theory-class.org/index.html

. . . and especially the excellent collection of material and resources that have been put together in a more anarchic way at:
http://www.gametheory.net/students.html

Giacomo Bonanno's excellent open source book: *GAME THEORY: An open access textbook with 165 solved exercises*. The recent version of the book can be obtained from his web page, http://www.econ.ucdavis.edu/faculty/bonanno/

For the more advanced student, the free book by Ariel Rubinstein and Martin Osborne, MIT Press, which can be downloaded for free from Ariel Rubinstein's web page: http://arielrubinstein.tau.ac.il.

My main advice is: do many exercises, but strive for understanding. I have marked a good sample with 's in the margin. For all of them, you must justify your answer.

Finally, if I may, a caution for all students:- I sometimes hear the sentiment, "I understand the material and the lectures, but it's only that I have trouble with the questions". Simply memorising questions and answers is not satisfying for the curious mind (*"I have no special talents. I am only passionately curious." - A. Einstein*). Being able to do things with ideas is the entire point. The quotes below express this sentiment admirably. They state what is the key to both the beauty and the power of mathematics and the sciences in general.

There is nothing so practical
as a good theory
– Kurt Lewin (1951)
Field theory in social science;
selected theoretical papers.
D. Cartwright (ed.). New York: Harper & Row.

I know of no better tactic
than the illustration of exciting principles
by well-chosen particulars.
–Stephen Jay Gould

If you really wish to learn
then you must mount the machine
and become acquainted with its tricks
by actual trial.
–Wilbur Wright

Dr Yordan P. Raykov
Office: MB313Q
Mathematics Group
Aston University
y.raykov@aston.ac.uk

Maths Symbols

$\Re$, $\Re^+$, $\Re^n$	real numbers, reals greater than 0, n-tuples of reals
	natural numbers: $0, 1, 2, \ldots$
	complex numbers
.....	set of ... such that ...
$(a\,..\,b)$, $[a\,..\,b]$	interval (open or closed) of reals between a and b
...	sequence; like a set but order matters
V, W, U	vector spaces
$\vec{v}, \vec{w}$ or $\boldsymbol{v}, \boldsymbol{w}$	vectors
, $_V$ 0r $\mathbf{0}, \mathbf{0}_V$	zero vector, zero vector of V
B, D	bases
$_n = \boldsymbol{e}_1, \ldots, \boldsymbol{e}_n$	standard basis for $\Re^n$
$\vec{\beta}, \vec{\delta}$ or $\boldsymbol{\beta}, \boldsymbol{\delta}$	basis vectors
$\vec{v}B$	matrix representing the vector
$_n$	set of n-th degree polynomials
$_{nm}$	set of nm matrices
S	span of the set S
MN	direct sum of subspaces
VW	isomorphic spaces
h, g	homomorphisms, linear maps
H, G	matrices
t, s	transformations; maps from a space to itself
T, S	square matrices
hB, D	matrix representing the map h
$h_{i,j}$	matrix entry from row i, column j
Z_{nm}, Z, I_n, I	zero matrix, identity matrix
T	determinant of the matrix T
h, h	rangespace and nullspace of the map h
h, h	generalized rangespace and nullspace
$a \succeq b$	a is "preferred" over b
Ω	A probability space

Lower case Greek alphabet

name	character	name	character	name	character
alpha	α	iota	ι	rho	ρ
beta	β	kappa	κ	sigma	σ
gamma	γ	lambda	λ	tau	τ
delta	δ	mu	μ	upsilon	υ
epsilon	ϵ	nu	ν	phi	ϕ
zeta	ζ	xi	ξ	chi	χ
eta	η	omicron	o	psi	ψ
theta	θ	pi	π	omega	ω

Cover. The cover shows a picture from one of Paul Cézanne's series of paintings on "The Card Players"

Glossary

[This will be developed and expanded as the course develops.]

Backward induction

Backward induction is a technique to solve a game of perfect information. It first considers the moves that are the last in the game, and determines the best move for the player in each case. Then, taking these as given future actions, it proceeds backwards in time, again determining the best move for the respective player, until the beginning of the game is reached.

Common knowledge

A fact is common knowledge if all players know it, and know that they all know it, and so on. The structure of the game is often assumed to be common knowledge among the players.

Dominating strategy

A strategy dominates another strategy of a player if it always gives a better payoff to that player, regardless of what the other players are doing. It weakly dominates the other strategy if it is always at least as good.

Extensive game

An extensive game (or extensive form game) describes with a tree how a game is played. It depicts the order in which players make moves, and the information each player has at each decision point.

Game

A game is a formal description of a strategic situation. Game theory Game theory is the formal study of decision-making where several players must make choices that potentially affect the interests of the other players.

Mixed strategy

A mixed strategy is an active randomization, with given probabilities, that determines the playerâs decision. As a special case, a mixed strategy can be the deterministic choice of one of the given pure strategies.

Nash equilibrium

A Nash equilibrium, also called strategic equilibrium, is a list of strategies, one for each player, which has the property that no player can unilaterally change his strategy and get a better payoff.

Payoff

A payoff is a number, also called utility, that reflects the desirability of an outcome to a player, for whatever reason. When the outcome is random, payoffs are usually weighted with their probabilities. The expected payoff incorporates the playerâs attitude towards risk.

Perfect information

A game has perfect information when at any point in time only one player makes a move, and knows all the actions that have been made until then.

Player

A player is an agent who makes decisions in a game.

Rationality

A player is said to be rational if he seeks to play in a manner which maximizes his own payoff. It is often assumed that the rationality of all players is common knowledge.

Strategic form

A game in strategic form, also called normal form, is a compact representation of a game in which players simultaneously choose their strategies. The resulting payoffs are presented in a table with a cell for each strategy combination.

Strategy

In a game in strategic form, a strategy is one of the given possible actions of a player. In an extensive game, a strategy is a complete plan of choices, one for each decision point of the player.

Zero-sum game

A game is said to be zero-sum if for any outcome, the sum of the payoffs to all players is zero. In a two-player zero-sum game, one playerâs gain is the other playerâs loss, so their interests are diametrically opposed.

Contents

1 Introduction **1**

1.1 An Overview. . . . 1

1.1.1 History and impact of game theory . . . 1

1.1.2 What is Game Theory? . . . 1

1.1.3 Game Descriptions . . . 2

1.1.3.1 Noncooperative and Cooperative Games . . . 3

1.1.3.2 Strategic Games and Extensive Games . . . 3

1.1.3.3 Games with Perfect and Imperfect Information . . . 3

1.1.4 Rational Behavior . . . 3

1.1.4.1 Preferences and Payoff Functions . . . 4

1.1.5 Theory of rational choice . . . 5

2 Strategic Games **6**

2.1 An Overview of Strategic Games. . . . 6

2.1.0.1 Comments on Interpretation . . . 7

2.2 Examples . . . 7

2.2.1 Prisoner's Dilemma . . . 8

2.2.2 Working on a Joint Poster Assessment . . . 9

2.2.3 Duopoly . . . 9

2.2.4 Battle of the Sexes . . . 10

2.2.5 Mathing Pennies . . . 11

2.2.6 The Stag Hunt . . . 11

2.3 Exercises for Week 1 . . . 12

2.3.1 Exercise . . . 12

2.3.2 Exercise . . . 12

2.3.3 Exercise . . . 13

2.3.4 Exercise . . . 14

2.3.5 Exercise . . . 15

2.3.6 Exercise . . . 16

2.3.7 Exercise . . . 17

2.3.8 Exercise . . . 18

3 Two-person Games **20**
3.1 2 person zero-sum games . . . 25
3.1.1 Saddle Points . . . 25
3.2 Exercises for Week 2 . . . 29
3.2.1 Exercise . . . 29
3.2.2 Exercise . . . 31
3.2.3 Exercise . . . 32
3.2.4 Exercise . . . 33
3.2.5 Exercise . . . 34
3.2.6 Exercise . . . 34

4 Mixed Strategies **36**
4.1 Two Person Zero Sum Game (Without a saddle point) . . . 36
4.1.1 Mixed Strategies: Example . . . 36
4.2 Mixed Strategies: Formal Explanation . . . 39
4.2.1 Solution of All 2 by 2 zero sum Matrix Games . . . 43
4.3 Exercises: Mixed Strategies . . . 45
4.3.1 Exercise . . . 45
4.3.2 Exercise . . . 46
4.3.3 Exercise . . . 46
4.3.4 Exercise . . . 47
4.3.5 Exercise . . . 48
4.3.6 Exercise . . . 49
4.3.7 Exercise [O'Neill's card game] . . . 50

5 Dominance **53**
5.1 Dominated Strategies . . . 53
5.2 Exercises on Dominance . . . 59
5.2.1 Exercise . . . 59
5.2.2 Exercise . . . 60
5.2.3 Exercise . . . 62
5.2.4 Exercise . . . 63
5.2.5 Exercise . . . 64
5.2.6 Exercise: Choose a Number . . . 65

6 Extensive Form Games **68**
6.1 Perfect Information Extensive Form Games . . . 69
6.1.1 (Pure) Strategies . . . 71
6.1.2 Extensive Form ⇒ Strategic Form . . . 75
6.1.3 Backward Induction . . . 75
6.2 Exercises: Perfect Information Extensive Form Games . . . 79
6.2.1 Exercise (Entry Game) . . . 79

6.2.2 Exercise: (Service Provider) . . . 81
6.2.3 Exercise: A Stackelberg Duopoly Example . . . 85
6.2.4 Exercise: Contract Negotiation: Team vs Footballer . . . 86
6.2.5 Exercise: Redistribution of Marks . . . 88
6.3 Perfect Information Extensive Form Games: Equilibria . . . 90
6.4 Subgame Perfect Equilibrium . . . 90
6.4.1 Subgame perfect equilibria and Nash equilibria . . . 92
6.4.2 Chess . . . 93
6.5 Exercises for Extensive form Subgame Perfect Equilibria . . . 94
6.5.1 Exercise: (Asymmetric Game) . . . 94
6.5.2 Exercise: (Voting by Alternating Veto) . . . 97
6.5.3 Exercise: (Multiple Optimal Actions) . . . 100
6.5.4 Exercise: (Stackelberg's Model of Duopoly) . . . 102

7 Incomplete Information: Bayesian Games **105**
7.1 Strategic Games . . . 105
7.1.1 Preliminaries . . . 105
7.1.2 Bayes-Nash Equilibrium . . . 107
7.1.2.1 Existence of Bayes–Nash Equilibria . . . 107
7.2 Bayesian Extensive Form Games: (Dynamic Games with Imperfect Information) . . . 109
7.2.1 Information Sets . . . 109
7.2.2 Imperfect Information Extensive Form Games: Illustrations . . . 110
7.2.2.1 Illustration . . . 111
7.3 Exercises for Chapter VII . . . 115
7.3.1 Strategic Games– Incomplete Information . . . 115
7.3.2 (2 Generals Problem –Harsanyi) . . . 119
7.3.3 Extensive form Games– Incomplete Information . . . 122
7.3.3.1 Law and Order
Consider the following game of political intrigue and scandal. . . . 122

8 Auctions **127**
8.1 Example: Second Price Auction (with Complete Information) . . . 127
8.2 Different types of auctions and terminology . . . 128
8.3 Modelling Auctions . . . 129
8.4 Second Price Auctions . . . 130
8.4.1 Expected payment for second price auction . . . 130
8.5 First Price Auctions . . . 131
8.5.1 First Price Auctions: Payments and Revenues . . . 133

8.5.2 Specific Example: Uniform distribution . . . 133
8.6 Revenue Equivalence . . . 134
8.6.1 General Theorem . . . 136
8.7 Scaling Distributions . . . 136
8.8 Exercises for Chapter 8 . . . 138
8.8.1 Exercise: First Price Auction with Reserve Price . . . 138
8.8.2 Exercise: (Nash Equilibrium in a first price sealed bid auction) . . . 139
8.8.3 Exercise: (All-pay auction) . . . 140
8.8.4 Exercise: (Exam Question 2013) . . . 141
8.8.5 Exercise (Random price, sealed-bid auction) . . . 142

Chapter 1

Introduction

1.1 An Overview.

1.1.1 History and impact of game theory

For an extensive historical survey of the development of game theory ideas, see the excellent site of Paul Walker at:
http://www.econ.canterbury.ac.nz/personal_pages/paul_walker/gt/hist.htm

However, perhaps the earliest "modern" example of a formal game-theoretic analysis is the study of a duopoly by Antoine Cournot in 1838. The mathematician Emile Borel suggested a formal theory of games in 1921, which was furthered by the mathematician John von Neumann in 1928 in a "A theory of parlor games". Game theory was established as a field in its own right after the 1944 publication of the monumental volume Theory of Games and Economic Behavior by von Neumann and the economist Oskar Morgenstern. This book provided much of the basic terminology and problem setup that is still in use today. In 1950, John Nash demonstrated that finite games have always have an equilibrium point, at which all players choose actions which are best for them given their opponents' choices. This central concept of noncooperative game theory has been a focal point of analysis since then. In the 1950s and 1960s, game theory was broadened theoretically and applied to problems of war and politics. Since the 1970s, it has driven a revolution in economic theory. Additionally, it has found applications in sociology and psychology, and established links with evolution and biology. Game theory received special attention in 1994 with the awarding of the Nobel prize in economics to Nash, John Harsanyi, and Reinhard Selten. At the end of the 1990s, a high-profile application of game theory has been the design of auctions. Prominent game theorists have been involved in the design of auctions for allocating rights to the use of bands of the electromagnetic spectrum to the mobile telecommunications industry. Most of these auctions were designed with the goal of allocating these resources more efficiently than traditional governmental practices, and additionally raised billions of £s/€s /$s in the UK, Europe and the USA.

1.1.2 What is Game Theory?

Game theory is the formal study of conflict and cooperation. Game theoretic concepts apply whenever the actions of several agents are interdependent. These agents may be individuals, groups, firms, or any

combination of these. The concepts of game theory provide a language to formulate, structure, analyze, and understand strategic scenarios.

As a mathematical tool for the decision-maker the strength of game theory is the methodology it provides for structuring and analyzing problems of strategic choice. The process of formally modelling a situation as a game requires the decision-maker to enumerate explicitly the players and their strategic options, and to consider their preferences and reactions. The discipline involved in constructing such a model already has the potential of providing the decision-maker with a clearer and broader view of the situation. This is a "prescriptive" application of game theory, with the goal of improved strategic decision making.

The object of study in game theory is the game, which is a formal model of an interactive situation. It typically involves several players; a game with only one player is usually called a decision problem. The formal definition lays out the players, their preferences, their information, the strategic actions available to them, and how these influence the outcome. Games can be described formally at various levels of detail. A coalitional (or cooperative) game is a high-level description, specifying only what payoffs each potential group, or coalition, can obtain by the cooperation of its members. What is not made explicit is the process by which the coalition forms. As an example, the players may be several parties in parliament. Each party has a different strength, based upon the number of seats occupied by party members. The game describes which coalitions of parties can form a majority, but does not delineate, for example, the negotiation process through which an agreement to vote en bloc is achieved. Cooperative game theory investigates such coalitional games with respect to the relative amounts of power held by various players, or how a successful coalition should divide its proceeds. This is most naturally applied to situations arising in political science or international relations, where concepts like power are most important. For example, John Nash proposed a solution for the division of gains from agreement in a bargaining problem which depends solely on the relative strengths of the two parties' bargaining position. The amount of power a side has is determined by the usually inefficient outcome that results when negotiations break down. Nash's model fits within the cooperative framework in that it does not delineate a specific timeline of offers and counteroffers, but rather focuses solely on the outcome of the bargaining process. In contrast, noncooperative game theory is concerned with the analysis of strategic choices. The paradigm of noncooperative game theory is that the details of the ordering and timing of players' choices are crucial to determining the outcome of a game. The term "noncooperative" means this branch of game theory explicitly models the process of players making choices out of their own interest. Cooperation can, and often does, arise in noncooperative models of games, when players find it in their own best interests.

Branches of game theory also differ in their assumptions. A central assumption in many variants of game theory is that the players are rational. A rational player is one who always chooses an action which gives the outcome he most prefers, given what he expects his opponents to do. The goal of game-theoretic analysis in these branches, then, is to predict how the game will be played by rational players, or, relatedly, to give advice on how best to play the game against opponents who are rational. We will initially focus on noncooperative game theory with rational players.

1.1.3 Game Descriptions

A game is a description of strategic interaction that includes the con- straints on the actions that the players can take and the players' interests, but does not specify the actions that the players do take. A

solution is a systematic description of the outcomes that may emerge in a family of games. Game theory suggests reasonable solutions for classes of games and examines their properties.

It is convenient to divide Games up into four groups of game theoretic models: Strategic games, Extensive games with and without perfect information, and Coalition games.

1.1.3.1 Noncooperative and Cooperative Games

In all game theoretic models the basic entity is a player. A player may be interpreted as an individual or as a group of individuals making a decision. Once we define the set of players, we may distinguish between two types of models: those in which the sets of possible actions of individual players are primitives and those in which the sets of possible joint actions of groups of players are primitives (coalition games). Sometimes models of the first type are referred to as "noncooperative", while those of the second type are referred to as "cooperative" (though these terms do not express well the differences between the models).

1.1.3.2 Strategic Games and Extensive Games

A strategic game is a model of a situation in which each player chooses his plan of action once and for all, and all players' decisions are made simultaneously (that is, when choosing a plan of action each player is not informed of the plan of action chosen by any other player). By contrast, the model of an extensive game specifies the possible orders of events; each player can consider his plan of action not only at the beginning of the game but also whenever he has to make a decision.

The strategic form (also known as the *normal form*) is the basic type of game studied in non-cooperative game theory, and the one we will focus on shortly. A game in strategic form lists each player's strategies, and the outcomes that result from each possible combination of choices. An outcome is represented by a separate payoff for each player, which is a number (also called utility) that measures how much the player likes the outcome. The extensive form, (also known as a game tree), is more detailed than the strategic form of a game. It is a complete description of how the game is played over time. This includes the order in which players take actions, the information that players have at the time they must take those actions, and the times at which any uncertainty in the situation is resolved. A game in extensive form may be analyzed directly, or can be converted into an equivalent strategic form. Later examples will illustrate in detail the interpretation and analysis of games in strategic and extensive form.

1.1.3.3 Games with Perfect and Imperfect Information

The third distinction that we make is between the models in which the participants are fully informed about each others' moves, or they may be imperfectly informed. The former models have firmer foundations. The latter were developed intensively only in the 1980s.

1.1.4 Rational Behavior

The models we study assume that each decision-maker is "rational" in the sense that she is aware of her alternatives, forms expectations about any unknowns, has clear preferences, and chooses her action deliberately after some process of optimisation.

1.1.4.1 Preferences and Payoff Functions

We assume that a decision-maker, when faced with a choice of two alternatives, knows which one of the two she prefers, or is indifferent between the two, and she is consistent in this choice. Therefore, if she prefers choice a over choice b and choice b over choice c then she prefers choice a over choice c. These choices could be altruistic (a person's preference could also include her desire for the welfare of others for example, rather than a selfish attribution). We can represent these preferences by *payoff functions* which associate a number with each choice in such a way that choices with higher numbers are preferred. Sometimes a choice a will have a consequence $g(a)$ and the decision-makers preference will be on the consequence, rather than the choice. We will sometimes suppress the consequence function $g(a)$ and use simply the choice a itself when its meaning is clear.

Later we will discuss the decision-maker's preferences specified by giving a utility function $U : C \to \Re$, which defines a preference relation by the condition $x \succeq y$ if and only if $U(x) \geq U(y)$. Given any set $B \subseteq A$ of actions that are feasible in some particular case, a rational decision-maker chooses an action a^* that is feasible (belongs to B) and optimal in the sense that $g(a^*) \succeq g(a)$ for all $a \in B$; alternatively she solves the problem $max_{a \in B} U(g(a))$. An assumption upon which the usefulness of this model of decision-making depends is that the individual uses the same preference relation when choosing from different sets B.

In the absence of uncertainty the following elements constitute a model of rational choice

Definition 1. *A model of rational choice consists of:*

- *A set A of actions from which the decision-maker makes a choice.*
- *A set C of possible consequences of these actions.*
- *A consequence function $g : A \to C$ that associates a consequence with each action.*
- *A preference relation $\succeq$ (a complete, transitive, reflexive, binary relation) on the set C.*

In the models we study later, individuals often have to make decisions under conditions of uncertainty. The players may be

- uncertain about the objective parameters of the environment,
- imperfectly informed about events that happen in the game,
- uncertain about actions of the other players that are not deterministic,
- uncertain about the reasoning of the other players.

Note that these preference functions are *ordinal* in the sense that they only convey information about the *order* of choices, and not by how much one is preferred over another. ie the preference function only tells us that a is preferred over b (ie $a \succeq b$) if $u(a) > u(b)$, and not that a is preferred "twice as much" as b for example, even if $u(a) = 2u(b)$. This is known as an ordinal relationship. Therefore a decision-maker's preferences can be represented by many different preference functions provided it retains the same ordering over choices.

1.1.5 Theory of rational choice

In any given situation the decision-maker chooses the member of the available subset of all possible choices A that are best according to her preferences.

Theorem 1 (Theory of Rational Choice)**.** *The choice selected by a decision-maker is at least as good, according to her preferences, as every other possible available choice.*

This Theorem is enormously successful in Game Theory and also pervades much of economic theory.

We will now go on to consider **Strategic Games**, how to set up examples as "matrix games" and how to seek equilibrium (i.e. steady state) solutions.

Chapter 2

Strategic Games

2.1 An Overview of Strategic Games.

A strategic game is a model of interactive decision-making in which each decision-maker chooses her plan of action once and for all, and these choices are made simultaneously. The model consists of a finite set N of players and, for each player i, a set A_i of actions and a preference relation on the set of action profiles. We refer to an action profile $a = (a_j)_{j \in N}$ as an outcome, and denote the set $\times_{j \in N} A_j$ of outcomes by A. The requirement that the preferences of each player i be defined over A, rather than A_i , is the feature that distinguishes a strategic game from a decision problem: each player may care not only about her own action but also about the actions taken by the other players. To summarise, our definition is the following.

Definition 2 (A strategic game)**.** *A strategic game (with ordinal preferences) consists of:*

- *a finite set N (the set of players)*
- *for each player $i \in N$ a nonempty set A_i (the set of actions available to player i)*
- *for each player $i \in N$ a preference relation $succeq_i$ on $A = \times_{j \in N} A_j$ (the preference relation of player).*

If the set A_i of actions of every player i is finite then the game is finite.

A wide variety of situations may be modelled as strategic games. Time is absent from the model. The idea is that the player makes her choice once and for all, all players choose simultaneously in the sense that no player is informed of the other player's choice before they make their own choice. An action may be interpreted as a "strategy" in this context. We abstract all complications that may arise from a player deciding to alter their choices once the game is underway. We look at this case later.

The high level of abstraction of this model allows it to be applied to a wide variety of situations. A player may be an individual human being or any other decision-making entity like a government, a board of directors, the leadership of a revolutionary movement, or even a flower or an animal. The model places no restrictions on the set of actions available to a player, which may, for example, contain just a few elements or be a huge set containing complicated plans that cover a variety of contingencies. However, the range of application of the model is limited by the requirement that we associate with each player a preference

	L	R
T	w_1, w_2	x_1, x_2
B	y_1, y_2	z_1, z_2

Figure 2.1: A general 2-person payoff matrix representing a strategic game

relation. A player's preference relation may simply reflect the player's feelings about the possible outcomes or, in the case of an organism that does not act consciously, the chances of its reproductive success. The fact that the model is so abstract is a merit to the extent that it allows applications in a wide range of situations, but is a drawback to the extent that the implications of the model cannot depend on any specific features of a situation.

A convenient representation of a two-player strategic game in which each player has two actions is a so-called *matrix game* that is shown in Figure 2.1. Under a wide range of circumstances the preference relation $\succeq_i$ of player i in a strategic game can be represented by a payoff function $u_i : A \to \Re$ (which are also the same as a utility function).

A finite strategic game in which there are two players can be described conveniently in a table like that in Figure 2.1. One player's actions are identified with the rows and the other player's with the columns. The two numbers in the box formed by row r and column c are the players' payoffs when the row player chooses r and the column player chooses c, the first component being the payoff of the row player. Thus in the game in Figure 2.1 the set of actions of the row player is $\{T, B\}$ and that of the column player is $\{L, R\}$, and for example the row player's payoff from the outcome (T, L) is w_1 and the column player's payoff is w_2. If the players' names are "1" and "2" then the convention is that the row player is player '1' and the column player is player '2'.

2.1.0.1 Comments on Interpretation

A common interpretation of a strategic game is that it is a model of an event that occurs only once; each player knows the details of the game and the fact that all the players are "rational" and the players choose their actions simultaneously and independently. Under this interpretation each player is unaware, when choosing her action, of the choices being made by the other players; there is no information (except the primitives of the model) on which a player can base her expectation of the other players' behavior. Another interpretation, is that a player can form her expectation of the other players' behavior on the basis of information about the way that the game or a similar game was played in the past. A sequence of plays of the game can be modelled by a strategic game only if there are no strategic links between the plays. That is, an individual who plays the game many times must be concerned only with her instantaneous payoff and ignore the effects of her current action on the other players' future behavior. For a situation to be modelled as a strategic game it is important only that the players make decisions independently, no player being informed of the choice of any other player prior to making her own decision.

2.2 Examples

The following classical games represent a variety of strategic situations. The games are very simple: in each game there are just two players and each player has only two possible actions. Nevertheless, each

game captures the essence of a type of strategic interaction that is frequently present in more complex situations.

2.2.1 Prisoner's Dilemma

Perhaps the most well known strategic game is the *Prisoner's Dilemma*.

The Prisoner's Dilemma is a game in strategic form between two players. Each player has two strategies, called "cooperate" (with their accomplice and stay Quiet - so we use Q to denote this choice) and "defect" (rat on their accomplice so we denote this by R).

Player I chooses a row, either Q or R, and simultaneously player II chooses one of the columns Q or R. The strategy combination (Q, Q) has payoff 2 for each player, and the combination (R, R) gives each player payoff 1. The combination (Q, R) results in payoff 0 for player I and 3 for player II, and when (R, Q) is played, player I gets 3 and player II gets 0.

Any two-player game in strategic form can be described by a table like the one in Figure 2.2, with rows representing the strategies of player I and columns those of player II. (A player may have more than two strategies). Each strategy combination defines a payoff pair, like $(3, 0)$ for (R, Q), which is given in the respective table entry. Each cell of the table shows the payoff to player I at the left, and the payoff to player II at the right. The game is symmetric between the two players. Symmetry means that the game stays the same when the players are exchanged, corresponding to a reflection along the diagonal. Note that in the strategic form, there is no order between player I and II since they act simultaneously (that is, without knowing the other's action), which makes the symmetry possible.

In the Prisoner's Dilemma game, "Rat" is a strategy that dominates Quiet Strategy R of player I dominates Q since if player II chooses Q, then player I's payoff is 3 when choosing R and 2 when choosing Q; if player II chooses R, then player I receives 1 for R as opposed to 0 for Q. Hence, R is indeed always better and dominates Q. In the same way, strategy R dominates Q for player II.

No rational player will choose a 'weak', or dominated strategy since the player will always be better off when changing to the strategy that dominates it. The unique outcome in this game, as recommended to utility-maximizing players, is therefore (R, R) with payoffs $(1, 1)$. Somewhat paradoxically, this is less than the payoff $(2, 2)$ that would be achieved when the players chose (Q, Q).

The story behind the name "Prisoner's Dilemma" is that of two prisoners held suspect of a serious crime. There is no judicial evidence for this crime unless one of the prisoners testifies against the other. If one of them testifies, he will be rewarded with immunity from prosecution (payoff 3), whereas the other will serve a long prison sentence (payoff 0). If both testify, their punishment will be less severe (payoff 1 for each). However, if they both cooperate with each other (and stasy Quiet) by not testifying at all, they will only be imprisoned briefly, for example for illegal weapons possession (payoff 2 for each). The "defection" from that mutually beneficial outcome is to testify, which gives a higher payoff no matter what the other prisoner does, with a resulting lower payoff to both. This constitutes their "dilemma".

Prisoner's Dilemma games arise in various contexts where individual "defections" at the expense of others lead to overall less desirable outcomes. Examples include arms races, litigation instead of settlement, environmental pollution, or cut-price marketing, where the resulting outcome is detrimental for the players. Its game-theoretic justification on individual grounds is sometimes taken as a case for treaties and laws, which enforce cooperation.

Game theorists have tried to tackle the obvious "inefficiency of the outcome of the Prisoner's Dilemma game. For example, the game is fundamentally changed by playing it more than once. In such a repeated game, patterns of cooperation can be established as rational behavior when players fear of punishment in the future outweighs their gain from defecting today.

	Q	R
Q	$2, 2$	$0, 3$
R	$3, 0$	$1, 1$

Figure 2.2: Payoff matrix for Prisoner's Dilemma

2.2.2 Working on a Joint Poster Assessment

You are working together with a friend on a joint poster for assessment. Each of you can either work hard, or skive off. If your friend works hard then you prefer to skive off since the poster will be good enough for a decent mark and you will not have had to do any work. So you have a 'quality of life' payoff of 3. But if you also work hard on it, the poster will be better, though the incremental value of the poster due to you working hard is not worth all the effort. So your quality of life payoff goes down to 2. If you both skive off the poster does not get done and you get a bad joint mark, so you have a payoff of 1. So you prefer the option of you both working hard to you both skiving. But the worst outcome for you personally (even worse than getting a bad mark!) is if you work hard and your friend skives off, since you do not like being exploited. In this case your quality of life payoff is 0.

Your friend has exactly the same preferences, so the payoff matrix can be produced as shown in Figure 2.3.

	Work Hard	Skive Off
Work Hard	$2, 2$	$0, 3$
Skive Off	$3, 0$	$1, 1$

Figure 2.3: Payoff matrix for the Joint Poster Project problem

This is of the same form as the Prisoner's Dilema and has the same unsatisfactory, and anti-common sense result that you both choose to skive off giving a worse overall payoff than if you had jointly worked hard (total 'quality of life' value of 4 versus 2)!

2.2.3 Duopoly

In a simple model of a duopoly, two firms produce the same goods, for which each firm can choose to charge either a low price or a high price. Each firm wants to achieve the highest possible profit. If both firms choose "High" then each earns a proifit of £1000. If one firm chooses "High" and the other firm chooses "Low", then the firm that chooses "High" gets no customers and makes a loss of £200 and the other firm earns a profit of £1200 (its volume if high but its unit cost at sale is low). If both firms choose "Low" then each earns a profit of £600. Since each firm is concerned only about its profit, we can represent their payoff preferences directly in terms of profit as shown in Figure 2.4.

	High	Low
High	$1000, 1000$	$-200, 1200$
Low	$1200, -200$	$600, 600$

Figure 2.4: Simple model of a Duopoly payoff.

Since what matters are simply the ordered preferences rather than the magnitudes of the payoffs, it is simple to see that from firm 1's perspective, the best they can do is to choose to go "Low", since no matter what Firm 2 do, Firm 1 is always better to choose "Low" over "High". To see this, note that if Firm 2 were to choose "High" then Firm 1 would choose "Low" to get a profit of £1200. Alternatively, if Firm 2 were to choose "Low" then Firm 1 would choose "Low" since that would give them a profit of £600 as opposed to getting zero had they chosen "High". Hence in both cases, Firm 1 would choose the "Low" strategy. In other words, this is the same outcome as in Prisoner's Dilemma in that the combined profit of the companies is less than had they both chosen to sell "High".

2.2.4 Battle of the Sexes

In the Prisoner's Dilemma, the issue was whether the prisoners would choose the 'best' option of remaining Quiet. In the Battle of the Sexes (BoS) example, the players agree that it is better to cooperate, but they disagree about the best outcome for each of them.

Two people who really like each other (think of a recently engaged couple) wish to go out together to an event. The man would like to go to a football match. The woman likes football, but prefers going to the cinema. The man doesnt mind going to the cinema. If they go to separate events by themselves, they are both equally very unhappy as they are alone. So the 'quality-of-life' payoff for the man would be '2' if he can go to the football with his friend, and '1' if he goes to the cinema with her instead. But if he goes to the football whilst she chooses the cinema he has a payoff of zero. Similarly, she has a payoff of '2' if she can go to the cinema with him, a payoff of '1' if she goes to the football with him, but zero if she goes to the cinemas without him (and zero if she goes to the football and he goes to the cinema).

We can represent this strategic game as shown in Figure 2.5.

	Football	Cinema
Football	$2, 1$	$0, 0$
Cinema	$0, 0$	$1, 2$

Figure 2.5: Payoff matrix for the Battle of the Sexes

We will see later when we have the tools of finding so-called Nash equilibria, that this game has 2 stable choices in which each player would reduce their payoff by changes their choices. In this case both players prefer $(Football, Football)$ or $(Cinema, Cinema)$ equally, although their individual preferences differ.

2.2.5 Mathing Pennies

The previous games had at least some element of cooperation between the players. But Matching Pennies is a game that is purely about conflict.

Two people each have a coin and choose, simultaneously, whether to show the Head or the Tail of their coin. If they show the same side then Player 1 gets £1 from Player 2. If they are different then Player 1 pays player 2 £1. So one player can only gain at the expense of the other player. Each person cares only about the amount of money they get.

A strategic game that models this situation is shown in Figure 2.6.

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Figure 2.6: Payoff matrix for Matching Pennies

This is known as a strictly competitive game since the players interests are diametrically opposed. Also note that this is also an example of a 'zero sum' game - an important class of game. For the time being we note that there is no obvious stable situation to decide how to best play this game, for either player. We return to this game later when we introduce random elements into game strategies.

2.2.6 The Stag Hunt

This game has multiple equilibria but different outcome values.

A collection of hunters are hunting for food; either a stag or a hare. They individually prefer a share of a stag, but a hare is also acceptable (they would share the stag but whoever catches a hare keeps the hare to themselves). So, each of the hunters individually decides on one of two courses of action: she may remain focussed on the task of working together to catch a stag, or, she may catch a hare. If all hunters pursue the stag, they will with certainty catch it and share the stag equally. If any hunter diverts her energy to catching a hare, which is much easier, she will keep the hare but the stag will escape and the rest of the group will not eat anything.

	Stag	Hare
Stag	$2, 2$	$0, 1$
Hare	$1, 0$	$1, 1$

Figure 2.7: Payoff matrix for Stag-Hare game

Figure 2.7 gives one version of this game in a strategic form payoff matrix where there are just two hunters. In this Figure, there is a benefit (payoff) of zero if a hunter ends up with no food, a benefit of "1" if the hunter catches a hare and a benefit of "2" for a share in a stag. We will see later that the two equilibrium solutions are `[Hare, Hare]` and `[Stag, Stag]`. (ie if one hunter remains focussed on the task of stag-catching, the other hunter prefers to remain focussed, and similarly if one hunter prefers to hunt for hare, then so does the other hunter.)

2.3 Exercises for Week 1

2.3.1 Exercise

The simplest strategic story that yields the Prisoners' Dilemma arises when Adam and Eve both have access to a pot of money. Both are independently allowed either to give their opponent £2 from the pot, or to put £1 from the pot in their own pocket.

a) Write down the payoff table of the game on the assumption that the players care only about how much money they make.

b) Which strategy is strongly dominant? (We have yet to formally define it, but a dominant strategy is one where the player's preference payoff function cannot be improved by switching to another strategy.)

Answer:

	Give	Take
Give	$2, 2$	$0, 3$
Take	$3, 0$	$1, 1$

Figure 2.8: Give or Take?

It matters that the players are said to care only about money so that we can take their payoffs to be how much money they end up with. As in the regular Prisoners' Dilemma, the strategy `Take` strongly dominates the strategy `Give`.
In most games a player's best strategy usually depends on which strategy their opponent adopts. Prisoner's Dilemma is special in that the dominant strategy is independent of the opponent's choice (hence "strongly" dominates).

2.3.2 Exercise

The classic toy game called *Chicken* derives from the James Dean movie *"Rebel without a Cause"*, in which two teenage boys drive cars towards a cliff edge to see who chickens out first. (The same game is played by middle-aged drivers who approach each other in streets too narrow for them to pass without someone slowing down.)

	Slow	Speed
Slow	$2, 2$	$0, 3$
Speed	$3, 0$	$-1, -1$

Figure 2.9: Chicken

a) Explain why the payoff table in Figure 2.9 fits both stories. (ie determine each player's preferences of the different strategies).

b) Enclose the payoffs that correspond to best replies in a circle for player I and a square for player II.

c) Explain why neither player has a dominant strategy.

Answer:

a) For teenage boys, `Slow` corresponds to "chicken out" and `Speed` to "go for broke".
For this version of Chicken to fit the story, player I's preferences must satisfy

$$(Speed.Slow)_I(Slow, Slow)_I(Slow, Speed)_I(Speed, Speed)$$

Player II's preferences must satisfy

$$(Slow, Speed)_{II}(Slow, Slow)_{II}(Speed, Slow)_{II}(Speed, Speed)$$

	Slow	Speed
Slow	$2, 2$	$0, 3$
Speed	$3, 0$	$-1, -1$

Figure 2.10: Chicken

b) In terms of circles and squares, no row has all player I's payoffs enclosed in circles and no column has all player II's payoffs enclosed in squares. The strategy pairs `(Slow , Speed )` and `(Speed , Slow )` are Nash equilibria (we will discuss Nash equilibria next week) because each strategy in a pair is a best reply to the other strategy in the pair.

c) Neither player has a dominant strategy because what is best for a player depends on the choice made by the other.

2.3.3 Exercise

The favourite toy game of evolutionary biologists is called the Hawk-Dove Game. Two birds of the same species are competing for a scarce resource. Each can behave aggressively or passively. Payoffs are measured in terms of a bird's fitness such as the extra number of offspring the bird will have on average as a result of the way the game was played. If one bird is aggressive and the other is passive, the aggressive bird takes the entire resource. The aggressive bird then gets a payoff of $V > 0$, and the passive bird gets 0. If both birds are passive, the resource is shared and each bird gets a payoff of $V/2$. If both birds are aggressive, there is a fight and both birds receive a payoff of W . If $0 < W < V/2$, show that the Hawk-Dove Game is an example of the Prisoners' Dilemma. (Hint: Use $V = 4$ and $W = 1$ to illustrate the point.)
If the damage a bird is likely to receive in a fight is sufficiently large, then $W < 0$. Show (by using $W = -1$ for example) that the Hawk-Dove Game then reduces to a version of the game Chicken introduced in the previous exercise..

Answer: Figure 2.11 and Figure 2.12 show versions of the game when $V = 4$. When $W = 1$ we get the version of Prisoner's Dilemma shown on the left. When $W = -1$ we get the version of Chicken shown on the right.

	Dove	Hawk
Dove	$2, 2$	$0, 4$
Hawk	$4, 0$	$1, 1$

Figure 2.11: Hawk-Dove with $V = 4$ and $W = 1$

	Dove	Hawk
Dove	$2, 2$	$0, 4$
Hawk	$4, 0$	$-1, -1$

Figure 2.12: Hawk-Dove with $V = 4$ and $W = -1$

2.3.4 Exercise

Monopoly in Wonderland:
We will next consider a situation on the relevance of Prisoner's Dilemma for the economics of imperfect competition. But first we look at an example of a Monopoly.

Alice is a monopolist who manufactures Hats. One sheet of hat material can be bought for £1. Alice can make 1 hat out of 1 sheet of hat material, But she needs 4 sheets to make 2 hats, 9 sheets to make 3 hats and generally h^2 sheets to make h hats. So the Cost function for making h hats is given by h^2. (Labour is free in Wonderland).

If Alice can sell her hats at a price of £p each, her profit from selling h_A hats is

$$\pi = p \times h_A - h_A^2.$$

In economics, to know what price maximises profit, it is necessary to know a demand expression, ie how many hats would be bought at a given price. In Wonderland the demand equation is given by

$$p \times h_A = 30.$$

ie, the higher the price the fewer hats would be bought, not surprisingly.

How many hats should Alice make to maximise her profit, what price should she charge for her hats, and what profit will she make?

Answer:
The demand equation is the constraint that restricts her profit. Since Alice is the sole supplier of hats, she can meet any arbitrary demand. But whatever the demand and price, the demand equation tells Alice that if she makes h hats she should be able to sell them for

$$p = £\frac{30}{h_A}.$$

Putting this into her profit equation we find that her profit is

$$\pi = 30 - h_A^2.$$

So her strategy to maximise her profit is to make as few hats as possible. Indeed, her profit is maximised by making just 1 hat since this maximise π and from the demand equation she sells it for £30 and her profit is £29.
This example illustrates the principle of a monopoly, even in a demand-constrained envi-

ronment. In a monopoly, the supplier is the *price maker*.
So what happens when another supplier enters the market? We now explore the case of a Duopoly.

2.3.5 Exercise

Duopoly Example

Bob decides to enter the hat-making market in Wonderland. He will sell h_B hats to Alice's h_A hats. The demand equation still holds though, so the cost per hat will be restricted to $30/(h_A + h_B)$.

To simplify the situation, let us consider the case when Alice and Bob need to decide on the two strategies of selling either 1 or 2 hats. They need to make this decision independent of each other.

a) Assuming that demand is such that both Alice and Bob can sell all their hats they produce (ie, 1 or 2), show that this situation may be expressed as a strategic game in which the payoff or preference matrix may be expressed as Figure 2.13)(where we have chosen to represent the two strategies as `Dove` and `Hawk`). Note: to identify a strategic Game, recall from the lectures that a Strategic Game is characterised by identifying (i) Players, (ii) strategies, and (iii) a preference relation. The payoff matrix is a convenient way to find out about its solution.

	Dove	Hawk
Dove	$14, 14$	$9, 16$
Hawk	$16, 9$	$11, 11$

Figure 2.13: Duopoly payoff matrix example.

b) What will happen? ie how many hats will be sold, for how much and how much profit will Bob and Alice each make?

Answer:

a) To put this as a strategic game we identify the Players as Alice and Bob, their 2 Strategies as `(Dove = Make 1 hat)` or `(Hawk = Make 2 hats)`, and their preference function is to maximise their individual profit. Therefore if Profit is the Preference function then if they each make 1 hat (ie `Dove,Dove)` strategy), there are a total of 2 hats in Wonderland that they can sell for £15 each (from the demand equation). So their individual profit will be $£15 \times 1 - 1^2 \equiv £14$. Similarly, if they each make 2 hats (`(Hawk,Hawk)` strategy) there will be 4 hats in existence which they can sell for £7.50. So their profit is $2 \times £7.50 - 2^2 \equiv £11$. If Alice chooses `Hawk` (ie make 2 hats) and Bob chooses `Dove` then there are 3 hats in existence which sell for $30/3 = £10$ each. So Alice's profit is $2 \times £10 - 2^2 \equiv £16$, and Bob's profit is $1 \times £10 - 1^2 \equiv £9$. Similarly if Alice chooses `Dove` and Bob chooses `Hawk` the profits will be reversed.
Therefore Figure 2.14 shows the payoff matrix with each player's best payoff identified, for each different strategy of the other player. It can be seen that the `(Hawk, Hawk)` strategy is strictly dominating, in that any other choice would mean less individual profit independent of which strategy the other player took. So, each player makes £11 profit, by selling 2 hats each for £7.50, whereas

had they colluded to price-fix they could have made a profit of £14. This is a Prisoner's Dilemma game of course.

	Dove	Hawk
Dove	14, 14	9, 16
Hawk	16, 9	11, 11

Figure 2.14: Duopoly payoff matrix example.

b) Neither player maximises their individual potential profit. `Hawk` strongly dominates `Dove`. So no matter how many hats Alice chooses to make, Bob is always better off choosing the `Hawk` strategy. The same goes for Alice.
This outcome indicates why competition is good for the consumer in that more goods are produced at a cheaper per-item cost, and indicates why the initial "common sense" solution is not necessarily a good thing! Compare this situation with that in the previous Monopoly Example.

2.3.6 Exercise

Duopolies and "Prisoner's Delight"!

Duopolies do not always give rise to a Prisoner's Dilemma Game.

a) Show that in the case that the Demand Equation used previously now becomes

$$p \times (h_A + h_B) = 12$$

then the pay-off matrix for Alice and Bob now becomes as shown in Figure 2.15.

	Dove	Hawk
Dove	5, 5	3, 4
Hawk	4, 3	2, 2

Figure 2.15: Prisoner's Delight Duopoly payoff matrix example.

b) What is the dominating strategy for each player and hence why is this called a Prisoner's Delight game?

Answer:

a) As before, if they each make 1 hat (ie `Dove,Dove`) strategy), there are a total of 2 hats in Wonderland that they can sell for £6 each (from the new demand equation). So their individual profit will be $£6 \times 1 - 1^2 \equiv £5$. Similarly, if they each make 2 hats (`(Hawk,Hawk)` strategy) there will be 4 hats in existence which they can sell for £3. So their profit is $2 \times £3 - 2^2 \equiv £2$. If Alice chooses `Hawk` (ie make 2 hats) and Bob chooses `Dove` then there are 3 hats in existence which sell for $12/3 = £4$ each. So Alice's profit is $2 \times £4 - 2^2 \equiv £4$, and Bob's profit is $1 \times £4 - 1^2 \equiv £3$. Similarly if Alice chooses `Dove` and Bob chooses `Hawk` leading to the new payoff matrix..

b) Figure 2.16 shows the best strategies for each player highlighted as a circle for Player 1 and a rectangle for Player 2. It can be seen that `Dove` strongly dominates `Hawk`. So, the common agreed best strategy is `(Dove,Dove)`. It is called a Prisoner's Delight since this equilibrium is also a best possible joint outcome for the players and Alice and Bob extract the maximum possible amount of money from the customers!

	Dove	Hawk
Dove	$5, 5$	$3, 4$
Hawk	$4, 3$	$2, 2$

Figure 2.16: Prisoner's Delight Duopoly payoff matrix showing best strategies.

2.3.7 Exercise

Using the same scenario as in the previous two exercises, show that the situation gives rise to a Prisoner's Dilemma whenever the demand equation exceeds 18 and is a Prisoner's Delight whenever the demand equation satisfies $p \times (h_A + h_B) < 18$. [Hint: Evaluate the Profit in the `Dove` strategy as well as the `Hawk` strategy and note that for Prisoner's Delight we must have `Dove` $\succeq$ `Hawk`, and for Prisoner's Dilemma we need `Hawk` $\succeq$ `Dove`.]

Answer:

From the Hint, we want to check when `Hawk` will dominate `Dove` always, as a strategy. Since the preference function is profit, for the `Dove` strategy we have

$$\pi_{Dove} = p \times 1 - 1^2$$

and for the `Hawk` strategy we have

$$\pi_{Hawk} = p \times 2 - 2^2.$$

We need a value for p, the cost of a hat. This we can get from the demand equation

$$p \times (h_A + h_B) = X$$

where X is the value of the demand equation. The query is, what value of X do we need to make it Prisoner's Delight or Prisoner's Dilemma. Substituting into the expression for the profits, and noting that for Prisoner's Dilemma for `Hawk` to dominate `Dove` regardless of the choice of the other player, we must have
`(Hawk,Dove)` $\succeq$ `(Dove,Dove)`, and `(Hawk,Hawk)` $\succeq$ `(Dove,Hawk)`.
So we can construct the generic payoff matrix in Figure 2.17:

For `Hawk` to dominate we can see that we need

$$2 \times \frac{X}{3} - 4 > \frac{X}{2} - 1,$$

	Dove	Hawk
Dove	$\frac{X}{2}-1, \frac{X}{2}-1$	$\frac{X}{3}-1, 2\times\frac{X}{3}-4$
Hawk	$2\times\frac{X}{3}-4, \frac{X}{3}-1$	$2\times\frac{X}{4}-4, 2\times\frac{X}{4}-4$

Figure 2.17: Duopoly with general demand equation.

and

$$2\times\frac{X}{4}-4>\frac{X}{3}-1.$$

The consistent solution of both of these equations is $X > 18$ as required. So, if $X > 18$ we have a Prisoner's Dilemma situation. If $X < 18$ we have a Prisoner's Delight.

2.3.8 Exercise

Social Dilemmas

The last example we consider is a little different and is known in the political science literature as the **Tragedy of the Commons** which is a dilemma arising from the situation in which multiple individuals, acting independently and rationally consulting their own self interest, will ultimately deplete a shared limited resource, even when it is clear that it is not in anyone's long-term interest for this to happen. It is an example of the situation your parents warned you about "what if everyone did that !!"

This dilemma was described in an influential article titled "The Tragedy of the Commons", written by ecologist Garrett Hardin and first published in the journal Science in 1968.

Hardin's Commons Theory is frequently cited to support the notion of sustainable development, meshing economic growth and environmental protection, and has had an effect on numerous current issues, including the debate over global warming.

Consider 10 familes who have the 'common' right to herd goats on one square mile of grazing land. The milk a goat gives depends on how much grass they eat. A goat that grazes on a fraction a of the common land will produce

$$b = e^{[1-1/(10a)]}$$

buckets of milk a day. Note that a goat grazing on one tenth of the land will produce 1 bucket of milk per day. As the fraction of land for a goat decreases, the milk yield will also decrease eventually down to zero when $a = 0$.

A town planner decides that if there are to be N goats on this land, then each goat has $a = 1/N$ of the land to graze giving total milk production of

$$M = N \times b = Ne^{[1-N/10]}.$$

a) Show that this is maximised when $N = 10$ and hence the town planner recommends that each family would have 1 goat each and get 1 bucket of milk each.

b) However, each family can choose itself how many goats they want. So, assuming each family decides independently to maximise their own "goatage", show that the total milk production would be just $100 \times e^{-9} \approx 0.012$.
(Hint: If one family chooses g goats independently of their neighbours, then their individual milk

quota from their own personal goats would be $m = g \times b$. So, if the total of all the other families' goats was G, say, write an expression to maximise to find the optimal value of g.)

Answer:

a) To find the maximum, find out where the derivative is zero, or where the turning points are. The expression is of the form $y = \alpha x e^{-x}$ for some constant α. Since

$$\frac{dy}{dx} = \alpha e^{-x} - \alpha x e^{-x}$$

then $dy/dx = 0$ when $e^{-x} - xe^{-x} = 0$, or when $x = 1$. So, back in the original expression, where $x = N/10$, then if $x = 1$ then $N = 10$, and hence each family would get 1 full bucket of milk per day.

b) So if the total number of goats were $G + g$ (ie, 9 families have a total number of goats G and the one family wants to optimise their number of goats g), then the milk yield to the family would be

$$m = gb = ge^{[1-(g+G)/10]} = e^{-G/10} \times ge^{[1-g/10]}$$

Since $e^{-G/10}$ is a constant and there is nothing the family can do about the other families choice, their own optimisation function is exactly the same form as before. So maximising, we find that $g = 10$. But all the families would also reason this way. Hence, 10 families would each decide to have 10 goats. 100 goats grazing on the land would give a total milk yield of

$$M = 100e^{-9} \approx 0.012.$$

The strategy of keeping 10 goats is strongly dominant over the other strategies of keeping 1 or 2 or 3 .. etc goats. Nevertheless, each family would have been better off if they had taken the townplanner's advice and stuck to just 1 goat each.

Chapter 3

Two-person Games

Previously we have discussed simple examples of games that we can express in terms of matrices. We now introduce more rigour into the analysis. The first part we focus on two-person zero-sum games. We will show how equilibria in such games can be determined. First let us re-examine some concepts already introduced. We saw that several two-person games could be expressed in terms of a payoff *bimatrix* , ie a matrix with two numbers in each cell, one for each player. We will separate the payoffs and strategies of individual players, as we can then generalise to multiple players easier later, and introduce some different notation.

Previously we defined a normal form game in terms of the triple *(Players,Strategies/Actions, Preference relation/Payoff function).* ie a normal form game can be interpreted as $G = \{\boldsymbol{P}, \Sigma, \pi\}$ where $\boldsymbol{P} = players$, Σ =the set of all strategies of all players, and π = the payoff function which acts on the set of Actions representing the preferences $\succeq_i$ for each of the N players, $i \in \{1, \ldots, N\}$.

The set of all Strategies, or Actions may be decomposed into those of each player, individually (Σ_i) as

$$\Sigma = \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_N$$

and the payoff relation then defines a separate payoff value for each player, ie the payoff relation

$$\pi : \Sigma \to \Re^N$$

for each player is

$$\pi_i : \Sigma \to \Re.$$

Then we have the modified definition of a game in normal form, as

Definition 3 (Normal Form)**.** *Let $\boldsymbol{P}$ be a set of players, $\Sigma = |\Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_N$ be a a strategy space and let $\pi : \Sigma \to \Re^N$ be strategy payoff function, then $G = (\boldsymbol{P}, \Sigma, \pi)$ is a game in normal form.*

Remark 1. *If $G = (\boldsymbol{P}, \Sigma, \pi)$ is a normal form game, then the function $\pi_i : \Sigma \to \Re$ is the payoff function for Player P_i and returns the i-th component of the function π.*

We previously said that a normal form and a strategic form game were equivalent. Here we make explicit what their difference is, and we will define a strategic game as a normal form game when the payoff values are expressed as matrix games.

Definition 4 (Strategic Form 2 Player Games)**.** *Let $G = \{\boldsymbol{P}, \Sigma, \pi\}$ be a normal form game with $\boldsymbol{P} = \{P_1, P_2\}$ and $\Sigma = \Sigma_1 \times \Sigma_2$. Let us arrange the strategies in Σ_i as $\{\sigma_1^i, \ldots \sigma_{n_i}^i\}$ for $i \in \{1, 2\}$.*

Then for each player there is a matrix $\boldsymbol{A}_i \in \Re^{n_1 \times n_2}$ where the row-column value (r, c) in $\boldsymbol{A}_i$ is given by $\pi_i(\sigma_r, \sigma_c)$. Then the tuple $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}_1, \boldsymbol{A}_2)$ is a two-player game in strategic form.

Remark 2. *Games with two players given in strategic form are also sometimes called matrix games because they are defined completely by matrices. Note also that by convention, Player P_1's strategies correspond to the rows of the matrices, while Player P_2's strategies correspond to the columns of the matrices.*

Remark 3. *Definition 4 can be extended to N player games. However, we no longer have simple matrices with payoff values for various strategies. Instead, we construct N, N-dimensional arrays (or tensors). So a game with 3 players yields 3 arrays with dimension 3. They have multiple indices, instead of just 1 index like a vector or 2 indices like a matrix. The elements of the array for Player i store the various payoffs for Player i under different strategy combinations of the different players. If there are three players, then there will be three different arrays, one for each player.*

Definition 5. *(Constant / General Sum Game).*
Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form. If there is a constant $C \in \Re$ so that for all tuples $(\sigma_1, \ldots, , \sigma_N) \in \Sigma$ we have:

$$\sum_{i=1}^{N} \pi_i(\sigma_1, \ldots, , \sigma_N) = C$$

then G is called a constant sum game. *If $C = 0$, then G is called a* zero sum game.

We also note a definition of a symmetric game, when Player 1 and Player 2's preferences may be interchanged without altering the game.

Definition 6. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}_1, \boldsymbol{A}_2)$ be a 2 person strategic game. If $\boldsymbol{A}_1^T = \boldsymbol{A}_2$ then G is a* symmetric game.

We also now specify an equilibrium formally as:

Definition 7. *(Equilibrium).*
A strategy profile $(\sigma_1^*, \ldots, , \sigma_N^*) \in \Sigma^N$ is an equilibrium if for all i.

$$\pi_i(\sigma_1^*, \ldots, \sigma_i^*, \ldots, \sigma_N^*) \geq \pi_i(\sigma_1^*, \ldots, \sigma_i, \ldots, \sigma_N^*)$$

where $\sigma_i \in \Sigma_i$.

Example Recall the Example given in the previous chapter on the competitive game known as **matching pennies** in which two people each had a coin and chose, simultaneously, whether to show the Head or the Tail of their coin. If they showed the same side then Player 1 received £1 from Player 2. If they were different then Player 1 paid player 2 £1.

The previous payoff matrix that we used to model this situation was shown as in Figure 3.1.

In this case we have

- $\boldsymbol{P} = \{Player\ 1, Player\ 2\}$

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Figure 3.1: Payoff matrix for Matching Pennies

	Dove	Hawk
Dove	$5, 5$	$3, 4$
Hawk	$4, 3$	$2, 2$

Figure 3.2: Prisoner's Delight Duopoly payoff matrix example.

- $\Sigma_1 = \{Head, Tail\}$
  $\Sigma_2 = \{Head, Tail\}$.
- $$\boldsymbol{A}_1 = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$
  $$\boldsymbol{A}_2 = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}$$

  since $\pi_1(Heads, Heads) = +1$, but $\pi_2(Heads, Heads) = -1$ for example.

This is a *zero-sum* game since $\boldsymbol{A}_1 + \boldsymbol{A}_2 = 0$.

Example Recall the exercise in the previous chapter on Prisoner's Delight which had the payoff bimatrix given in Figure 2.15, reproduced in Figure 3.2 for ease. In that example we denoted the two strategies as `Dove` (ie make 1 hat) and `Hawk` (ie make 2 hats).

In this case we have

- $\boldsymbol{P} = \{Alice, Bob\}$
- $\Sigma_1 = \{Dove\ (make\ 1\ hat), Hawk\ (make\ 2\ hats)\}$
  $\Sigma_2 = \{Dove\ (make\ 1\ hat), Hawk\ (make\ 2\ hats)\}$

- $$\boldsymbol{A}_1 = \begin{bmatrix} 5 & 3 \\ 4 & 2 \end{bmatrix}$$
$$\boldsymbol{A}_2 = \begin{bmatrix} 5 & 4 \\ 3 & 2 \end{bmatrix}$$

  since in this case the preference function is simply profit, and $\pi_1(Dove, Hawk) = £3$ and $\pi_2(Dove, Hawk) = £4$, for example.

This is not a constant sum game since $\boldsymbol{A}_1 + \boldsymbol{A}_2 \neq constant$.

In the next section, to discuss techniques for evaluating matrix games we need a few extra vector and matrix definitions.

Definition 8 (Standard Basis Vector)**.** *The standard basis vector $\boldsymbol{e}_i \in \Re^n$ is:*

$$\boldsymbol{e}_i = (\underbrace{0, 0, \ldots, 0}_{i-1}, 1, \underbrace{0, 0, \ldots, 0}_{n-i-1})$$

Note, this definition is only valid for $n \geq i$. Further the standard basis vector $\boldsymbol{e}_i$ is also the i-th row or column of $\boldsymbol{I}_n$, the n-dimensional identity matrix.

Definition 9 (Unit and Zero Vectors)**.** *The vector $\mathbf{1} \in \Re^n$ is the one vector*

$$\mathbf{1} = (1, 1, ..., 1)$$

Similarly, the zero vector $\mathbf{0}$ is

$$\mathbf{0} = (0, 0, ..., 0) \in \Re^N$$

We can use standard basis vectors to compute the payoff to Player P_i when a specific set of strategies are used. Our next proposition relates the strategy set Σ to pairs of standard basis vectors and reduces computing the payoff function to simple matrix multiplication.

Proposition 1. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}_1, \boldsymbol{A}_2)$ be a two-player game in strategic form with*

$$\Sigma_1 = (\sigma_1^1 \ldots \sigma_m^1)$$
$$\Sigma_2 = (\sigma_1^2 \ldots \sigma_n^2)$$

If Player P_1 chooses strategy σ_r^1 and Player P_2 chooses strategy σ_c^2 , then

$$\pi_1(\sigma_r^1, \sigma_c^2) = \boldsymbol{e}_r^T \boldsymbol{A}_1 \boldsymbol{e}_c \tag{3.1}$$
$$\pi_2(\sigma_r^1, \sigma_c^2) = \boldsymbol{e}_r^T \boldsymbol{A}_2 \boldsymbol{e}_c \tag{3.2}$$

Proof. For any matrix $\boldsymbol{A} \in \Re^{n \times m}$, $\boldsymbol{A}\boldsymbol{e}_c$ returns column c of matrix $\boldsymbol{A}$, that is, $\boldsymbol{A}_{\cdot,c}$. Likewise $\boldsymbol{e}_r^T \boldsymbol{A}_{,c}$ is the rth element of this vector. Thus, $\boldsymbol{e}_r^T \boldsymbol{A} \boldsymbol{e}_c$ is the (r, c)th element of the matrix $\boldsymbol{A}$. By definition, this

must be the payoff for the strategy pair (σ_r, σ_c) for Player P_1. A similar argument follows for Player P_2 and matrix $\boldsymbol{A}_2$.

□

Remark 4. *What Proposition 1 says is that for two-player matrix games, we can relate any (pure) choice of (a single) strategy that Player P_i makes with a unit vector. Thus, we can actually define the payoff function in terms of vector and matrix multiplication. We will see that this can be generalised to cases when the strategies of the players are not represented by standard basis vectors.*

Example Re-consider the Prisoner's Delight Duopoloy example. If Alice adopts the Hawk strategy and makes 2 hats, and Bob adopts the Dove strategy and makes 1 hat, then Alice and Bob's strategy vectors can be expressed as

$$\boldsymbol{e}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad \& \quad \boldsymbol{e}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

Recall the payoff matrices for these two players were

$$\boldsymbol{A}_1 = \begin{bmatrix} 5 & 3 \\ 4 & 2 \end{bmatrix}$$

$$\boldsymbol{A}_2 = \begin{bmatrix} 5 & 4 \\ 3 & 2 \end{bmatrix}$$

So we can compute

$$\pi_1(Hawk, Dove) = \boldsymbol{e}_2^T \boldsymbol{A}_1 \boldsymbol{e}_1 = \begin{bmatrix} 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 5 & 3 \\ 4 & 2 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 4$$

$$\pi_2(Hawk, Dove) = \boldsymbol{e}_2^T \boldsymbol{A}_2 \boldsymbol{e}_1 = \begin{bmatrix} 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 5 & 4 \\ 3 & 2 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 3$$

Remark 5. *The next proposition relates the definition of Equilibria and the properties of matrix games and strategies.*

Proposition 2. *(Equilibrium).*
Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}_1, \boldsymbol{A}_2)$ be a two-player game in strategic form with $\Sigma = \Sigma_1 \times \Sigma_2$.
$(\sigma_i, \sigma_j) \in \Sigma_1 \times \Sigma_2$ is an equilibrium strategy if and only if (iff)

$$\boldsymbol{e}_i^T \boldsymbol{A}_1 \boldsymbol{e}_j \geq \boldsymbol{e}_k^T \boldsymbol{A}_1 \boldsymbol{e}_j \quad \forall k \neq i \tag{3.3}$$

$$\boldsymbol{e}_i^T \boldsymbol{A}_2 \boldsymbol{e}_j \geq \boldsymbol{e}_i^T \boldsymbol{A}_2 \boldsymbol{e}_l \quad \forall l \neq j \tag{3.4}$$

Proof. From Proposition 1, we know that:

$$\pi_1(\sigma_r^1, \sigma_c^2) = \boldsymbol{e}_r^T \boldsymbol{A}_1 \boldsymbol{e}_c \tag{3.5}$$
$$\pi_2(\sigma_r^1, \sigma_c^2) = \boldsymbol{e}_r^T \boldsymbol{A}_2 \boldsymbol{e}_c \tag{3.6}$$

Therefore, using these in equations (3.3),(3.4), we have

$$\pi_1(\sigma_i^1, \sigma_j^2) \geq \pi_1(\sigma_k^1, \sigma_j^2) \ \forall k \neq i \tag{3.7}$$
$$\pi_2(\sigma_i^1, \sigma_j^2) \geq \pi_2(\sigma_i^1, \sigma_l^2) \ \forall l \neq j \tag{3.8}$$

Thus from the previous Definition 7 of an equilibrium, it is clear that $(\sigma_i^1, \sigma_j^2) \in \Sigma$ is an equilibrium strategy. □

Remark 6. *We can now think of relating a strategy choice k for player i,$\sigma_k^i \in \Sigma_i$ with the unit vector $\boldsymbol{e}_k$.*

3.1 2 person zero-sum games

We focus for the time being on 2 person zero sum games in which $\boldsymbol{A}_1 = -\boldsymbol{A}_2 = \boldsymbol{A}$. We can also assume that $\Sigma = \Sigma_1 \times \Sigma_2$ and $\Sigma_1 = \{\sigma_1^1, \ldots, \sigma_m^1\}$ and $\Sigma_2 = \{\sigma_1^2, \ldots, \sigma_n^2\}$. Then a 2-person zero-sum game in strategic form can be expressed as the tuple $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ with $\boldsymbol{A} \in \Re^{m \times n}$.

3.1.1 Saddle Points

Occasionally it is easy to solve a 2-person zero-sum game. If some entry a_{ij} of the payoff matrix $\boldsymbol{A}$ has the property that

1. a_{ij} is the minimum of the ith row, and
2. a_{ij} is the maximum of the jth column,

then we say a_{ij} is a *saddle point*. If a_{ij} is a saddle point, then Player I can then win at least a_{ij} by choosing row i, and Player II can keep her loss to at most a_{ij} by choosing column j. Hence in this case a_{ij} is called the *value* of the game (to be defined later).

Before we come to the definition and proof of the saddle point theorem, let us consider the following 2-person, 3 strategy examples.

Example Consider the following payoff matrix:

$$\boldsymbol{A} = \begin{bmatrix} 4 & 1 & -3 \\ 3 & 2 & 5 \\ 0 & 1 & 6 \end{bmatrix}$$

The central entry, 2, is a saddle point, since it is a minimum of its row and maximum of its column. Thus it is optimal for I to choose the second row, and for II to choose the second column. The value of the game is 2, and `(0, 1, 0)` is an optimal strategy for both players.

Example Consider the following payoff matrix for Player 1 in a zero sum game. (Therefore, Player 2's payoff matrix is the negative of this one).

$$\boldsymbol{A} = \begin{bmatrix} -15 & -35 & 10 \\ -5 & 8 & 0 \\ -12 & -36 & 20 \end{bmatrix}$$

There is an equilibrium strategy in this case, and specifically `(0,1,0)` for Player 1. Why is this the equilibrium strategy?

Consider the situation from Player 1's perspective. Since it is a zero sum game, any gain of Player 2 must be at the expense of Player 1. So Player 1 scans each row to see what is the list of worst losses in each row. In this example, this would give us the vector $(-35, -5, -36)$. Since Player 1 wishes to reduce her potential losses, she would choose row 2 since this has the smallest loss.

Now consider Player 2's position. By the same argument, she also wishes to maximise her gain, given that Player 1 gets to select the row. So Player 2 looks at the columns to see in each column what would be Player 1's best score if that location was selected. In this case Player 2 would select $(-5, 8, 20)$ from each column as the best possible outcomes for Player 1. Since Player 2's best outcome is at the expense of Player 1's outcome. Player 2 would choose the column that would minimise the best payoff to Player 1 (which would therefore maximise Player 2's payoff). In this case, she would therefore choose column 1.

So both players decide on the strategy pair $(r, c) = (2, 1)$ which has a payoff value of -5 to Player 1. They both agree on this point, and hence this is an equilibrium strategy choice.

We can embody the results of these examples as a Theorem:

Theorem 2. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ be a zero-sum two player game. A strategy pair $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ is an equilibrium strategy if and only if:*

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = \max_{k \in \{1,\ldots m\}} \min_{l \in \{1,\ldots n\}} A_{k,l} = \min_{l \in \{1,\ldots n\}} \max_{k \in \{1,\ldots m\}} A_{k,l} \tag{3.9}$$

Proof. ($\Rightarrow$) Suppose that $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ is an equilibrium solution. Then we know that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j \geq \boldsymbol{e}_k^T \boldsymbol{A} \boldsymbol{e}_j \quad k \in \{1, \ldots m\} \tag{3.10}$$

and

$$\boldsymbol{e}_i^T (-\boldsymbol{A}) \boldsymbol{e}_j \geq \boldsymbol{e}_i^T (-\boldsymbol{A}) \boldsymbol{e}_l \quad l \in \{1, \ldots n\}$$

or

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j \leq \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_l \quad l \in \{1, \ldots n\} \tag{3.11}$$

Since $\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = A_{ij}$ and that Equation (3.10) holds if and only if: $A_{ij} \geq A_{kj}$ for all $k \in \{1, \ldots, m\}$. From this we deduce that element i must be a maximal element in column $\boldsymbol{A}_{\cdot,j}$

Based on this, we know that for each row $k \in \{1, \ldots, m\}$:

$$A_{ij} \geq \min\{A_{kl} : l \in \{1, \ldots, n\}\} \tag{3.12}$$

To see this, note that for a fixed row $k \in \{1, \ldots, m\}$:

$$A_{kj} \geq \min\{A_{kl} : l \in \{1, \ldots, n\}\}$$

This means that if we compute the minimum value in a row k, then the value in column j, $\boldsymbol{A}_{kj}$ must be at least as large as that minimal value. But, Expression 3.12 implies that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = A_{ij} = \max_{k \in \{1,\ldots,m\}} \min_{l \in \{1,\ldots,n\}} A_{kl} \tag{3.13}$$

Likewise, Equation 3.11 holds if and only if

$$A_{ij} \leq A_{il} \quad \forall l \in \{1, \ldots, n\}. \tag{3.14}$$

From this we deduce that element j must be a minimal element in row $A_{i,.}$. Based on this, we know that for each column $l \in \{1, \ldots, n\}$

$$A_{ij} \leq \max\{A_{kl} : k \in \{1, \ldots, m\}\} \tag{3.15}$$

To see this, note that for a fixed column $l \in \{1, \ldots, n\}$:

$$A_{il} \leq \max\{A_{kl} : k \in \{1, \ldots, m\}\}$$

This means that if we compute the maximum value in a column l, then the value in row i, A_{il} must not exceed that maximal value. But Expression 3.15 implies that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = A_{ij} = \min_{l \in \{1,\ldots,n\}} \max_{k \in \{1,\ldots,m\}} A_{kl} \tag{3.16}$$

Thus it follows that:

$$A_{ij} = \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = \max_{k \in \{1,\ldots m\}} \min_{l \in \{1,\ldots n\}} A_{kl} = \min_{l \in \{1,\ldots n\}} \max_{k \in \{1,\ldots m\}} A_{kl}$$

($\Leftarrow$) To prove the converse, suppose that

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j = \max_{k \in \{1,\ldots m\}} \min_{l \in \{1,\ldots n\}} A_{ij} = \min_{l \in \{1,\ldots n\}} \max_{k \in \{1,\ldots m\}} A_{ij}$$

Consider

$$\boldsymbol{e}_k^T \boldsymbol{A} \boldsymbol{e}_j = A_{kj}$$

The fact that

$$A_{ij} = \max_{k \in \{1,\ldots m\}} \min_{l \in \{1,\ldots n\}} A_{kl}$$

implies that $A_{ij} \geq A_{kj}$ for any $k \in \{1, \ldots, m\}$. To see this note that

$$C_{max} = \{c^* = \max\{A_{kl} : k \in \{1, \ldots, m\}\} : l \in \{1, \ldots, n\}\} \tag{3.17}$$

and $A_{ij} \in C_{max}$ by construction. Thus it follows that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j \geq \boldsymbol{e}_k^T \boldsymbol{A} \boldsymbol{e}_j$$

for any $k \in \{1, \ldots, m\}$.

By a similar argument we know that:

$$A_{ij} = \min_{l \in \{1,\ldots n\}} \max_{k \in \{1,\ldots m\}} A_{kl}$$

implies that $A_{ij} \leq A_{il}$ for any $l \in \{1, \ldots, n\}$.

To see this remember:

$$R_{min} = \{r_k^* = \min\{A_{kl} : l \in \{1, \ldots, n\}\} : k \in \{1, \ldots, m\}\}$$

and $A_{ij} \in R_{min}$ by construction. Thus it follows that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j \leq \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_l$$

for any $l \in \{1, \ldots, n\}$. Thus $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ is an equilibrium solution. This completes the proof. □

This equilibrium, if it exists, is unique, as the following theorem dictates:

Theorem 3. *Suppose that $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ be a zero-sum two player game. Let $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ be an equilibrium strategy pair for this game. If $(\boldsymbol{e}_k, \boldsymbol{e}_l)$ is a second equilibrium strategy pair, then*

$$A_{ij} = A_{kl} = A_{il} = A_{kj}$$

Definition 10. *(Saddle Point).*
Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ be a zero-sum two player game. If $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ is an equilibrium, then it is called a saddle point.

Definition 11. *(Game Value).*
Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ be a zero-sum game. If there exists a strategy pair $(\boldsymbol{e}_i, \boldsymbol{e}_j)$ so that:

$$\max_{k \in \{1,\ldots,m\}} \min_{l \in \{1,\ldots,n\}} A_{kl} = \min_{l \in \{1,\ldots,n\}} \max_{k \in \{1,\ldots,m\}} A_{kl}$$

then:

$$V_G = \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_j \tag{3.18}$$

is the value *of the game.*

Remark 7. *We will see that we can define the value of a zero-sum game even when there is no equilibrium point in strategies. Using Theorem 3 we can see that this value is unique, that is any equilibrium pair for a game will yield the same value for a zero-sum game. This is not the case in a general-sum game.*

Remark 8. *Not all zero sum games have saddle points.*

3.2 Exercises for Week 2

Today we practice some examples to illustrate the principles from the week's lectures. We focus on saddle point solutions. Recall the following that you should now have learned:

For player A:
The minimum value in each row represents the least gain to him if he chooses that particular strategy. These are written in the matrix by row minima. He will then select the strategy that gives maximum gain among the row minimum values. This choice of player A is called the *maximin criterion* and the corresponding gain is called the *maximin value* of the game.

Similarly, for player B who is assumed to be the loser, the maximum value in each column represents the maximum loss to him if he chooses his particular strategy. These are written as column maxima. He will select that strategy which gives minimum loss among the column maximum values. This choice of player B is called the *minimax criterion*, and the corresponding loss is the *minimax value* of the game.

When these are the same, we said that this was an equilibrium solution of the game, the corresponding strategies were optimal, that the solution was a saddle point, and the value at this point was the value of the game.

3.2.1 Exercise

Consider a two person zero sum game in Figure 3.3 with the given payoff matrix for player A. The rows indicate the payoffs for Player A's three possible strategies. Player B has 2 possible strategies given by the columns and Player B's 'gains' are the negative of Player A's gains.

What strategies should the players adopt and does the game have a value?

	B1	B2
A1	4	3
A2	8	6
A3	2	4

Figure 3.3: Exercise 1

Answer: Suppose that player A starts the game knowing that whatever strategy she adopts, B will select that particular counter strategy which will minimize the payoff to A. So, if A selects the strategy A1 then B will select B2 as this corresponds to the minimum payoff to A having chosen A1. Similarly, if A chooses the strategy A2, she may gain 8 or 6 depending upon the choice of B. If A chooses A2 then A can guarantee a gain of at least `min (8,6) = 6` irrespective of the choice of B. Obviously A would prefer to maximise her minimum assured gains. Choosing strategy A3 would result in less gain than choosing A2 and so A3 would not be considered. (Actually for this specific example we can more easily argue that both strategies A1 and A3 are strictly dominated by strategy A2 and so player A would choose A2 regardless of considering B's options! However the point of this exercise is to use the saddle point argument.) In this example the selection of strategy A2

gives the maximum of the minimum gains to A.
This gain is called the maximin value of the game and the corresponding strategy as the maximin strategy.
On the other hand, player B prefers to minimise his losses. If he plays strategy B1 his loss is at most `max (4,8,2) = 8` regardless of the strategy selected by A. If B plays B2 then he loses no more than `max (3,6,4) = 6`. B wishes to minimise his maximum possible losses. In this example, the selection of strategy B2 gives the minimum of the maximum losses of B, this loss is called the minimax value of the game and the corresponding strategy the minimax strategy. In this example,
`Maximin value (V) = minimax value (V)`
i.e. `maximum (row minima) = minimum value (column maxima)` or
`max (row) = 6 = min (column)`
or
`max ( min [` A_{ij} `]) = 6 = min (max [` A_{ij} `]).`

	B1	B2
A1	4	3
A2	8	6
A3	2	4

Figure 3.4: Exercise 1 solution

Summary of approach:

STEP 1: Select the minimum elements of each row of the payoff matrix and mark them with a circle. These are the row minima of the respective rows.

STEP 2: Select the greatest element of each column of the payoff matrix and mark them with a square. These are the column maxima of the respective columns.

STEP 3: If there appears an element in the payoff matrix marked with both a circle and a square, then this element represents the value of the game and its position is a saddle point of the payoff matrix.

3.2.2 Exercise

For the game with payoff matrix given in Figure 3.5, determine the Player A nd Player B's best strategies and determine if the game is biased to any specific player and whether the game has a value.

	B1	B2	B3
A1	-2	3	-4
A2	7	5	-5

Figure 3.5: Exercise 2

Answer:

	B1	B2	B3
A1	-2	3	-4
A2	7	5	-5

Figure 3.6: Exercise 2 solution

For Player A the row minima are `(-4, -5)`.
For Player B the column maxima are `(7,5,-4)`.
So the largest of the row minima is `-4` which is the maximin value, $\overline{V}$.
The smallest of the column maxima is `-4` which is the minimax value $\underline{V}$.
So in this case the maximin and the minimax are the same and so the game *value* is `-4`. This is not a fair game, and is biased in favour of Player B.

3.2.3 Exercise

Find the range of values of p and q that will make the payoff element A_{22} a saddle point for the game whose payoff matrix is as given in Figure 3.7:

	B1	B2	B3
A1	2	4	5
A2	10	7	q
A3	4	p	6

Figure 3.7: Exercise 3

Answer: Determining the maximin and the minimax values first of all by ignoring the values of p and q we find that the row minima are `(2,7,4)` and the column maxima are `(10,7,6)`. So the maximin is `7` and the minimax is `6`. So the two solutions are not the same and so there is not a saddle point. However if we choose p and q to change one of the solutions, we can make A_{22} a saddle as long as $p \leq 7$ and $q \geq 7$.

3.2.4 Exercise

A company's management and the workforce's union are negotiating a new 3 year pay settlement. Each of the sides has 4 strategies to choose from:

I: Hard and aggressive bargaining.

II: Reasoning and logical approach.

III: Legalistic strategy.

IV: Conciliatory approach.

The gains to the *union* (the negatives of which are the costs to the company) are given in Figure 3.8. Note that the Union Strategies are the row choices and the Company choices are the COLUMNs.

	I	II	III	IV
I	20	15	12	35
II	25	14	8	10
III	40	2	10	5
IV	−5	4	11	0

Figure 3.8: Exercise 4

What strategies will the two sides adopt, and what is the value of the game? Is it biased to any particular side?

Answer: Applying the rules to determine the maximin and minimax values we find the solutions as given in Figure 3.9.
The company incurs costs and hence its strategy is to minimise its maximum losses. For the union, negotiation results in gain, hence the union strategy aims at maximising the minimum gains.

	I	II	III	IV
I	20	15	12	35
II	25	14	8	10
III	40	2	10	5
IV	−5	4	11	0

Figure 3.9: Exercise 4 solution.

So the minima of the rows are `(12,8,2,-5)` and the maxima of the columns are `(40,15,12,35)`. Hence the maximin is `12` and the minimax is 12. The company will choose strategy III (legalistic), and the union will adopt strategy I (aggressive bargaining). Therefore the two parties strategies lead to $A_{1,3}$ which is a saddle point, and so the value of the game is 12, and it is an unfair game biased in preference to the union.

3.2.5 Exercise

The company Sneaky Ltd has developed a sales forecasting function for its products and the products of its competitor, Plinky Ltd. There are four strategies S1, S2, S3 and S4 available to Sneaky Ltd. and three strategies P1, P2 and P3 to Plinky Ltd. The payoffs to Sneaky in units of thousands of pounds corresponding to all the twelve combinations of the strategies are given in Figure 3.10. From the table we can see that, for example, if strategy S1 is employed by Sneaky Ltd. and strategy P1 by Plinky Ltd., then there shall be a gain of £30,000 in quarterly sales to the former. Other entries can be similarly interpreted. Considering this information, state what would be the optimal strategy for Sneaky Ltd and Plinky Ltd? Is the game fair?

		Plinky		
		P1	P2	P3
Sneaky	S1	30	−21	1
	S2	18	14	12
	S3	−6	28	4
	S4	18	6	2

Figure 3.10: Exercise 5

Answer: To determine the optimal strategies for the players we shall determine if saddle point exist in this game.

		Plinky		
		P1	P2	P3
Sneaky	S1	30	−21	1
	S2	18	14	12
	S3	−6	28	4
	S4	18	6	2

Figure 3.11: Exercise 5 solution

We note that the row minima are `(-21, 12, -6, 2)`.
The column maxima are `(30,28,12)`.
So the maximin solution is `12` and the minimax solution is `12`. Therefore there is a saddlepoint solution which is the equilibrium optimum set of strategies. The value of the game is 12 (since maximin = minimax), and it is biased in favour of Sneaky.

3.2.6 Exercise

Construct a payoff matrix for the Rock–Paper–Scissors game and show it does not have a saddle point solution.

Answer:

There are many payoff matrices you could construct, as long as the ordering is maintained. One such example is shown in Figure 3.12.

		Player 2		
		R	P	S
Player 1	R	0	−1	1
	P	1	0	−1
	S	−1	1	0

Figure 3.12: Exercise 6: Rock-Paper-Scissors

The set of row minima are `(-1,-1,-1)`, the set of column maxima are `(1,1,1)`, and the maximin = `-1` and the minimax = `(+1)`, and so the game has no equilibrium pure strategy solution, ie it has no saddle point solution.

Chapter 4

Mixed Strategies

4.1 Two Person Zero Sum Game (Without a saddle point)

We have already observed that in game situations which have a saddle point there is an adequate theory of how best to play the game. However it is possible that there is no saddle point of a game and then it is then not possible to find its solution in terms of the pure strategies â the maximin and the minimax techniques.

Games without saddle point cannot be strictly determined. The solution to such problems calls for employing *mixed strategies* i.e. to solve such games both the players must determine an optimal mixture of strategies to find a saddle point. A mixed strategy therefore represents a combination of two or more strategies that are selected according to pre-determined probabilities. Thus in employing a mixed strategy, a player decides to mix his choices among several alternatives in a certain ratio with a random element determined by probability theory.

Consider this with the help of an example

4.1.1 Mixed Strategies: Example

Example Determine the optimal strategies for the players and value of the game from the following payoff matrix.

	B1	B2
A1	8	-7
A2	-6	4

Figure 4.1: Payoff matrix for Game with no saddlepoint

The given problem does not have a saddle point. Therefore, the method of saddlepoints is not sufficient to determine optimal strategies. If A plays `A1` then B would play `B2` while if A plays `A2` then B would choose `B1` to play. So if B knows what choice A will make then B can ensure that she gains by choosing a strategy, opposite to the one desired by A. Thus it is important for A to make it difficult for B to guess as to what choice she is going to make. Similarly, B would like to make it very difficult for A to assess the

strategy she is likely to adopt. We do this by setting the randomisation of strategies in such a proportion that the opposing player is ambivalent between selecting any strategy. We do this by ensuring the payoff to the opposing player is always constant irrespective of whichever strategy is chosen.

Now suppose that A plays strategy `A1` with probability p_1 and plays strategy `A2` with probability $p_2 = 1 - p_1$. If B plays strategy `B1`, then B's expected payoff from selecting this strategy can be determined by the figures in the first column of the payoff matrix as (remembering this is a zero sum game so B's payoffs are the negative of A's payoffs):

$$\begin{aligned} \text{Expected payoff to B (if B plays B1)} &= -8p_1 + 6p_2 \\ &= -8p_1 + 6(1 - p_1) \end{aligned}$$

Similarly, If B plays `B2`, the expected payoff (to B) can be determined as:

$$\begin{aligned} \text{Expected pay off (if B plays B2)} &= 7p_1 - 4p_2 \\ &= 7p_1 - 4(1 - p_1) \end{aligned}$$

Now, if player A wants to determine a value of p_1 so that the expected pay-off to B is independent of the choice that B might make (ie is the same *irrespective* of the strategy adopted by B), then we must have equality between these two payoffs:

$$\begin{aligned} -8p_1 + 6(1 - p_1) &= 7p_1 - 4(1 - p_1) \\ -8p_1 + 6 - 6p_1 &= 7p_1 - 4 + 4p_1 \\ \Rightarrow p_1 &= \frac{10}{25} = \frac{2}{5} \\ \Rightarrow p_2 &= 1 - p_1 = \frac{3}{5} \end{aligned}$$

So player A would do best to randomise her plays between the two strategies in the ratio of `2:3` with probabilities $2/5$ for `A1` and $3/5$ for `A2`.

Using this strategy we can work out the expected payoff to B. So if player B chooses strategy `B1` then player B has expected payoff of:

$$\begin{aligned} \sigma_B(A1, B1) \times \frac{2}{5} + \sigma_B(A2, B1) \times \frac{3}{5} \\ -8 \times \frac{2}{5} + 6 \times \frac{3}{5} \\ = \frac{2}{5}. \end{aligned}$$

And if player B chooses strategy `B2` the expected payoff for player B would be:

$$\begin{aligned}\sigma_B(A1, B2) \times \frac{2}{5} + \sigma_B(A2, B2) \times \frac{3}{5} \\ = 7 \times \frac{2}{5} - 4 \times \frac{3}{5} \\ = \frac{2}{5}.\end{aligned}$$

and so it is the same irrespective of the strategy selected, ie player B can gain no advantage from trying to guess which strategy player A will choose. Indeed, Player A could announce her mixed strategy *in advance* of any move by Player B and player B would still not be able to benefit at all from this knowledge!

So, irrespective of whichever strategy player B chooses, player B's payoff is fixed at $\frac{2}{5}$.

Now, we determine the mixed strategy that B should choose in a similar manner. Let us suppose B plays strategy `B1` with probability q_1 and strategy `B2` with probability $q_2 = 1 - q_1$. Then,

$$\begin{aligned}\text{Expected payoff to A (if A plays } \texttt{A1}\text{)} &= 8q_1 - 7q_2 \\ &= 8q_1 - 7(1 - q_1)\end{aligned}$$

Similarly, If A plays `A2`, the expected payoff (to A) can be determined as:

$$\begin{aligned}\text{Expected pay off to A (if A plays } \texttt{A2}\text{)} &= -6q_1 + 4q_2 \\ &= -6q_1 + 4(1 - q_1)\end{aligned}$$

So again, if B chooses her probability q_1 such that the expected payoff to A is independent of the choice of whatever player A chooses to do, then we find:

$$\begin{aligned}8q_1 - 7(1 - q_1) &= -6q_1 + 4(1 - q_1) \\ 8q_1 - 7 + 7q_1 &= -6q_1 + 4 - 4q_1 \\ \Rightarrow q_1 &= \frac{11}{25} \\ \Rightarrow q_2 = 1 - q_1 &= \frac{14}{25}\end{aligned}$$

So if player B uses this strategy we can work out player A's expected payoff. So if player A happens to choose strategy `A1` then player A has expected payoff of:

$$\begin{aligned}\sigma_A(A1, B1) \times \frac{11}{25} + \sigma_A(A1, B2) \times \frac{14}{25} \\ 8 \times \frac{11}{25} - 7 \times \frac{14}{25} \\ = \frac{-2}{5}.\end{aligned}$$

And if player A chooses strategy `A2` instead then the the expected payoff for player A would be:

$$\begin{aligned}\sigma_A(A2, B1) \times \frac{11}{25} + \sigma(A2, B2) \times \frac{14}{25} \\ -6 \times \frac{11}{25} + 4 \times \frac{14}{25} \\ = \frac{-2}{5}.\end{aligned}$$

So player A should expect a payoff of -2/5 independent of which strategy she selects.

Therefore we conclude that players A and B should both use a randomised, mixed strategy of selecting their respective choices with probabilities of $(2/5, 3/5)$ for player A and $(11/25, 14/25)$ for player B. In this case, Player A has an expected loss of $-2/5$ and player B has an expected gain of $+2/5$. The value of the game is therefore $-2/5$.

This approach of dealing with a mutual solution is known as the method of *equalising strategies*.

4.2 Mixed Strategies: Formal Explanation

Definition 12 (Mixed Strategy)**.** *Let $G = (P, \Sigma, \pi)$ be a game in normal form with $P = \{P_1, \ldots, P_N\}$. A mixed strategy for Player $P_i \in P$ is a discrete probability distribution function ρ_i defined over the sample space Σ. That is, we can define a discrete probability space $(\Sigma_i, \mathcal{F}_{\Sigma_i}, \rho_i)$ where Σ_i is the discrete sample space, $\mathcal{F}_{\Sigma_i}$ is the power set of Σ_i and ρ_i is the discrete probability function that assigns probabilities to events in $\mathcal{F}_{\Sigma_i}$.*

Remark 9. *We assume that players choose their mixed strategies independently. Thus we can compute the probability of a strategy element as a product of individual probabilities, $(\boldsymbol{\sigma}^1, ..., \boldsymbol{\sigma}^N) \in \Sigma$ as:*

$$\rho(\boldsymbol{\sigma}^1, \ldots, \boldsymbol{\sigma}^N) \equiv \rho(\boldsymbol{\sigma}^1)\rho(\boldsymbol{\sigma}^2) \ldots \rho(\boldsymbol{\sigma}^N) \tag{4.1}$$

Using this, we can define a discrete probability distribution over the sample space Σ as: $(\Sigma, \mathcal{F}_\Sigma.\rho)$. Define Π_i as a random variable that maps Σ into $\Re$ so that Π_i returns the payoff to Player P_i as a result of the random outcome $(\sigma^1, \ldots, \sigma^N)$. Therefore, the expected payoff for Player P_i for a given mixed strategy $(\rho^1 \ldots . \rho^N)$ is given as:

$$E[\Pi_i] = \sum_{\sigma^1 \in \Sigma_1} \sum_{\sigma^2 \in \Sigma_2} \cdots \sum_{\sigma^N \in \Sigma_N} \pi_i(\sigma^1, \ldots, \sigma^N)\rho^1(\sigma^1)\rho^2(\sigma^2) \ldots \rho^N(\sigma^N) \tag{4.2}$$

Definition 13 (Mixed Strategy Vectors)**.** *Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form with $\boldsymbol{P} = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots \sigma_{n_i}^i\}$ be the set of n_i pure strategies for player P_i. To any mixed strategy for Player P_i we may associate a vector $\boldsymbol{x}^i = [x_1^i, \ldots x_{n_i}^i]^T$. Provided that it satisfies the properties:*

$$\begin{aligned}x_j^i \geq 0 \text{ for } j = 1, \ldots, n_i \\ \sum_{j=1}^{n_i} x_j^i = 1\end{aligned}$$

These two properties ensure we are defining a mathematically correct probability distribution over the strategies set Σ_i.

Definition 14 (Player Mixed Strategy Space)**.** *Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form with $\boldsymbol{P} = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots \sigma_{n_i}^i\}$ be the pure strategy space for player i. Then the set:*

$$\Delta_{n_i} = \left\{ [x_1^i, \ldots, x_{n_i}^i]^T \in \Re^{n \times 1} s.t. \sum_{j=1}^{n_i} x_j^i = 1; x_j^i \geq 0, j = 1, \ldots, n_i \right\} \tag{4.3}$$

is the mixed strategy space in n_i dimensions for Player P_i.

Definition 15 (Pure Strategy)**.** *Let Σ_i be the strategy set for Player P_i in a game. If $\Sigma_i = \{\sigma_1^i, \ldots \sigma_{n_i}^i\}$ then the standard basis vectors, $\boldsymbol{e}^i = \{e_j^i : j = 1, \ldots, n_i\} \in \Delta_{n_i}$ are the* pure strategies *in Δ_{n_i} and e_j^i corresponds to a pure strategy choice $\sigma_j^i \in \Sigma_i$.*

Hence, a pure strategy can be thought of as the limit of a mixed strategy when the full probability mass is concentrated on just one of the choices.

Definition 16 (Mixed Strategy Space)**.** *Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form with $\boldsymbol{P} = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots \sigma_{n_i}^i\}$. Then the* mixed strategy space *for the game G is the set*

$$\Delta = \Delta_{n_1} \times \Delta_{n_2} \times \cdots \times \Delta_{n_N} \tag{4.4}$$

The values of the mixed strategy have a neat geometrical interpretation as points on a *simplex*. In 3 dimension, if the components of a mixed strategy vector have the values (x_1, x_2, x_3) then the simplex is the 2D surface defined by all values of $\{x_1, x_2, x_3\}$ such that

$$\begin{aligned} 0 \leq x_1 \leq 1 \\ 0 \leq x_2 \leq 1 \\ 0 \leq x_3 \leq 1 \\ x_1 + x_2 + x_3 = 1 \end{aligned}$$

Definition 17 (Mixed Strategy Payoff Function)**.** *Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form with $\boldsymbol{P} = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots, \sigma_{n_i}^i\}$. The expected payoff can be written in terms of a tuple of mixed strategy vectors, $(\boldsymbol{x}^1, \boldsymbol{x}^2, \ldots, \boldsymbol{x}^N)$:*

$$u_i(\boldsymbol{x}^1, \boldsymbol{x}^2, \ldots, \boldsymbol{x}^N) = \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_N=1}^{n_N} \pi_i(\sigma_{i_1}^1, \ldots, \sigma_{i_N}^N) \boldsymbol{x}_{i_1}^1 \boldsymbol{x}_{i_2}^2 \ldots, \boldsymbol{x}_{i_N}^N \tag{4.5}$$

where $\boldsymbol{x}_j^i$ is the j-th element of the $i-th$ player's mixed strategy vector $\boldsymbol{x}^i$. The function $u_i : \Delta \to \Re$ defined in (4.5) is the mixed strategy payoff function *for player P_i.*

Example For the Rock-Paper-Scissors game, since each player has 3 strategies, then $n = 3$ and Δ_3 consists of those vectors $[x_1, x_2, x_3]^T$ such that $x_1, x_2, x_3 \geq 0$ and $x_1 + x_2 + x_3 = 1$. For example, the

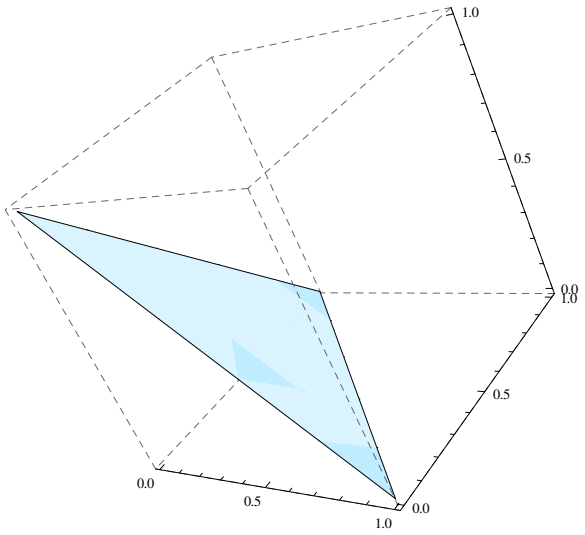

Figure 4.2: Simplex in 3 dimensions

vectors

$$\begin{bmatrix} 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}$$

is a mixed strategy vector that instructs the player to play each of the three strategies $1/3$ of the time, randomly.

We can now finally properly define the important concept of a *Nash equilibrium*. Intuitively a Nash equilibrium is a mixed strategy choice in which no player has any beneficial reason to deviate unilaterally from her mixed strategy.

Definition 18 (Nash Equilibrium)**.** *Let $G = (\boldsymbol{P}, \Sigma, \pi)$ be a game in normal form with $\boldsymbol{P} = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots, \sigma_{n_i}^i\}$. A* Nash equilibrium *is a tuple of mixed strategies $(\boldsymbol{x}^{1*}, \ldots, \boldsymbol{x}^{N*}) \in \Delta$ such that for all $i = 1, \ldots, N$:*

$$u_i(\boldsymbol{x}^{1*}, \ldots, \boldsymbol{x}^{i*}, \ldots, \boldsymbol{x}^{N*}) \geq u_i(\boldsymbol{x}^{1*}, \ldots, \boldsymbol{x}^{i}, \ldots, \boldsymbol{x}^{N*}) \tag{4.6}$$

for all $\boldsymbol{x}^i \in \Delta_{n_i}$.

The following proposition tells us that the expected payoffs in a mixed strategy game can be conveniently calculated using the mixed strategy vectors similar to the previous calculation using pure strategies and the unit basis vectors.

Proposition 3. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}, \boldsymbol{B})$ be a two-player matrix game. Let $\Sigma = \Sigma_1 \times \Sigma_2$ where $\Sigma_1 = \{\sigma_1^1, \ldots, \sigma_m^1\}$ and $\Sigma_2 = \{\sigma_1^2, \ldots, \sigma_n^2\}$. Let $\boldsymbol{x} \in \Delta_m$ and $\boldsymbol{y} \in \Delta_n$ be mixed strategies for Player 1 and*

Player 2 respectively. Then:

$$u_1(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{y} \tag{4.7}$$

$$u_2(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{y} \tag{4.8}$$

The proof is simply to expand the definition using matrix multiplication and collecting terms.

Proof. Let $\boldsymbol{x} = [x_1, \ldots, x_m]^T$ and $\boldsymbol{y} = [y_1, \ldots, y_n]^T$. We know that $\pi_1(\sigma_i^1, \sigma_j^2) = \boldsymbol{A}_{ij}$. Matrix multiplication therefore yields

$$\begin{aligned}
\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{y} &= [x_1, \ldots, x_m] \begin{bmatrix} A_{11} & \ldots & A_{1n} \\ \vdots & \ddots & \vdots \\ A_{m1} & \ldots & A_{mn} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} \\
&= [x_1, \ldots, x_m] \begin{bmatrix} \sum_{j=1}^n A_{1,j} y_j \\ \sum_{j=1}^n A_{2,j} y_j \\ \vdots \\ \sum_{j=1}^n A_{m,j} y_j \end{bmatrix} \\
&= \sum_{i=1}^m \sum_{j=1}^n A_{i,j} x_i y_j \\
&= \sum_{i=1}^m \sum_{j=1}^n \pi_1(\sigma_i^1 \sigma_j^2) x_i y_j \\
&\equiv u_1(\boldsymbol{x}, \boldsymbol{y})
\end{aligned}$$

A similar argument shows that $u_2(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{y}$. This completes the proof. □

A two-person zero-sum game $G = (P, \Sigma, \boldsymbol{A}, \boldsymbol{B})$ is said to be a finite game if both strategy sets Σ_1 and Σ_2 are finite sets. The following is one statement of one of the most well-known theorems in game theory due to John von Neumann which stems from the ability to use mixed strategies, and tells us that *every* two person zero sum game has an equilibrium solution in terms of mixed strategies.

Theorem 4. ***The Minimax Theorem.***

For every finite two-person zero-sum game,

1. *there is a number V, called the value of the game,*
2. *there is a mixed strategy for Player I such that I's average gain is at least V no matter what II does, and*
3. *there is a mixed strategy for Player II such that II's average loss is at most V no matter what I does.*

This is one form of the minimax theorem to be disussed in more detail later. If V is zero we say the game is fair. If V is positive, we say the game favours Player I, while if V is negative, we say the game favours Player II.

An equivalent way to express the theorem is

Theorem 5. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A})$ be a zero-sum game with $\boldsymbol{A} \in \Re^{m \times n}$. Then there exists a Nash equilibrium (x^*, y^*).*

The various proofs of this theorem are involved and we will consider one such proof later.

4.2.1 Solution of All 2 by 2 zero sum Matrix Games

Consider the general 2×2 game matrix

$$\boldsymbol{A} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

To solve this game (ie to find the value and at least one optimal strategy for each player) we proceed as follows.

1. Test for a saddle point.
2. If there is no saddle point, solve by finding equalising strategies.

We now prove the method of finding equalising strategies works whenever there is no saddle point by deriving the value and the optimal strategies.

Assume there is no saddle point. If $a \geq b$, then $b < d$, since otherwise b would be a saddle point. Since $b < d$, we must have $d > c$, as otherwise d is a saddle point. Continuing thus, we see that $c < a$ and $a > b$. In other words, if $a \geq b$, then $a > b < d > c < a$. By symmetry, if $a \leq b$, then $a < b > d < c > a$. This shows that:
if there is no saddle point, then either

$$a > b, b < d, d > c \text{ and } c < a,$$

or

$$a < b, b > d, d < c \text{ and } c > a.$$

Now consider the situation if Player I selects strategy `A1` with probability p and strategy `A2` with probability $(1 - p)$. Similarly Player II chooses strategy `B1` with probability q and strategy `B2` with probability $(1-q)$. Now for Player I to ensure no information can be given to player II which would allow player II to second guess player I's strategy, player I will randomly select her strategies in such a way that the payoff to player II remains the same *independent* of whatever strategy player II decides to select.

So for player II to have the same expected payoff irrespective of her choice of strategy means that we must have (recall this is a zero sum game so we have replaced all payoffs with their negatives - appropriate for the payoffs to player II)

$$p \times (-a) + (1 - p) \times (-c)c = p \times (-b) + (1 - p) \times (-d) \tag{4.9}$$

$$\Rightarrow \quad p \equiv \frac{d - c}{(a - b) + (d - c)} \tag{4.10}$$

Since there is no saddle point then either $(a-b)$ and $(d-c)$ are both positive, or both negative and hence we can be assured that $0 < p < 1$.

Using this value of p the expected payoff for player II would be

$$v = (-a) \times p + (-c) \times (1-p) = -\frac{(ad-bc)}{(a-b)+(d-c)}$$

By a similar process, we find that by equalising strategies, Player II plays her strategies `(B1,B2)` with probabilities $q, (1-q)$ in such a way that Player I's payoffs are the same irrespective of her choice of strategies. In this way we find that Player II must mix her strategies with probabilities

$$q = \frac{d-b}{(a-b)+(d-c)}$$

and in so doing constrains Player I to have an expected payoff of

$$v = \frac{(ad-bc)}{(a-b)+(d-c)}$$

which is the same as Player II's payoff in magnitude but opposite in sign (as it should be of course for a zero sum game in which one player's gain can only be at the other player's expense!). So the game has a value (v) and the players have optimal strategies as predicted by the minimax theorem.

Example Assuming the payoff matrix can be expressed in the form

$$\boldsymbol{A} = \begin{bmatrix} -2 & 3 \\ 3 & -4 \end{bmatrix}$$

then we find the optimal strategy for player I is $(7/12, 5/12)$ and for player II it is also $(7/12, 5/12)$. This gives a value of the game as

$$v = \frac{8-9}{-2-3-4-3} = \frac{1}{12}$$

and so it is a biased game in favour of Player I (the value of the game is positive).

4.3 Exercises: Mixed Strategies

We have recently explored the issues of strategic games in matrix form in which there is no saddle point equilibrium in terms of pure strategies. Instead we introduced the notion of mixed strategies and why selecting strategies randomly (from a given distribution) is the correct approach in many games. It was seen that a mixed strategy vector generalises the pure strategy vector, and it was seen how the formalism could be used to calculate the value of the game, the payoffs to individual players and how to determine strategies in two-person zero sum games. Also the famous Minimax Theorem was introduced and the concept that every 2 person zero sum game had a Nash equilibrium.

The following exercises illustrate some of these ideas. Note that the idea of using mixed strategies to make a player's choices indifferent to their competitor's choices is not restricted to zero sum games (but the Minimax Theorem is!).

To reiterate, the main idea in mixed strategies is that each player chooses to play one of their strategies, but with a given probability. The skill is in figuring out what probabilities to choose, in such a way as to nullify the effect of your opponent's choice of strategy. You do this by figuring out what are the *expected* payoffs your opponent could hope for and then choosing your own probabilities to make your opponents choices irrelevant - ie they become indifferent to any of their potential choices.

4.3.1 Exercise

What is the correct strategy to play the game "Matching Pennies"? ie for the following payoff matrix find optimal strategies for the players and find the value of the game. Note that most of our examples present the payoff matrix for the row player who we will generically assume is called Player A.

$$\boldsymbol{A} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

Figure 4.3: Exercise 1

Answer: Clearly the payoff matrix does not contain any saddle points. Therefore the players need to use mixed strategies. Using the formulae we find

$$p = \frac{A_{22} - A_{21}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{1 - (-1)}{1 + 1 - (-1 - 1)} = \frac{2}{4}$$

and

$$q = \frac{A_{22} - A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{1 - (-1)}{1 + 1 - (-1 - 1)} = \frac{2}{4}$$

The expected value of the game is

$$V = \frac{A_{11}A_{22} - A_{21}A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{1 \times 1 - (-1) \times (-1)}{1 + 1 - (-1 - 1)} = \frac{0}{4}$$

Therefore the game is fair, and each player should randomly choose heads or tails with equal probability.

4.3.2 Exercise

Solve the following zero sum 2 player game. (ie find optimal strategies for the players and find the value of the game).

$$\boldsymbol{A} = \begin{bmatrix} 4 & 1 \\ 2 & 3 \end{bmatrix}$$

Figure 4.4: Exercise 2

Answer: Clearly the payoff matrix does not contain any saddle points. Therefore the players use mixed strategies. Using the formulae we find

$$p = \frac{A_{22} - A_{21}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{3-2}{4+3-(2+1)} = \frac{1}{4}$$

and

$$q = \frac{A_{22} - A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{3-1}{4+3-(2+1)} = \frac{1}{2}$$

The expected value of the game is

$$V = \frac{A_{11}A_{22} - A_{21}A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{4 \times 3 - 2 \times 1}{4+3-(2+1)} = \frac{5}{2}$$

Therefore, player A should play her strategy A1 25% of the time, and strategy A2 75% of the time. Similarly, player B should play his two strategies 50% of the time each.

4.3.3 Exercise

Solve the following zero sum 2 player game. (ie find optimal strategies for the players and find the value of the game).

$$\boldsymbol{A} = \begin{bmatrix} 5 & 1 \\ 3 & 4 \end{bmatrix}$$

Figure 4.5: Exercise 3

Answer: Clearly the payoff matrix does not contain any saddle points. Therefore the method of saddle points is insufficient to obtain a solution. Therefore the players use mixed strategies. Using the formulae we find

$$p = \frac{A_{22} - A_{21}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{4-3}{5+4-(1+3)} = \frac{1}{5}$$

So player A should play her strategy A1 with probability 0.2 and her strategy A2 with probability 0.8. For player B, assuming he plays strategy B1 with probability q we find from the formula:

$$q = \frac{A_{22} - A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{4-1}{5+4-(1+3)} = \frac{3}{5}$$

So he plays his two strategies randomly in the ratio of $3:2$.
The expected value of the game is

$$V = \frac{A_{11}A_{22} - A_{21}A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{5 \times 4 - 1 \times 3}{5+4-(1+3)} = \frac{17}{5}$$

4.3.4 Exercise

Find the mixed strategy of the following 2 person zero sum game and the value of the game. Comment on your answer.

$$\boldsymbol{A} = \begin{bmatrix} 0 & -10 \\ 1 & 2 \end{bmatrix}$$

Answer: Using the formulae we find

$$p = \frac{2-1}{0+10+2-1} = \frac{1}{11}$$

and

$$q = \frac{2+10}{0+10+2-1} = \frac{12}{11}$$

But these are meant to be probabilites and so cannot be greater than 1! (Note, $q > 1$). So there is an issue.
The issue is that we did not check to see if there was a saddle point first, and there is. The lower left corner of this matrix is a saddle point and hence $p = 0$ and $q = 1$ and the value of the game is in fact $v = 1$.

4.3.5 Exercise

As a prelude to the next installment, consider a non zero sum game as follows:
Two Firms, A and B must decide whether to put one of their restaurants in a shopping mall.

- The strategies are to `Enter` or `Not Enter`.
- If either firm plays `Not Enter` it earns £0 profits.
- If one firm plays `Enter` and the other plays `Not Enter`, the Firm that plays `Enter` earns £300,000 per year in profits (`Not Enter` always yields £0 profits).
- If both firms choose to play `Enter`, both lose £100,000 per year as there is not enough demand for two restaurants to make positive profits.

Use the method of equalising strategies to analyse this game. (ie explore any equilibria and debate what each company should do to maximise their payoff. Note that this is not a zero sum game, but it is a symmetric game.

Answer: We can construct an appropriate payoff matrix for Company A as follows (payoffs in units of £100,000):

		Company B	
		Enter	Not Enter
Company A	Enter	−1	3
	Not Enter	0	0

Figure 4.6: Market Entry payoff matrix for company A

Similarly, the payoff matrix for Company B is given in Table 4.7

		Company B	
		Enter	Not Enter
Company A	Enter	−1	0
	Not Enter	3	0

Figure 4.7: Market Entry payoff matrix for company B

Firstly note that there are 2 strategic equilibria in the upper right and bottom left cells. But are there any mixed strategies?
To find this, we note that Company A want to choose strategy `Enter` with probability p, such that it makes company B's choice of strategy indifferent (ie the outcome would remain the same).
Therefore assuming Company A choose `Enter` with probability p and `Not Enter` with probability $1-p$, and equating Company B's payoff from choosing either of the strategies to each other (to make them an indifferent choice - ie the method of equalising strategies) gives:

$$p \times (-1) + (1-p) \times 3 \equiv p \times 0 + (1-p) \times 0 = 0$$

So we must have that p must satisfy

$$3 - 4 \times p = 0 \Rightarrow p = \frac{3}{4}.$$

By a similar argument, or noting that the game is symmetric between companies A and B, we also reason that company B should choose the `Enter` strategy with probability $q = 3/4$ as well.
Given these probabilities of mixed strategies we may calculate the expected payoff to Company A (say) as

$$(-1) \times \frac{3}{4} + 3 \times \frac{1}{4} = 0$$

which is the same payoff had strategy 2 been selected. Similarly for Company B.
Note that

- Probability both `Enter` is $(3/4) \times (3/4) = 9/16$, (ie about half the time).
- Probability A enters and B does not is $(3/4) \times (1/4) = 3/16$.
- Probability A does not enter and B does is $(1/4) \times (3/4) = 3/16$.
- Probability that neither firm Enters: $(1/4) \times (1/4) = 1/16$.
- Note that $9/16 + 3/16 + 3/16 + 1/16 = 1$

.

4.3.6 Exercise

Now consider the same situation, but suppose we change the payoff matrix so that Company A has a competitive advantage over Company B in one situation: if A is the sole entrant, it earns profits of $£400,000$. Otherwise the game is the same as before. So this is now an asymmetric game. Construct a payoff matrix and analyse the game for the best mixed strategies.

Answer: The payoff matrix can be expressed as shown in Figure 4.8:

		Company B	
		Enter	Not Enter
Company A	Enter	-1	4
	Not Enter	0	0

Figure 4.8: Asymmetric Market Entry payoff matrix for company A

Similarly, the payoff matrix for Company B remains the same as before:

The pure strategy equilibria remain the same, (`Entry, Not Entry`) or (`Not Entry, Entry`).
What happens to the mixed strategy probability?
Since Company B's payoffs have not changed, Company A's mixture probability choice for p does not change. It chooses p so that $(-1) \times p + 3 \times (1 - p) = 0$, or $3 = 4p \Rightarrow p = 3/4$, as before.
However, since Company A's payoffs have changed, B's mixture probability choice for q must change to

		Company B	
		Enter	Not Enter
Company A	Enter	-1	0
	Not Enter	3	0

Figure 4.9: Asymmetric Market Entry payoff matrix for company B

accommodate the difference.

- B chooses q so as to make A just indifferent between entering and earning $(-1) \times q + 4 \times (1 - q)$ and not entering and earning 0: $\hat{a}q + 4 - 4 \times q = 0$, or $4 = 5q$, so $q = 4/5$.
- Note that the probability that B enters the market goes up from 3/4 to 4/5. If it did not, then A would choose the pure strategy $p = 1$. (Why?)
- Company A's expected payoff remains zero: $-1 \times (4/5) + 4 \times (1/5) = 0$.

4.3.7 Exercise [O'Neill's card game]

Alice and Bob each have the playing cards A,K,Q,J from a suit from a deck of cards. Simultaneously they each show one of these cards. Alice wins if both shown cards are Aces or if there is a mismatch between the picture cards. Bob wins otherwise. If we assign a payoff of 1 to the winner and 0 to the loser, list the strategies and hence write down the payoff matrices for both Alice and Bob.
There are no pure strategy solutions to this game (check!), and so find and equilibrium solution $(\boldsymbol{x}, \boldsymbol{y})$ in mixed strategies for both players by seeking probability choices over strategies that makes their opponent indifferent between all their pure strategies.

Answer: The payoff matrix $\boldsymbol{A}$ for Alice can be expressed as:

	Bob			
	A	K	Q	J
A	1	0	0	0
K	0	0	1	1
Q	0	1	0	1
J	0	1	1	0

$\boldsymbol{A}$

Figure 4.10: Payoff matrix for Alice in the O'Neill card game

Similarly the payoff matrix $\boldsymbol{B}$ for Bob can be expressed as:
Alice is indifferent to all her pure strategies when Bob plays $\boldsymbol{y} = (y_A, y_K, y_Q, y_J)$ as follows.
If Bob plays $\boldsymbol{y}$ then the payoff to Alice for each of her strategies would be (in vector matrix notation as

per the lecture notes) $\boldsymbol{Ay}$. If the payoff for each of these strategies is determined to be the same, then we must have that that this quantity is simply proportional to the vector of '1'': $\boldsymbol{Ay} \equiv \alpha \mathbf{1}$ where α is a proportionality constant we need to determine. We can now proceed in one of two ways. This is a linear matrix equation so we can find the matrix inverse of $\boldsymbol{A}$ and multiply through to find out $\boldsymbol{y}$ in terms of α. Or we can use symmetry and solve a simpler algebraic problem. Inverse way first: Since $\boldsymbol{Ay} = \alpha \mathbf{1}$ then we must have $\boldsymbol{y} = \alpha \boldsymbol{A}^{-1}\mathbf{1}$. solving for the matrix inverse (I used mathematica but you could do it by hand if you wish!) we find

$$\boldsymbol{y} = \alpha \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1/2 & 1/2 & 1/2 \\ 0 & 1/2 & -1/2 & 1/2 \\ 0 & 1/2 & 1/2 & -1/2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \alpha \begin{bmatrix} 1 \\ 1/2 \\ 1/2 \\ 1/2 \end{bmatrix}$$

Now since these are probabilities and must add up to one, we must have $\sum_{i \in \{A,K,Q,J\}} y_i = 1$. Therefore we must have $\alpha = 2/5$ and hence finally

$$\boldsymbol{y} = (2/5, 1/5, 1/5, 1/5)^T$$

Or, by symmetry there is no difference or distinguishing characteristics between any of the picture cards and so whatever probability we play the King must also be the same for Queen and Jack. ie $y_K = y_Q = y_J = y_0$ (say). In which case expanding out $\boldsymbol{Ay} = \alpha \mathbf{1}$ gives just two equations to solve:

$$y_A = \alpha$$
$$2y_0 = \alpha$$

therefore we must have $y_A = 2y_0$ and so by normalisation, $5y_0 = 1$ and hence once again we find $\boldsymbol{y} = (2/5, 1/5, 1/5, 1/5)$.
A similar process but using $\boldsymbol{B}$ rather than $\boldsymbol{A}$ shows that Alice can choose her mixed strategy to ensure Bob is indifferent to his pure strategies if

$$\boldsymbol{x} = (2/5, 1/5, 1/5, 1/5)$$

	Bob			
	A	K	Q	J
A	0	1	1	1
K	1	1	0	0
Q	1	0	1	0
J	1	0	0	1

$\boldsymbol{B}$

Figure 4.11: Payoff matrix for Alice in the O'Neill card game

Chapter 5

Dominance

Intuitively, we say the ith row of a matrix $\boldsymbol{A} = (a_{i\cdot})$ dominates the kth row if $a_{ij} \geq a_{kj}$ for all j. We say the ith row of $\boldsymbol{A}$ strictly dominates the kth row if $a_{ij} > a_{kj}$ for all j. Similarly, the jth column of $\boldsymbol{A}$ dominates (or strictly dominates) the kth column if $a_{ij} \geq a_{ik}$ (resp. $a_{ij} > a_{ik}$) for all i.

Anything Player I can achieve using a dominated row can be achieved at least as well using the row that dominates it. Hence dominated rows may be deleted from the matrix. A similar argument shows that dominated columns may be removed. To be more precise, removal of a dominated row or column does not change the value of a game. We will formalise these ideas next.

5.1 Dominated Strategies

Definition 19 (Weak Dominance)**.** *Let $G = (P, \Sigma, \pi)$ be a game in normal form with $P = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots \sigma_N^i\}$. A mixed strategy $\boldsymbol{x}^i \in \Delta_{n_i}$ for Player P_i weakly dominates another strategy $\boldsymbol{y}^i \in \Delta_{n_i}$ for Player P_i if for all mixed strategies $\boldsymbol{z}^{-i}$ we have:*

$$u_i(\boldsymbol{x}^i, \boldsymbol{z}^{-i}) \geq u_i(\boldsymbol{y}^i, \boldsymbol{z}^{-i}) \tag{5.1}$$

and for at least one $\boldsymbol{z}^{-i}$ the inequality in this Equation is strict.

Recall the notation introduced previously, that, for example, $\boldsymbol{z}^{-i}$ means the mixed strategy for all Players EXCEPT for Player i.

Definition 20 (Strong Dominance)**.** *Let $G = (P, \Sigma, \pi)$ be a game in normal form with $P = \{P_1, \ldots, P_N\}$. Let $\Sigma_i = \{\sigma_1^i, \ldots \sigma_{n_i}^i\}$.*
A mixed strategy $\boldsymbol{x}^i \in \Delta_{n_i}$ for Player P_i strictly dominates another strategy $\boldsymbol{y}^i \in \Delta_{n_i}$ for Player P_i if for all mixed strategies $\boldsymbol{z}^{-i}$ we have:

$$u_i(\boldsymbol{x}^i, \boldsymbol{z}^{-i}) > u_i(\boldsymbol{y}^i, \boldsymbol{z}^{-i}) \tag{5.2}$$

Remark 10. *In a two player matrix game $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}, \boldsymbol{B})$ with $\boldsymbol{A}, \boldsymbol{B} \in \Re^{m \times n}$. A mixed strategy $\boldsymbol{x} \in \Delta_m$ for Player 1 weakly dominates a strategy $\boldsymbol{y} \in \Delta_m$ if for all $\boldsymbol{z} \in \Delta_n$ (ie mixed strategies for Player 2) we have:*

$$\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{z} \geq \boldsymbol{y}^T \boldsymbol{A} \boldsymbol{z} \tag{5.3}$$

and the inequality is strict for at least one $\boldsymbol{z} \in \Delta_n$. *If* $\boldsymbol{x}$ *strictly dominates* $\boldsymbol{y}$ *then we have:*

$$\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{z} > \boldsymbol{y}^T \boldsymbol{A} \boldsymbol{z} \tag{5.4}$$

for all $\boldsymbol{z} \in \Delta_m$.

Example (Bonnie and Clyde : Prisoner's Dilemma)

Recall the example for the earlier part of the course. We choose to rephrase the story as often told, as follows: Two prisoner's Bonnie and Clyde commit a bank robbery. They stash the cash and are driving around wondering what to do next when they are pulled over and arrested for a weapons violation. The police suspect Bonnie and Clyde of the bank robbery, but do not have any hard evidence. They separate the prisoners and offer them the following options to Bonnie:

1. If neither Bonnie nor Clyde confess, they will go to prison for 1 year on the weapons violation.
2. If Bonnie confesses, but Clyde does not and remains quiet, then Bonnie can go free while Clyde will go to jail for 10 years.
3. If Clyde confesses and Bonnie does not, then Bonnie will go to jail for 10 years while Clyde will go free.
4. If both Bonnie and Clyde confess, then they will go to jail for 5 years each.

A similar offer is made to Clyde. The following two-player matrix game describes the situation:

- $\boldsymbol{P} = \{Bonnie, Clyde\}$;
- $\Sigma_1 = \Sigma_2 \equiv \{\texttt{Stay Quiet}, \texttt{Confess}\}$.
- The matrices for this game are given below in figure 5.1:

$\boldsymbol{A}$:

Bonnie \ Clyde	Stay Quiet	Confess
Stay Quiet	-1	-10
Confess	0	-5

$\boldsymbol{B}$:

Bonnie \ Clyde	Stay Quiet	Confess
Stay Quiet	-1	0
Confess	-10	-5

Figure 5.1: Payoff matrices for the Bonnie and Clyde problem. Units are in the loss of years.

Thus, we see that if Bonnie does not confess and Clyde does (row 1, column 2), then Bonnie loses 10 years and Clyde loses 0 years.

We will now show that the strategy `Confess` dominates `Stay Quiet` for Bonnie. Pure strategies correspond to standard basis vectors. Thus we are claiming that $\boldsymbol{e}_2$ strictly dominates $\boldsymbol{e}_1$ for Bonnie. We can use remark 10 to see that we must show:

$$\boldsymbol{e}_2^T \boldsymbol{A}\boldsymbol{z} > \boldsymbol{e}_1^T \boldsymbol{A}\boldsymbol{z} \tag{5.5}$$

We know that $\boldsymbol{z}$ is a mixed strategy. Therefore this means that:

$$\boldsymbol{z} = \begin{bmatrix} z_1 \\ \\ z_2 \end{bmatrix}$$

and also $z_1 + z_2 = 1$ and $z_1, z_2 \geq 0$. For simplicity, let us define:

$$\boldsymbol{z} = \begin{bmatrix} z \\ \\ 1-z \end{bmatrix}$$

with $z \geq 0$.

Therefore,

$$\begin{aligned} \boldsymbol{e}_2^T \boldsymbol{A}\boldsymbol{z} &= [0,1] \begin{bmatrix} -1 & -10 \\ \\ 0 & -5 \end{bmatrix} \begin{bmatrix} z \\ \\ (1-z) \end{bmatrix} \\ &= [0,-5] \begin{bmatrix} z \\ \\ (1-z) \end{bmatrix} \\ &= 5z - 5. \end{aligned}$$

and

$$\begin{aligned} \boldsymbol{e}_1^T \boldsymbol{A}\boldsymbol{z} &= [1,0] \begin{bmatrix} -1 & -10 \\ \\ 0 & -5 \end{bmatrix} \begin{bmatrix} z \\ \\ (1-z) \end{bmatrix} \\ &= [-1,-10] \begin{bmatrix} z \\ \\ (1-z) \end{bmatrix} \\ &= 9z - 10. \end{aligned}$$

Therefore substituting back into equation (5.5) gives us the inequality

$$5z - 5 > 9z - 10$$

This means generally then that for the inequality to hold we must have $4z < 5$, or $z < 5/4$. Now since we know that $0 < z < 1$, then clearly it will satisfy the requirement that $z < 5/4$. Therefore we have shown

that inequality (5.5) is satisfied. Or $\boldsymbol{e}_2$ strictly dominates $\boldsymbol{e}_1$ for Bonnie.

Exercise Show that `Confess` strictly dominates `Stay Quiet` for Clyde.

Strict dominance can be extremely useful for identifying pure Nash equilibria. This is especially true in matrix games. This is summarized in the following theorems.

Theorem 6. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}, \boldsymbol{B})$ be a two player matrix game with $\boldsymbol{A}, \boldsymbol{B} \in \Re^{m \times n}$. If*

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_k > \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_k \tag{5.6}$$

for $k = 1, \ldots, n$, then $\boldsymbol{e}_i$ strictly dominates $\boldsymbol{e}_j$ for Player 1.

Remark 11. . *We know that $\boldsymbol{e}_i^T \boldsymbol{A}$ is the i-th row of $\boldsymbol{A}$. Theorem 6 says: if every element in $\boldsymbol{A}_{i,.}$ (the i-th row of $\boldsymbol{A}$) is greater than its corresponding element in $\boldsymbol{A}_{.,j}$, (the j-th row of $\boldsymbol{A}$), then Player 1's i-th strategy strictly dominates Player 1's j-th strategy.*

Proof. For all $k = 1, \ldots n$ we know that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_k > \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_k.$$

Suppose that $z_1, \ldots z_n \in [0, 1]$ with $z_1 + \cdots + z_n = 1$. Then for each z_k we know that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_k z_k > \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_k z_k$$

for $k = 1, \ldots, n$. This implies that:

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_1 z_1 + \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_2 z_2 + \cdots + \boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{e}_n z_n > \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_1 z_1 + \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_2 z_2 + \cdots + \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{e}_n z_n$$

Factoring gives

$$\boldsymbol{e}_i^T \boldsymbol{A} (\boldsymbol{e}_1 z_1 + \boldsymbol{e}_2 z_2 + \cdots + \boldsymbol{e}_n z_n) > \boldsymbol{e}_j^T \boldsymbol{A} (\boldsymbol{e}_1 z_1 + \boldsymbol{e}_2 z_2 + \cdots + \boldsymbol{e}_n z_n).$$

Defining $\boldsymbol{z} = [z_1 \boldsymbol{e}_1 + z_2 \boldsymbol{e}_2 + \cdots + z_n \boldsymbol{e}_n] \equiv [z_1, z_2, \ldots, z_n]^T$ allows the above to be written as

$$\boldsymbol{e}_i^T \boldsymbol{A} \boldsymbol{z} > \boldsymbol{e}_j^T \boldsymbol{A} \boldsymbol{z}.$$

Since the values of z_i were chosen arbitrarily, the above holds for all choices of mixed strategy $\boldsymbol{z}$ and hence is the definition of strict dominance of strategy i over strategy j. Thus $\boldsymbol{e}_i$ dominates $\boldsymbol{e}_j$ according to the definition 20. □

Theorem 7. *Let $G = (\boldsymbol{P}, \Sigma, \boldsymbol{A}, \boldsymbol{B})$ be a two player matrix game with $\boldsymbol{A}, \boldsymbol{B} \in \Re^{m \times n}$. Suppose pure strategy $\boldsymbol{e}_j \in \Delta_m$ for Player 1 is strictly dominated by pure strategy $\boldsymbol{e}_i \in \Delta_m$. If $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ is a Nash equilibrium, then $x_j^* = 0$. Similarly, if pure strategy $\boldsymbol{e}_j \in \Delta_n$ for Player 2 is strictly dominated by pure strategy $\boldsymbol{e}_i \in \Delta_n$, then $y_j^* = 0$.*

In other words, if choice j is strictly dominated by another choice i, then in any mixed strategy, the probability mass for option j must necessarily be zero.

Proof. We will prove the theorem for Player 1; the proof for Player 2 is completely analogous. We will proceed by contradiction. Suppose that $x_j^* > 0$. We know:

$$\boldsymbol{e}_i^T \boldsymbol{A}\boldsymbol{y} > \boldsymbol{e}_j^T \boldsymbol{A}\boldsymbol{y}$$

because $\boldsymbol{e}_i$ strictly dominates $\boldsymbol{e}_j$. We can express the Nash equilibrium $\boldsymbol{x}^*$ in terms of the pure strategies and the mixing probabilities of each pure strategy and therefore expand as follows:

$$\boldsymbol{x}^{*T} \boldsymbol{A}\boldsymbol{y} \equiv (x_1^* \boldsymbol{e}_1^T + \cdots + x_m^* \boldsymbol{e}_m^T)\boldsymbol{A}\boldsymbol{y}. \tag{5.7}$$

Since we are assuming that $x_j^* > 0$ and since $\boldsymbol{e}_j$ is dominated by $\boldsymbol{e}_i$, then we know for sure that $\boldsymbol{e}_i^T \boldsymbol{A}\boldsymbol{y} > \boldsymbol{e}_j^T \boldsymbol{A}\boldsymbol{y}$, and therefore it must follow that

$$x_j^* \boldsymbol{e}_i^T \boldsymbol{A}\boldsymbol{y} > x_j^* \boldsymbol{e}_j^T \boldsymbol{A}\boldsymbol{y}.$$

Therefore from (5.7) we must have

$$(x_1^* \boldsymbol{e}_1^T + \cdots + x_i^* \boldsymbol{e}_i^T + \cdots + x_j^* \boldsymbol{e}_i^T + \cdots + x_m^* \boldsymbol{e}_m^T)\boldsymbol{A}\boldsymbol{y} > (x_1^* \boldsymbol{e}_1^T + \cdots + x_i^* \boldsymbol{e}_i^T + \cdots + x_j^* \boldsymbol{e}_j^T + \cdots + x_m^* \boldsymbol{e}_m^T)\boldsymbol{A}\boldsymbol{y},$$

or

$$(x_1^* \boldsymbol{e}_1^T + \cdots + (x_i^* + x_j^*)\boldsymbol{e}_i^T + \cdots + 0 \times \boldsymbol{e}_j^T + \cdots + x_m^* \boldsymbol{e}_m^T)\boldsymbol{A}\boldsymbol{y} > \boldsymbol{x}^{*T} \boldsymbol{A}\boldsymbol{y}.$$

So if we define a vector $\boldsymbol{z}$ to be such that

$$\boldsymbol{z} = (x_1^*, x_2^*, \ldots, (x_i^* + x_j^*), \ldots, 0, \ldots, x_m^*),$$

then we must have

$$\boldsymbol{z}\boldsymbol{A}\boldsymbol{y} > \boldsymbol{x}^{*T} \boldsymbol{A}\boldsymbol{y}.$$

But, since $\boldsymbol{x}^*$ is meant to be a Nash equilibrium, there can be no other vector such that $\boldsymbol{z}\boldsymbol{A}\boldsymbol{y} > \boldsymbol{x}^{*T} \boldsymbol{A}\boldsymbol{y}$ from the definition of what a Nash equilibrium is. Therefore we have a contradiction and our original assumption that $x_j^* > 0$ must be false. Hence we must have that $x_j^* = 0$. □

Example Consider the 2 player zero sum game defined by the payoff matrix:

$$\boldsymbol{A} = \begin{bmatrix} 2 & 0 & 4 \\ 1 & 2 & 3 \\ 4 & 1 & 2 \end{bmatrix}$$

The last column is dominated by the middle column, so deleting the final column gives:

$$\begin{bmatrix} 2 & 0 \\ 1 & 2 \\ 4 & 1 \end{bmatrix}$$

Now the top row is dominated by the bottom row (note this was not the case in the original matrix) so deleting the top row gives:

$$\begin{bmatrix} 1 & 2 \\ 4 & 1 \end{bmatrix}$$

This 2×2 matrix has no saddle point, and so the optimal mixed strategies may be determined as $p = 3/4$, $q = 1/4$ and the game value is $v = 7/4$. So, Player I's optimal strategy is to play $(A1, A2, A3)$ with mixture $(0, 3/4, 1/4)$ and Player 2's strategy is to play his options $(B1, B2, B3)$ with mixture $(1/4, 3/4, 0)$.

Remark 12. *A row (column) may also be removed if it is dominated by a probability combination of other rows (columns).*

If for some $0 < p < 1$, and if $pa_{i_1,j} + (1-p)a_{i_2,j} \geq a_{k,j} \quad \forall j$, then the kth row is dominated by the mixed strategy that chooses row i_1 with probability p and row i_2 with probability $1 - p$. Player I can do at least as well using this mixed strategy instead of choosing row k. (In addition, any mixed strategy choosing row k with probability p_k may be replaced by the one in which k's probability is split between i_1 and i_2. A similar argument may be used for columns.

Consider the matrix:

$$\boldsymbol{A} = \begin{bmatrix} 0 & 4 & 6 \\ 5 & 5 & 4 \\ 9 & 7 & 3 \end{bmatrix}$$

The middle column is dominated by the outside columns taken with probability 1/2 each. With the central column deleted, the middle row is dominated by the combination of the top row with probability 1/3 and the bottom row with probability 2/3.

The reduced matrix,

$$\begin{bmatrix} 0 & 6 \\ 9 & 3 \end{bmatrix},$$

is easily solved. The game value is $V = 54/12 = 9/2$.

5.2 Exercises on Dominance

PRINCIPLES OF DOMINANCE The RULES to follow are:

Rule 1: If all the elements in a row (say i-th row) of a pay off matrix are less than or equal to the corresponding elements of another row (say the j-th row) then the player A will never choose the i-th strategy then we say the i-th strategy is dominated by the j-th strategy and we can delete the i-th row.

Rule 2: If all the elements in a column (say the r-th column) of a payoff matrix are greater than or equal to the corresponding elements of the other column (say s-th column) then the player B will never choose the r-th strategy or in the other words the r-th strategy is dominated by the s-th strategy and we can delete the r-th column .

Rule 3: A pure strategy may be dominated if it is inferior to a weighted average of two or more other pure strategies.

Now, consider some simple examples.

5.2.1 Exercise

Given the zero sum payoff matrix for player A in figure 5.2, obtain the optimum strategies for both the players and determine the value of the game.

		Player B		
		B1	B2	B3
Player A	A1	6	−3	7
	A2	−3	0	4

Figure 5.2: Exercise 1: Payoff matrix for Player A in a zero sum game.

Answer: For the column player, strategy B3 is always a bad choice as it is dominated by either of the other strategies, B1 or B2. (Remember, from Player B's perspective, this is the payoff matrix for Player A and so the payoffs need to be reversed as it is a zero sum game.) Therefore we can eliminate the third column and get the reduced 2 by 2 payoff matrix:

		Player B	
		B1	B2
Player A	A1	6	−3
	A2	−3	0

In this reduced matrix, the minimax value is zero, and the maximin value is -3. Therefore they are not the same and so a saddle point does not exist, so a Nash equilibrium must be in terms of a mixed strategy. If the probability of A choosing strategy A1 is p and choosing A2 is $(1-p)$, then the principle of equalising strategies gives the equality:

$$6p - 3(1-p) = -3p + 0 \times (1-p) \;\Rightarrow\; p = 1/4.$$

Similarly, if B choose B1 with probability q and chooses B2 with probability $(1-q)$ then to ensure balancing of strategies we must have

$$6q - 3 \times (1-q) = -3q + 0 \times (1q) \;\Rightarrow\; q = 1/4.$$

Therefore the strategies of the two players are to play their respective strategies using the strategy vectors

$$\begin{aligned} A;&\quad (x_1, x_2) = (1/4, 3/4), \\ B:&\quad (y_1, y_2, y_3) = (1/4, 3/4, 0). \end{aligned}$$

Using these mixed strategies we can substitute in the values of p and q into any of the expected payoffs, and work out the value of the game which is :

$$(6p - 3 \times (1-p)) \equiv (3q + 0 \times (1-q)) = -3/4.$$

5.2.2 Exercise

In an election campaign, the strategies adopted by the ruling and opposition party along with pay-offs (ruling party's share in scaled votes polled) are given in Table 5.1.

		Opposition Party Strategies		
		Target Cities	Target Towns	Target Rural
Ruling Party	Target Cities	55	40	35
	Target Towns	70	70	55
	Target Rural	75	55	65

Table 5.1: Political Strategies and Payoffs.

Assume a zero sum game. Find optimum strategies for both parties and expected payoff to ruling party.

Answer: Note that row 1 is dominated by row 2 (and in fact row 3) for the ruling party, and so they can remove that strategy from consideration. In addition, note that column 1 is dominated by column 2 (and indeed column 3) and so the first column can be removed from consideration by the opposition party.This leaves a two-by-two payoff matrix that we know can be solved.

		Opposition Party Strategies	
		Target Towns	Target Rural
Ruling Party	Target Towns	70	55
	Target Rural	55	65

Table 5.2: Political Strategies and Payoffs - reduced matrix.

Considering the strategy from the ruling party perspective, if they choose strategy `Target Towns` with probability p and strategy `Target Rural` with probability $(1-p)$ then if the opposition choose the strategy `Target Towns` the ruling party can expect a payoff of:

$$p \times 70 + (1-p) \times 55 \equiv 15p + 55.$$

However, if the Opposition were to choose strategy `Target rural` then the Ruling Party could expect a payoff of

$$p \times 55 + (1-p) \times 65 \equiv -10p + 65.$$

So if the Ruling Party want to choose a strategy with a payoff which cannot be influenced by the Opposition, then we can equate the above two payoffs and find

$$15p + 55 = -10p + 65 \;\Rightarrow\; 25p = 10 \;\Rightarrow\; p = 2/5.$$

Similarly the Opposition Party will choose strategy `Target Towns` with probability q and `Target Rural` with probability $(1-q)$. So the Opposition's expected payoffs depending on whether the Ruling Party chooses strategy `Towns` vs `Rural` would be

$$q \times 70 + (1-q) \times 55 \equiv 15q + 55,$$

and

$$q \times 55 + (1-q) \times 65 \equiv -10q + 65.$$

So if Opposition want to have a payoff uninfluenced by the Ruling Party, they should choose q such that

$$15q + 55 = -10q + 65 \;\Rightarrow\; 25q = 10 \;\Rightarrow q = 2/5.$$

The expected gain to the Ruling Party can be found by substituting the derived value of p into either of the expected payoff equations, giving, for example,

$$V = \frac{2}{5} \times 15 + 55 = 61.$$

Similarly, the expected loss to the Opposition Party is

$$\frac{2}{5} \times (-10) + 65 = 61.$$

(Recall this is a loss and so is negative, for the Opposition Party).

5.2.3 Exercise

Two firms, A and B manufacture scooters and they control the market in their local area. So if one makes a model improvement in one year and their competitor does not, then they will gain market share at the expense of their competitor. If they make model changes of the same type for this market segment in the same year, their respective market shares remain constant. Likewise, if neither makes model changes, then also their market shares remain constant. The pay-off matrix in terms of increased/decreased percentage market share under different possible conditions is given in Table 5.3:

		B Strategies		
		No Change	Minor Change	Major Change
A strategies	No Change	0	−4	−10
	Minor Change	3	0	5
	Major Change	8	1	0

Table 5.3: Payoffs for changes in scooter manufacturing.

1. Find the value of the game.
2. What change should manufacturer A consider if this information is available only to itself?

Answer:

1. Clearly, the game has no saddle point (check that minimax does not equal maximin). Making use of dominance principle, since the first row is dominated by the third row, we can delete the first row as a viable strategy. Similarly, the first column is dominated by the second column and hence we can delete the first column as a rational strategy. The reduced pay-off matrix will therefore be as shown in Table 5.4:

		B Strategies	
		Minor Change	Major Change
A strategies	Minor Change	0	5
	Major Change	1	0

Table 5.4: Reduced matrix game payoffs for changes in scooter manufacturing.

As this reduced matrix does not possess a saddle point, we can determine the probabilities for each manufacturer's mixed strategies and we find:

$$p = \frac{A_{22} - A_{21}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{0 - 1}{0 + 0 - (5 + 1)} = \frac{1}{6},$$

and

$$q = \frac{A_{22} - A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{0 - 5}{0 + 0 - (5 + 1)} = \frac{5}{6}.$$

The expected value of the game is

$$V = \frac{A_{11}A_{22} - A_{21}A_{12}}{A_{11} + A_{22} - (A_{21} + A_{12})} = \frac{0 \times 0 - 1 \times 5}{0 + 0 - (5 + 1)} = \frac{5}{6}.$$

So Manufacturer A's mixed strategy vector is

$$(x_1, x_2, x_3) = (0, 1/6, 5/6),$$

and Manufacturer B's mixed strategy vector would be

$$(y_1, y_2, y_3) = (0, 5/6, 1/6).$$

2. Therefore, Manufacturer A should consider to have major change with probability 5/6 and of minor change with probability 1/6 and choose randomly with these proportions year on year. Even if Manufacturer B does not know tha value of the game or the probabilities, Manufacturer A would not do better by switching their strategy since they would likely get a reduced payoff.

5.2.4 Exercise

Now consider a NON zero sum game. This is an extension of the Battle of the Sexes game where both players desire an outcome but there is no single saddle point solution.

The game payoff bimatrix is shown in Table 5.5. The first number in each cell refers to Player A's payoff and the second number refers to Player B's actual payoff. Player A has 2 choices, A1 and A2. Player B has 3 choices: B1,B2,B3.

		Player B		
		B1	B2	B3
Player A	A1	$4, 3$	$0, 0$	$2, 1$
	A2	$0, 0$	$3, 4$	$1, 1$

Table 5.5: Bimatrix Payoffs for Players A and B.

Show that Choice `B3` for Player B is dominated by his other two choices.

Answer: In this case it is not clear that either choices B1 or B2 uniquely dominates B3 for Player B. However, consider the mixed strategy of 50% of choice B1 and 50% of choice B2. Then this mixed strategy would give an 'equivalent' payoff of $(1.5, 2)$ to Player B. Note that this strictly dominates the payoff if choice B3 is followed (ie $((1, 1))$). Therefore, any mixed strategy which involves a non-zero contribution of choice B3 could be improved by replacing B3 with 50% of B1 and 50% of B2. Therefore the equivalent game reduces to a standard Battle of the Sexes game which has 2 equilibrium points.

5.2.5 Exercise

How do we find out if a mixture of strategies dominates a pure strategy? This question shows how to evaluate conditions on the probabilities of choosing strategies.

Consider the following game payoff bimatrix shown in Table 5.6. The first number in each cell refers to Player A's payoff and the second number refers to Player B's payoff. Player A has 2 choices, `u` and `d`. Player B has 3 choices: `L`,`M`,`R`.

		Player B		
		L	M	R
Player A	u	$2, 4$	$3, 0$	$1, -1$
	d	$3, 2$	$10, 3$	$0, 4$

Table 5.6: Bimatrix Payoffs for Players A and B.

a) We suspect that the action `M` of Player B is strictly dominated by a mixed strategy involving her other two actions *provided* she appropriately chooses the probability ρ of randomly selecting `L` and probability $(1 - \rho)$ of choosing action `R`. What range of values of ρ will ensure strict domination of `M`?
{ Hint: Work out the expected payoff to player B using the mixed strategy with unknown mixing probability ρ and see what happens if we insist these payoffs have to be greater than the corresponding payoffs in `M`.}

b) Hence determine the optimal strategies for each of the players.

Answer:

a) The expected payoff to player B using the mixed strategy needs to be greater than the payoffs she could gain if she were to choose action `M`. Therefore this gives us 2 equations to solve as follows:

$$\rho \times 4 - (1 - \rho) \times 1 > 0 \;\Rightarrow\; \rho > 1/5$$
$$\rho \times 2 + (1 - \rho) \times 4 > 3 \;\Rightarrow\; \rho < 1/2$$

Therefore we must have $1/5 < \rho < 1/2$. Any value of ρ in this range will ensure action `M` is strictly dominated by a mixed strategy combination of `L` and `R`.

b) Since no dominated strategy ever takes part in an optimal strategy, we note that the following remaining 2×2 payoff matrix has no Nash equilibria in pure strategies:

		B: L	B: R
A	u	(2, 4)	(1, −1)
	d	(3, 2)	(0, 4)

So, assuming player A chooses u with probability p and hence chooses d with probability $(1-p)$, and player B chooses L with probability q and R with probabiliy $(1-q)$, then we can use the method of equalising strategies to determine what the optimal values of p and q should be.
So to make sure that the column player, B, has no advantage in choosing her strategies, then the row player would choose p such that the expected payoffs to the column player are constant. This leads to the equation:

$$\begin{aligned} 4p + 2(1-p) &= -p + 4(1-p) \\ \Rightarrow 2p + 2 &= 4 - 5p \\ \Rightarrow 7p &= 2 \\ \Rightarrow p &= 2/7. \end{aligned}$$

Similarly, to ensure the row player cannot benefit from guessing which column player B will choose, player B chooses q in such a way to equalise the payoffs to the strategies chosen by player A. This means q satisfies:

$$\begin{aligned} 2q + (1-q) &= 3q + 0 \\ \Rightarrow 2p + 2 &= 4 - 5p \\ \Rightarrow q + 1 &= 3q \\ \Rightarrow q &= 1/2. \end{aligned}$$

Therefore player A's strategy is a mixed strategy with components $(2/7, 5/7)$, and B's strategy is also a mixed strategy with components $(1/2, 0, 1/2)$ (note, the zero is significant!!).

5.2.6 Exercise: Choose a Number

Each player simultaneously chooses a number from $\{1, 2, \ldots, 10\}$ and writes it down on a piece of paper; then the players compare the two numbers. If the numbers differ by one, the player with the higher number pays the other player £2. If the players' choices differ by two, the player with the lower number pays £1 to the other player. If the numbers differ by 3 or more then the player with the highest number pays £2 to the other player. In the event of a tie, no money changes hands.

Construct a game matrix which describes this game, and evaluate what each player's optimum strategies would be. **[10 marks]**

Answer: The payoff matrix for the game is:

$$\begin{bmatrix} 0 & 2 & -1 & 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ -2 & 0 & 2 & -1 & 2 & 2 & 2 & 2 & 2 & 2 \\ 1 & -2 & 0 & 2 & -1 & 2 & 2 & 2 & 2 & 2 \\ -2 & 1 & -2 & 0 & 2 & -1 & 2 & 2 & 2 & 2 \\ -2 & -2 & 1 & -2 & 0 & 2 & -1 & 2 & 2 & 2 \\ -2 & -2 & -2 & 1 & -2 & 0 & 2 & -1 & 2 & 2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -2 & -2 & -2 & -2 & -2 & -2 & -2 & 1 & -2 & 0 \end{bmatrix}$$

The game is symmetric, and is a zero sum game, so the game value is zero and the players have identical optimal strategies. We see that row 2 dominates row 5 and row 1 dominates rows 4, 6, 7,8,9,10, and so all dominated rows and columns cannot be part of the solution and so can be ignored. So we may restrict attention to the remaining 3 by 3 submatrix:

$$\begin{matrix} 0 & 2 & -1 \\ -2 & 0 & 2 \\ 1 & -2 & 0 \end{matrix}$$

Assuming player A chooses strategies $(A1 = 1, A2 = 2, A3 = 3)$ with probabilities (p_1, p_2, p_3), then since it is a symmetric game, player B will have the same values, ie $(q_1, q_2, q_3) = (p_1, p_2, p_3)$. From the principle of indifference (ie equalising strategies) we therefore have the following equations to solve:

$$\begin{aligned} -2p_2 + p_3 &= 0 \\ 2p_1 - 2p_3 &= 0 \\ -p_1 + 2p_2 &= 0 \end{aligned}$$

and we also know that

$$p_1 + p_2 + p_3 = 1$$

So we can solve the first equation to get $p_3 = 2p_2$ and $p_1 = p_3$, and substituting into the

last equation we find

$$2p_2 + p_2 + 2p_2 = 1$$

which gives $p_2 = 1/5$, and so $p_1 = 2/5$ and $p_3 = 2/5$.
This gives the solution

$$\boldsymbol{p} = \boldsymbol{q} = (2/5, 1/5, 2/5, 0, 0, 0, 0, 0, 0, 0)$$

which is optimal for both players, and the value of the game is zero.

Chapter 6

Extensive Form Games

Previously we assumed that a game is represented in normal form: effectively, as a big game table. In some sense, this is reasonable. The normal form is conceptually straightforward. While many other representations exist to describe finite games, we will see in this chapter that each of them has an "induced normal form": a corresponding normal form representation that preserves game-theoretic properties such as Nash equilibria. Thus the results given earlier hold for all finite games, no matter how they are represented; in that sense the normal form representation is universal. In this chapter we will look at extensive form games, a finite representation that does not always assume that players act simultaneously. This representation is in general exponentially smaller than its induced normal form, and furthermore can be much more natural to reason about. While the Nash equilibria of an extensive form game can be found through its induced normal form, computational benefit can be had by working with the extensive form directly. Furthermore, there are other solution concepts, such as subgame-perfect equilibrium which explicitly refer to the sequence in which players act and which are therefore not meaningful when applied to normal form games.

The normal form game representation does not incorporate any notion of sequence, or time, of the actions of the players. The extensive (or tree) form is an alternative representation that makes the temporal structure explicit. We start by discussing the special case of perfect information extensive form games, and then move on to discuss the more general class of imperfect information extensive form games. In both cases we will restrict the discussion to finite games, that is, to games represented as finite trees.

We will proceed with our study of games under the following assumptions:

1. There are a finite set of Players: $P = \{P_1, \ldots, P_N\}$.
2. Each player has a knowledge of the rules of the game (the rules under which the game state evolves) and the rules are fixed.
3. At any time $t \in \Re^+$ during game play, the player has a finite set of moves or choices to make. These choices will affect the evolution of the game. The set of all available moves will be denoted S.
4. The game ends after some finite number of moves.
5. At the end of the game, each player receives a prize. (We assume that these prizes can be ordered according to preference and that a utility function exists to assign numerical values to these prizes.)

In addition to these assumptions, some games may incorporate two other components:

1. At certain points, there may be chance moves which advance the game in a non- deterministic way. This only occurs in games of chance. (This occurs, e.g., in poker when the cards are dealt.)

2. In some games the players will know the entire history of moves that have been made at all times. (This occurs, e.g., in Chess, but not e.g., in Poker.)

6.1 Perfect Information Extensive Form Games

Informally speaking, a perfect information game in extensive form (or, more simply, a perfect information game) is a tree in the sense of graph theory, in which each node represents the choice of one of the players, each edge represents a possible action, or move, and the leaves represent final outcomes over which each player has a utility function. These are known simply as game trees.

Essentially, we will begin with some directed tree T. Each non terminal vertex of T will be controlled by a player who will make a move at the vertices she owns. If v is a vertex controlled by Player P , then out–edges from v will correspond to the possible moves Player P can take. The terminal vertices will represent endgame conditions (e.g., check–mate in chess). Each terminal vertex will be assigned a payoff (score or prize) amount for each player of the game. In this case, there will be no chance moves (all moves will be deliberately made by players) and all players will know precisely who is moving and what their move is.

(Player Vertex Assignment).
If $T = (V, E)$ is a directed tree, let $F \subseteq V$ be the terminal vertices and let $D = V \backslash F$ be the intermediate (or decision) vertices. An assignment of players to vertices is an onto function $\nu : D = V \backslash F \to P$ that assigns to each non-terminal vertex $v \in V \backslash F$ a player $\nu(v) \in P$. Then Player $\nu(v)$ is said to own or control vertex v.

(Move Assignment).
If $T = (V, E)$ is a directed tree, then a move assignment function is a mapping $\mu : E \to S$ where S is a finite set of player moves. So that if $v, u_1, u_2 \in V$ and $(v, u_1) \in E$ and $(v, u_2) \in E$, then $\mu(v, u_1) = \mu(v, u_2)$ if and only if $u_1 = u_2$.

(Payoff Function).
If $T = (V, E)$ is a directed tree, let $F \subseteq V$ be the terminal vertices. A payoff function is an onto mapping $\pi : F \to \Re^N$ that assigns to each terminal vertex of T a numerical payoff for each player in P.

(Game Tree with Complete Information and No Chance Moves).
A game tree with complete information and no chance moves is defined as the tuple $G = (T, P, S, \nu, \mu, \pi)$ such that T is a directed tree, ν is a player vertex assignment on intermediate vertices of T , μ is a move assignment on the edges of T and π is a payoff function on T.

(Strategy–Perfect Information).
Let $G = (T, P, S, \nu, \mu, \pi)$ be a game tree with complete information and no chance moves, with $T = (V, E)$. A pure strategy for Player P_i (in a perfect information game) is a mapping $\sigma_i : V_i \to S$ with the property that if $v \in V_i$ and $\sigma_i(v) = s$, then there is some $y \in V$ so that $(x, y) \in E$ and $\mu(x, y) = s$. (Thus σ_i will only choose a move that labels an edge leaving v.)

Remark 13. *(Rationality).*
A strategy tells a player how to play in a specific game at any moment in time. We assume that players are

rational and that at any time they know the entire game tree and that Player i will attempt to maximize her payoff at the end of the game by choosing a strategy function σ_i appropriately.

Remark 14. ***(Perfect Information:)***
All players know the game structure (including the payoff functions at every outcome).
Each player, when making any decision, is perfectly informed of all the events that have previously occurred.

Example: (Rock-Paper-Scissors).

Consider an odd version of rock-paper-scissors played between two people in which the first player plays first and then the second player plays. If we assume that the winner receives $+1$ points and the loser receives -1 points (and in ties both players win 0 points), then the game tree for this scenario is visualized in Figure 6.1. You may think this game is not entirely fair, which is not mathematically defined, because it looks like Player 2 has an advantage in knowing Player 1's move before making his own move. Irrespective of feelings, this is a valid game tree.

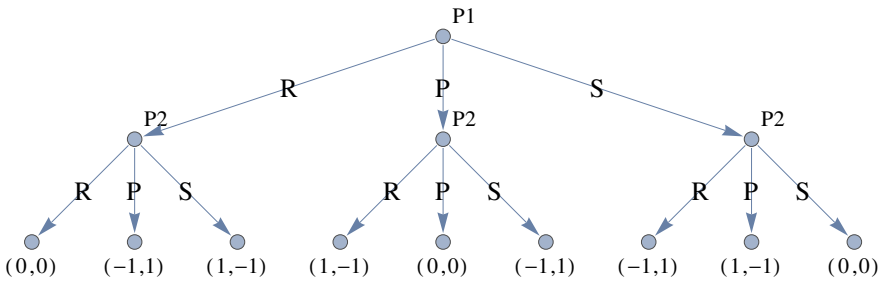

Figure 6.1: Example Game Tree for one-sided Rock-Paper-Scissors game.

Example: The Sharing Game.

Alice and Bob try to split two indivisible and identical gifts. First, Alice suggests a split: which can be "Alice keeps both", "they each keep one", and "Bob keeps both". Then, Bob chooses whether to Accept or Reject the split. If Bob accepts the split, they each get what the split specifies. If Bob rejects, they each get nothing.

- **Set of Players:**
  $P = \{Alice, Bob\}$.
- **Player Vertex Function,**
  $\nu(1) = Alice$,
  $\nu(2) = \nu(3) = \nu(4) = Bob$.
  (ie a vertex is 'owned' by a player).
- **Move Assignment:**
  $\mu(1,2) = 2 - 0$, $\mu(1,3) = 1 - 1$, $\mu(1,4) = 0 - 2$.
  (ie, a link corresponds to a chosen move.)

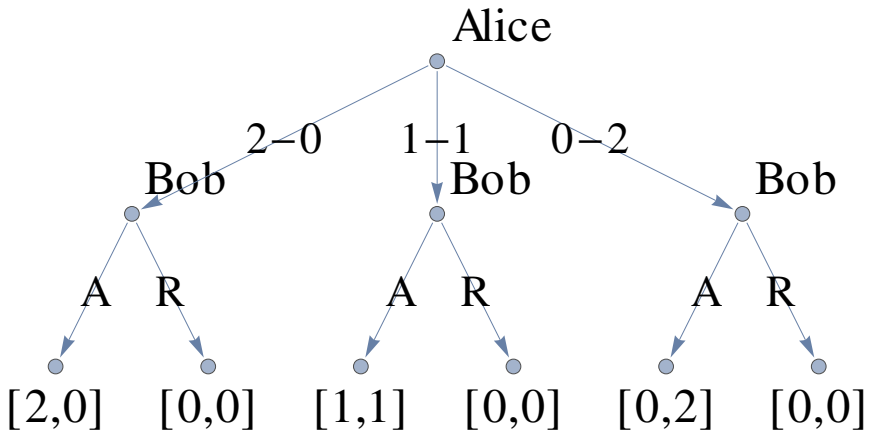

Figure 6.2: Extensive form of the Sharing Game

- **Payoff Function:**
  $\pi(5, Alice) = 2$, $\pi(5, Bob) = 0$, $\pi(6, Alice) = 0$, $\pi(6, Bob) = 0$, $\pi(7, Alice) = 1$, $\pi(7, Bob) = 1$, $\pi(8, Alice) = 0$, $\pi(8, Bob) = 0$, $\pi(9, Alice) = 0$, $\pi(9, Bob) = 2$,$\pi(10, Alice) = 0$, $\pi(10, Bob) = 0$.
  (ie terminal nodes have associated prizes).

6.1.1 (Pure) Strategies

To illustrate the notion of a strategy consider the game in Figure 6.3. Player I takes an action only after the initial history, so that we can identify each of her strategies with one of the two possible actions that she can take after this history: r or l. Player II takes an action after each of the two histories r,l, and depending on the case he has either two or three possible actions, either a choice from (R, L), or a choice from (R, M, L). For one of these choices, note that player II gets to play again, with two choices R or L. Then player I possibly has another go, depending on Player II's choices. So the pure strategies for each player is the set of all possible choices at each node that the player 'owns'. So if player i owns 4 nodes, and each node has 3 choices, the set of pure strategies will consist of a vector of values where each value can take on one of the 3 choices. The following example may make this clear.

Example: (Asymmetric)

Note: a tree graph does not have any loops or cycles.

Recall that a pure strategy for Player P_i (in a perfect information game) is a mapping $\sigma_i : V_i \to S$ with the property that if $v \in V_i$ and $\sigma_i(v) = s$, then there is some $y \in V$ so that $(x, y) \in E$ and $\mu(x, y) = s$. (Thus σ_i will only choose a move that labels an edge leaving v.)

In this case, Player I's pure strategies are:

`ll, lr, rl, rr`

(since Player I 'owns' two vertices, and there are two choices on each vertex), and Player II's pure strategies are:

`LLL, LLR, LML, LMR, LRL, LRR, RLL, RLR, RML. RMR, RRL, RRR`

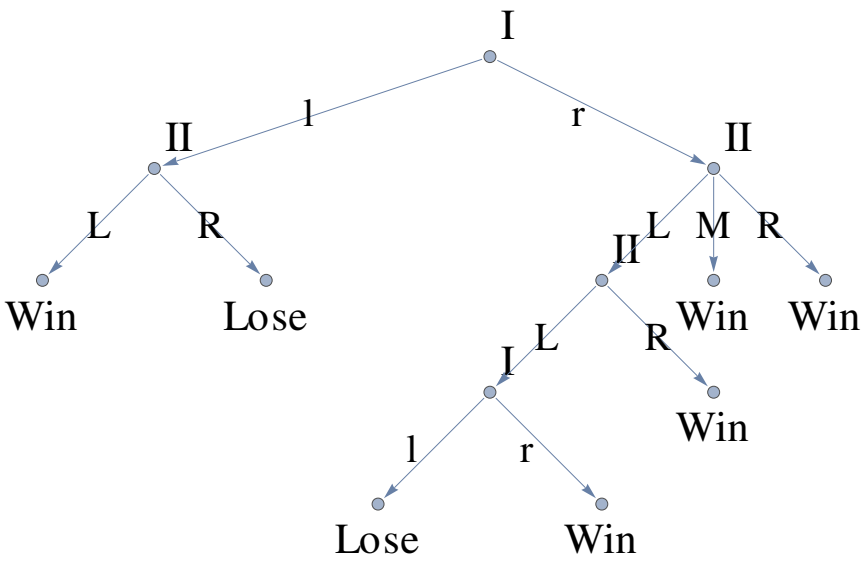

Figure 6.3: Asymmetric Game. A Game where player I moves both before and after player II.

since Player II owns three vertices and one of the vertices has three choices, the other two have two choices each.

The strategies in an extensive form game are regarded as *contingent* choices (ie a choice which depends on the previous move of the opponent).

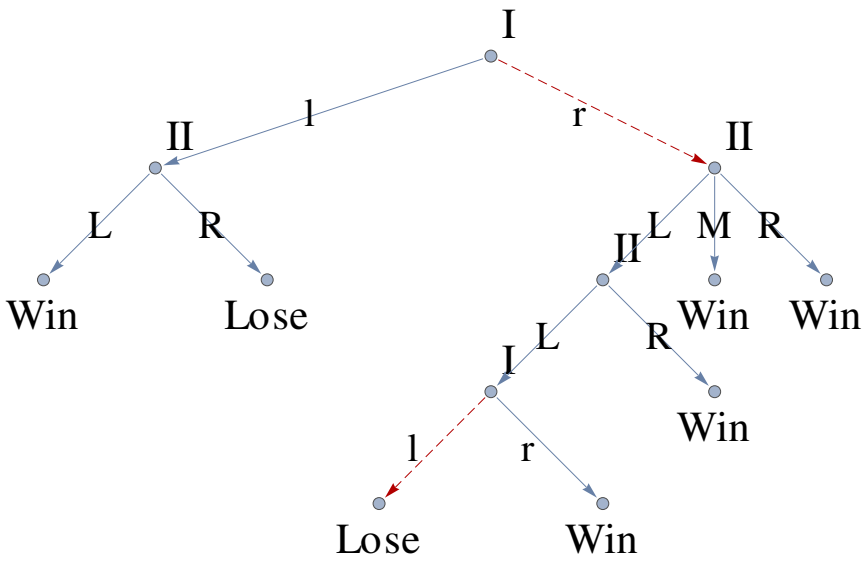

Figure 6.4: Asymmetric Game. Showing Player I's pure strategy `rl` shown as the dashed red line

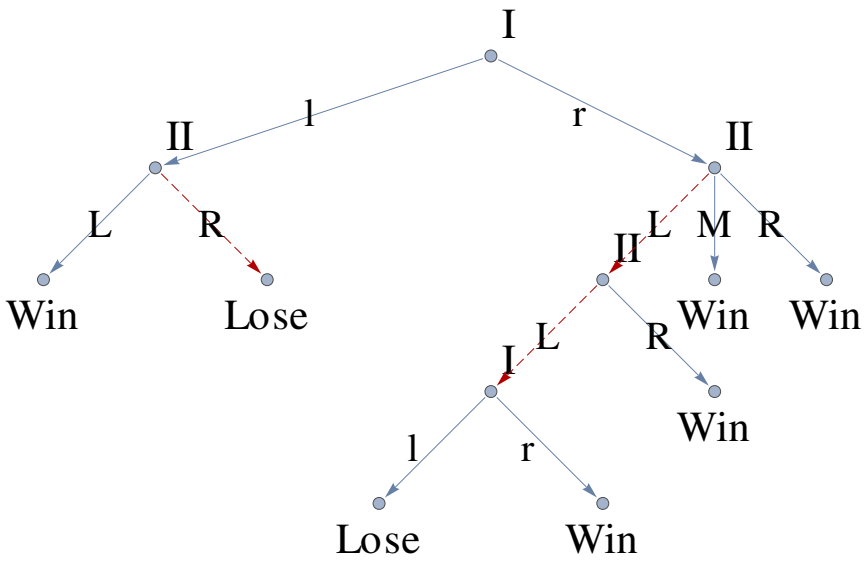

Figure 6.5: Asymmetric Game. Showing Player II's pure strategy `RLL` as the dashed red line.

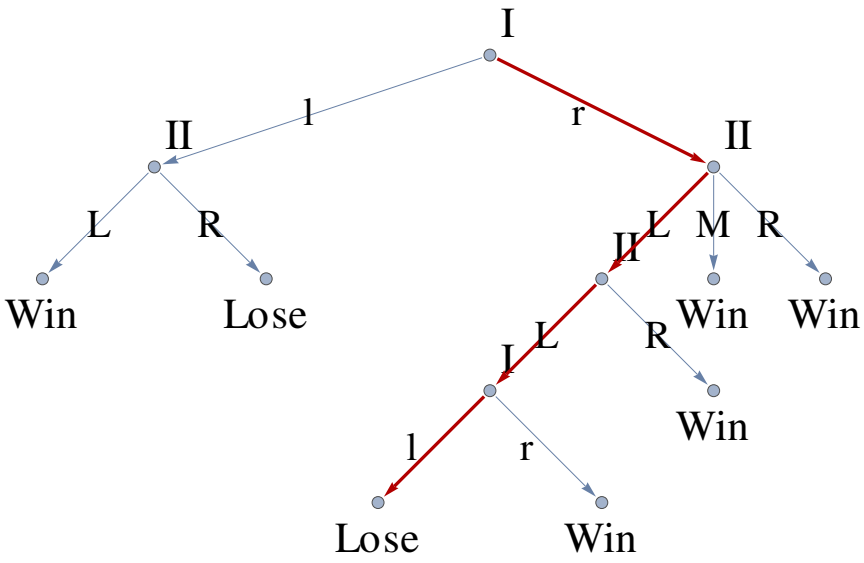

Figure 6.6: Asymmetric Game. Showing the resulting Game play of (`rLLl`)from the combined pure strategies as a solid thick red line.

6.1.2 Extensive Form ⇒ Strategic Form

An extensive form game can be placed into a strategic form, by finding all the pure strategies and assigning the payoffs to the resulting combinations.

So in the previous example we have already seen the list of pure strategies of Player I {`ll, lr, rl, rr`} and Player II {`LLL, LLR, LML, LMR, LRL, LRR, RLL, RLR, RML. RMR, RRL, RRR`}. So we can construct a strategic form matrix for the game as follows:

		Player II											
		LLL	LLR	LML	LMR	LRL	LRR	RLL	RLR	RML	RMR	RRL	RRR
Player I	ll	?	?	?	?	?	?	?	?	?	?	?	?
	lr	?	?	?	?	?	?	?	?	?	?	?	?
	rl	?	?	?	?	?	?	?	?	?	?	?	?
	rr	?	?	?	?	?	?	?	?	?	?	?	?

Table 6.1: Strategic form matrix game for the Asymmetric Game Example

By following the paths and pure strategies in the game tree we can find the payoffs for each combination to appear in this payoff table.

		Player II											
		LLL	LLR	LML	LMR	LRL	LRR	RLL	RLR	RML	RMR	RRL	RRR
Player I	ll	W	W	W	W	W	W	L	L	L	L	L	L
	lr	W	W	W	W	W	W	L	L	L	L	L	L
	rl	L	W	W	W	W	W	L	W	W	W	W	W
	rr	W	W	W	W	W	W	W	W	W	W	W	W

Table 6.2: Strategic form matrix game for the Asymmetric Game Example with payoffs

We can see from this payoff matrix that Player I is sure to win if she plays `(r,r)`.

6.1.3 Backward Induction

Another way to determine who will win in an extensive form game, is a method known as Backward Induction. It links to a significant theorem on extensive form non-random win/lose games:

Theorem 8. *In every finite win-or-lose game (without chance moves), one player can force a win against any defence by the other player.*

We can illustrate backward induction explicitly by following the previous asymmetric game example. Backward induction works by considering sub-trees, starting from the terminal vertices and working back to the root vertex. Each sub-tree is chosen such that the optimal strategy is followed for whichever player owns the local 'root' node of the local sub-tree.

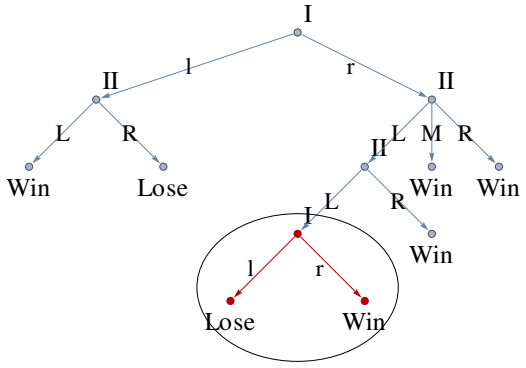

Figure 6.7: Asymmetric Game. The first sub-tree is shown in red and ringed with a circle.

Consider Figure 6.7 which shows the full game with a sub-tree highlighted in red and ringed with a circle. This subtree can be analysed independently. The root node is owned by Player I so they would choose option `r` to force a `Win`. So we replace the entire sub-tree by a terminal vertex whose payoff is `Win`.

This gives Figure 6.8. This in turn may be analysed by selecting a sub-tree (again highlighted in red and ringed with a circle) as shown in figure 6.9. The subtree in Figure 6.9 is owned by Player II, but unfortunately both options for Player II lead to a `Win` for Player I. Therefore we can replace this subtree with a new terminal vertex with a `Win` payoff.

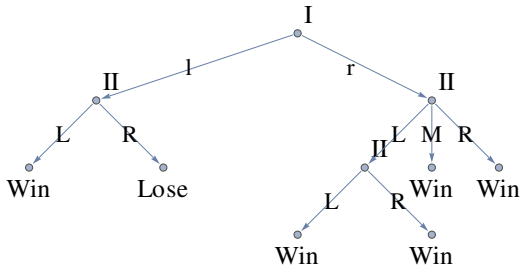

Figure 6.8: Asymmetric Game. The reduced extensive form game

This leads on to the revised tree graph in Figure 6.10, and its corresponding subtree in figure 6.11.

In this case, the subtree is owned by Player II and Player II has a winning strategy here by selecting option `R` (which gives a `Lose` for Player I). Therefore we replace this subtree with a terminal vertex with payoff `Lose` for Player I.

This leads on to Figure 6.12 as a reduced game tree and its final subtree shown in Figure 6.13.

This last subtree is owned by Player II, but none of his options lead to a winning payoff for him. So we can replace this final subtree with a terminal vertex of payoff `Win`. The result of all this subtree reduction on the extensive form game leaves us with the equivalent game shown in Figure 6.14. This is a subtree owned by Player I. As can be seen in this figure, Player I has a winning strategy by selecting option `r`, which agrees with the comment at the end of the discussion on the equivalent strategic form game matrix in the previous subsection.

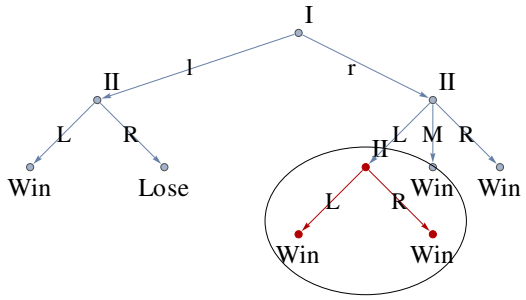

Figure 6.9: Asymmetric Game. New subtree highlighted.

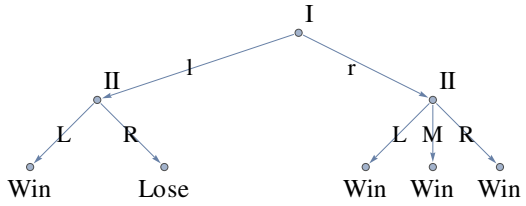

Figure 6.10: Asymmetric Game. The reduced extensive form game

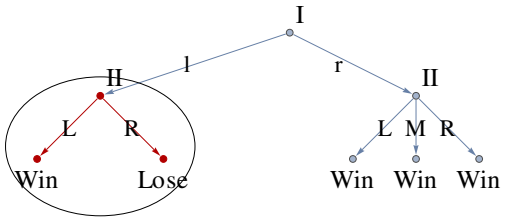

Figure 6.11: Asymmetric Game. New subtree highlighted.

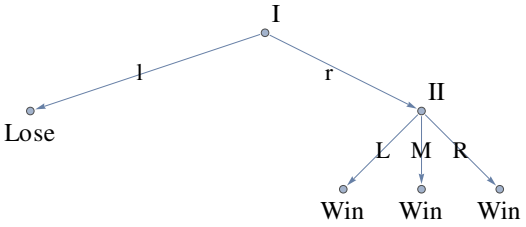

Figure 6.12: Asymmetric Game. The reduced extensive form game

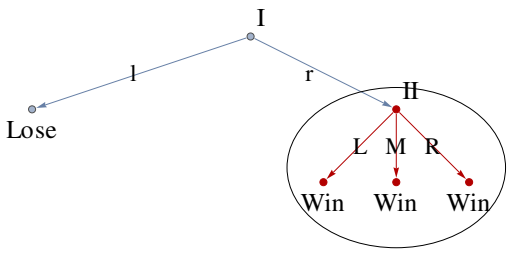

Figure 6.13: Asymmetric Game. New subtree highlighted.

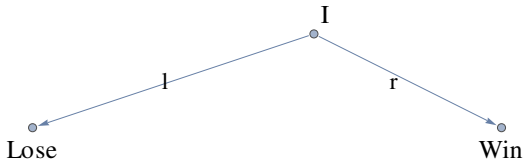

Figure 6.14: Asymmetric Game. Final reduced extensive form game after the process of backward induction leading to a winning strategy for Player I

6.2 Exercises: Perfect Information Extensive Form Games

6.2.1 Exercise (Entry Game)

In the normal way of describing this game, an "incumbant" (ie someone who is dominating a market) faces the possibility of a challenger entering her market. The challenger may be a company for example considering entering a new market occupied by a monopolist, an animal competing for the right to mate, or a politician seeking a leadership battle. So, the challenger may enter or not, and the incumbent may fight, or acquiesce if the challenger chooses to enter the market. If the challenger does not enter the market, the incumbant does nothing.

This may be modelled as an extensive game with perfect information in which the game play possible sequences of actions are `(In, Acquiesce)` , `(In, Fight)`, and `Out`. ie the move assignment function μ associates the links of the tree with the strategies `In, Out, Acquiesce, Fight`.

The player function assigns the challenger to the initial root node and the incumbent to the second play.

Assume the payoff function π for the challenger is such that it assigns values of
π(In, Acquiesce) = 2, π(Out) = 1, π(In,Fight)=0,
and the for Incumbent, the payoff function assigns the values:
π(Out) = 2, π(In, Acquiesce) = 1, and π(In, Fight) = 0.

(i) Construct the extensive game tree for this problem.

(ii) Specify each player's pure strategies and write down the strategic form game matrix.

(iii) Use backward induction to show that the Challenger will enter and the Incumbent will Acquiesce.

Answer:

(i)

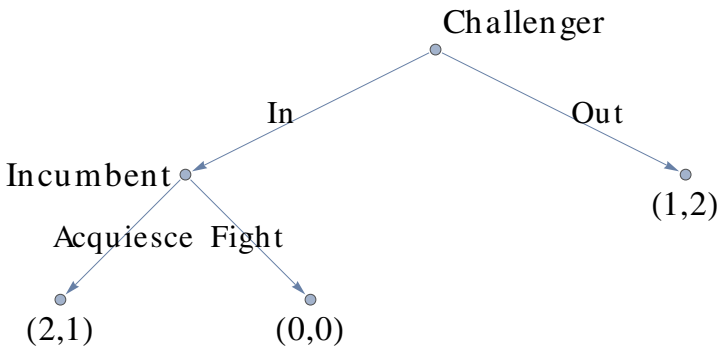

Figure 6.15: Extensive form game tree for the Entry Game.

(ii) The pure strategies for the Challenger are `(In, Out)`. The pure strategies for the Incumbent are `(Acquiesce, Fight)`. So the payoff matrix will be a 2×2 matrix with values as shown in Table 6.3.

		Incumbent	
		Acquiesce	Fight
Challenger	In	$(2, 1)$	$(0, 0)$
	Out	$(1, 2)$	$(1, 2)$

Table 6.3: Strategic form payoff matrix for the Entry Game.

(iii) There is one obvious subtree in the full treegraph. This is shown in Figure 6.16. This subtree is owned by the Incumbent and it is clear that the Incumbent's obvious choice is to select `Acquiesce` since this will give highest payoff (ie 1 as opposed to 0). Therefore we may replace the subtree with a terminal vertex as shown in Figure 6.17. This treegraph is owned by the Challenger and again, it is clear that the Challenger will choose to enter the market (ie chooses `In`) since this gives highest payoff (2 as opposed to 1).

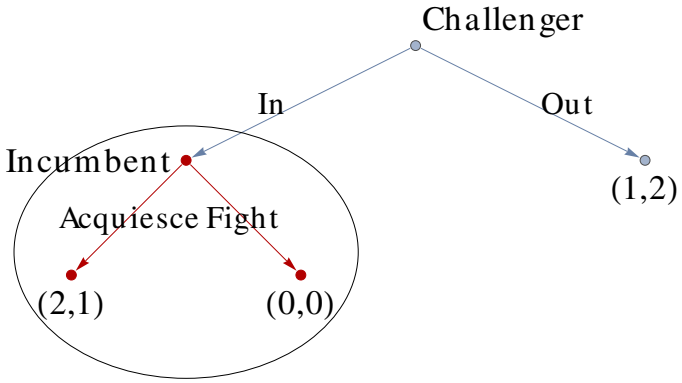

Figure 6.16: Subtree analysis for the Entry Game.

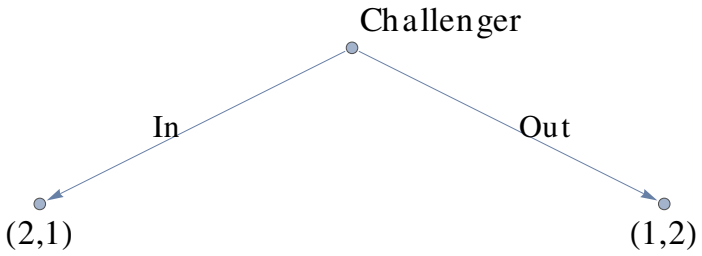

Figure 6.17: Reduced treegraph for the Entry Game by collapsing the subtree.

6.2.2 Exercise: (Service Provider)

Suppose player I is an internet service provider and player II a potential customer. They consider entering into a contract of service provision for a period of time. The provider can, for herself, decide between two levels of quality of service, `High` or `Low`. High-quality service is more costly to provide, and some of the cost is independent of whether the contract is signed or not. The level of service cannot be put verifiably into the contract. High-quality service is more valuable than low-quality service to the customer, in fact so much so that the customer would prefer not to buy the service if she knew that the quality was low. Her choices are to buy (`Y`) or not to buy (`N`) the service.

The customer prefers to buy a high quality service (assume a payoff of 2 for this choice). But buying a low quality service is the worst of all options (payoff of 0). All other options has a payoff of 1 to the customer.

The service provider's worst option is if they go to the expense of producing a high quality service that the customer does not buy (so has payoff of 0 to the provider). The next worst option to the provider is if they set up a low quality service that the customer does not buy (payoff of 1). The next preference to the provider is if they provide a high quality service that the customer does buy (payoff of 2) and the best option is if the provider offers a low quality cheap service that the customer buys (payoff of 3).

The provider goes first, and offers a service and the customer decides if he wants it.

(i) Construct an extensive form game tree for this game, and list the two players pure strategies.

(ii) Use backward induction to determine the game play assuming both players act rationally to maximise their individual payoffs.

(iii) Construct the equivalent strategic form game.

Answer:

(i)

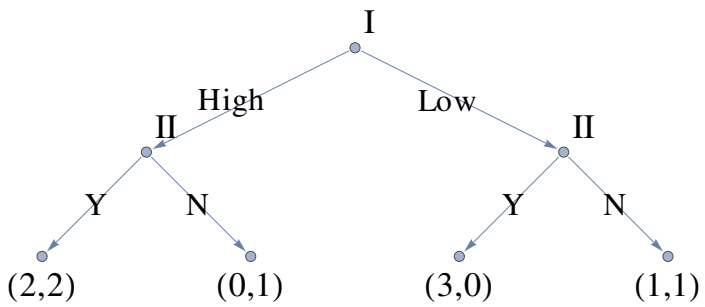

Figure 6.18: extensive form game tree for the service provider game.

The pure strategies are:

Provider: `(High, Low)`

Customer: `(Y,Y), (Y,N), (N,Y), (N,N)`

(ii) The game plan *solution* of the extensive form game is {`High, Y`} which we can see by using backward induction on the two obvious subtrees and analysing the payoffs.
Note that the game solution is how the game is played, and is *not* how we specify the result as an equilibrium in pre strategies. As an equilibrium in pure strategies, this is more correctly written as `SPE = (H,(Y,N))`.
Extensive games with perfect information can be analyzed by backward induction. This technique solves the game by first considering the last possible choices in the game. Here, player II (Customer) moves last. Since he knows the play will end after his move, he can safely select the action which is best for him. If player I (Service Provider) has chosen to provide high-quality service, then the customer prefers to buy, since his resulting payoff of 2 is larger than 1 when not buying. If the provider has chosen `Low`, then the customer prefers not to purchase. Once the last moves have been decided, backward induction proceeds to the players making the next-to-last moves (and then continues in this manner). In the following figures, player I (Provider) makes the next-to-last move, which in this case is the first move in the game. Being rational, she anticipates the subsequent choices by the customer. She therefore realises that her decision between `High` and `Low` is effectively between the outcomes with payoffs $(2, 2)$ or $(1, 1)$ for the two players, respectively. Clearly, she prefers `High`, which results in a payoff of 2 for her, to `Low`, which leads to an outcome with payoff 1. So the unique solution to the game, as determined by backward induction, is that player I offers high-quality service, and player II responds by buying the service. This sequence of decisions is revealed in the following figures of subtrees.

It can be seen in this last example in Figure 6.20, that Player I (the Provider) will choose `High`. and hence, the customer, Player II, will choose `Y` from figure 6.19.

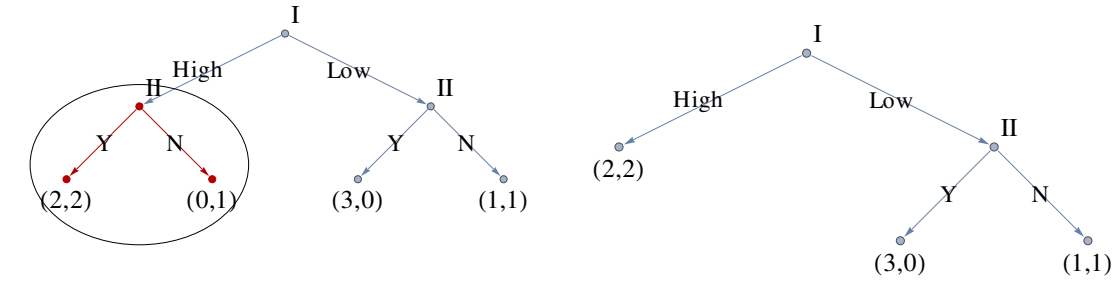

Figure 6.19: Service Provider Example: first level of decomposition

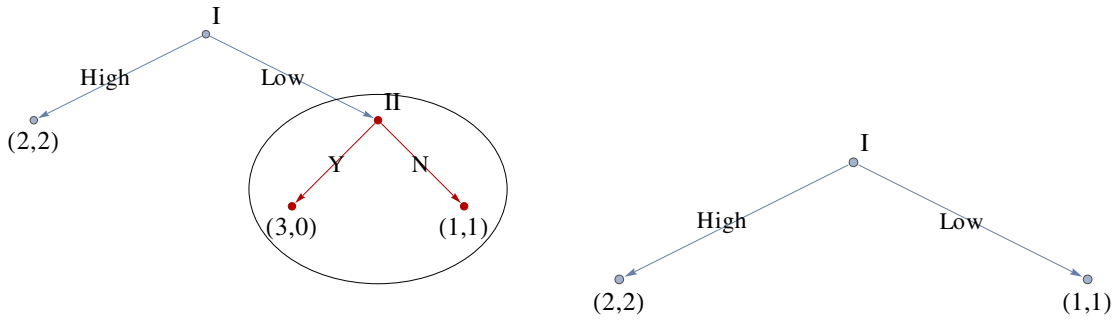

Figure 6.20: Service Provider Example: Final level of decomposition

(iii) The equivalent strategic form game is given by the payoff table 6.4:

		Customer			
		YY	YN	NY	NN
Provider	High	2, 2	2, 2	0, 1	0, 1
	Low	3, 0	1, 1	3, 0	1, 1

Table 6.4: Equivalent Strategic Form Payoff Table for the Service Provider Game

6.2.3 Exercise: A Stackelberg Duopoly Example

Recall the Cournot duopoly model in which two companies simultaneously choose a quantity to produce. The Stackelberg duopoly model considers the outcome when the quantity choices are made sequentially. Consider the following extensive form game:
Alice can choose not to enter a market, or, otherwise she chooses to manufacture q_A units for the market. Bob, having observed what Alice has done, can also choose to not enter the same market, or otherwise he chooses to manufacture q_B units. Each unit sells for $£q$. However the market demand for the products is such that the price is given by

$$q = q_0 - Q,$$

where q_0 is the maximum number of units the market will stand, and $Q = q_A + q_B$ is the total number of units produced by both companies. Assume the cost of producing one unit is c $(0 \leq c \leq q_0)$ and is the same for both Alice and Bob.

(a) Assuming both Alice and Bob wish to maximise their individual profits, determine how many units each will produce and what is the price that each unit will sell for. Remember that Alice realises that Bob will produce his units in response to what Alice chooses to do.

b) Determine who makes the most profit? (ie in this game work out if is there a "first player advantage"?

Answer:

(a) The profit made by Bob if he makes q_B units (and also assuming Alice produces q_A units !) is volume times (price-cost) ie

$$\pi_B[q_A, q_B] = q_B \times [q_0 - (q_A + q_B) - c].$$

To maximise profit, we take the derivative and equate to zero and solve, and we find the optimum choice for how many units Bob makes depends on q_A and is

$$q_B^* = \frac{q_0 - q_A - c}{2}.$$

Alice knows this. So to maximise her profit

$$\begin{aligned} \pi_A[q_A, q_B] &= q_A \times [q_0 - (q_A + q_B) - c] \\ &= q_A \times [q_0 - (q_A + \frac{q_0 - q_A - c}{2}) - c] \end{aligned}$$

Differentiating wrt q_A, setting to zero, and solving, gives the optimal choice of q_A for Alice as:

$$q_A^* = \frac{q_0 - c}{2}.$$

Substituting back into the expression for q_B gives the optimum choice for Bob to make as

$$q_B^* = \frac{q_0 - c}{4}.$$

Hence the total number of units made is $3(q_0 - c)/4$ and the price per unit is $(q_0 + 3c)/4$.
Since these quantities are positive, it is always advantageous for Alice and Bob to both enter the market and make q_A, q_B units respectively.

b)

$$\begin{aligned}\pi_A[q_A, q_B] &= q_A \times [q_0 - (q_A + q_B) - c] \\ &= \frac{q_0 - c}{2} \times [q_0 - \frac{q_0 - c}{2} - \frac{q_0 - c}{4} - c \\ &= \frac{(q_0 - c)^2}{8}\end{aligned}$$

and

$$\begin{aligned}\pi_B[q_A, q_B] &= q_B \times [q_0 - (q_A + q_B) - c] \\ &= \frac{q_0 - c}{4} \times [q_0 - \frac{q_0 - c}{2} - \frac{q_0 - c}{4} - c \\ &= \frac{(q_0 - c)^2}{16}\end{aligned}$$

so there is a first player advantage by a factor of 2.

6.2.4 Exercise: Contract Negotiation: Team vs Footballer

It is contract negotiation time between the Football Team and their star Player. The negotiation takes the form of up to 3 rounds.

Round 1 In round 1 both parties simultaneously choose whether the Player receives a Big or Small contract offer. If both agree the contract is signed and the game ends. Otherwise proceed to ound 2. We can summarise the payoffs to the Team and Player in Round 1 as follows:

		Player	
		Big	Small
Team	Big	$(20, 100)$	Round 2
	Small	Round 2	$(100, 20)$

Figure 6.21: Round 1 payoffs

Round 2 In this round of negotiation, the two parties simultaneously decide whether to walk away from negotiations or continue. If at least one of them walks away then the negotiations end with payoffs given in Table 6.22. If both continue with negotiations they move to Round 3. Payoffs for Round 2 can be summarised in Table 6.22.

Round 3 The Team now makes a High or Low offer. Seeing this offer the footballer can either accept or reject the revised offer. If he accepts the High offer, the contract is signed with a payoff of 15 utils

		Player	
		Walk	Bargain
Team	Walk	$(25, 25)$	$(10, 30)$
	Bargain	$(30, 10)$	Round 3

Figure 6.22: Round 2 payoffs

to the Team and 80 utils to the Player. If he accepts the Low offer, the contract is signed with a payoff of 80 utils to the Team and 15 utils to the Player. If he rejects, both parties have zero payoff.

Find the subgame perfect equilibrium of this negotiation game and determine the equilibrium payoffs to each party.

Answer:

Solve by Backward Induction. If we represent round 3 as an extensive form game then it can be represented as shown in Figure 6.23.

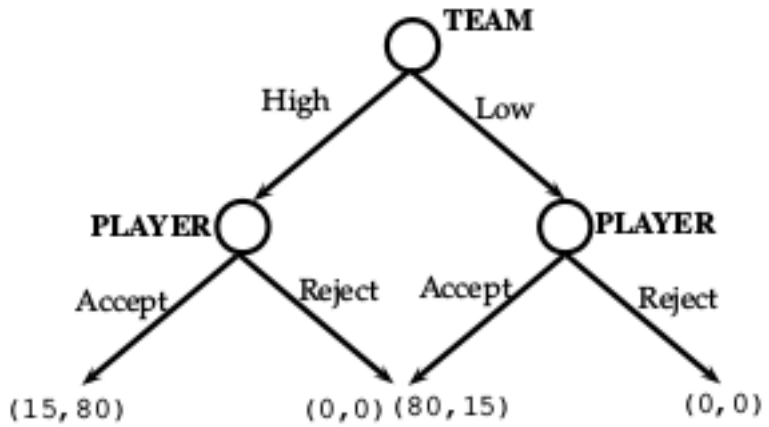

Figure 6.23: Round 3 extensive form game and payoffs

By exploring the two subgames controlled by the Player it is clear Payer will choose (Accept,Accept) and so in the subgame controlled by Team, Team will choose (Low) giving a payoff of (80,15) for the Round 3 game.

So we can substitute this back into the payoff matrix in Round 2 to get Figure 6.24:

		Player	
		Walk	Bargain
Team	Walk	$(25, 25)$	$(10, 30)$
	Bargain	$(30, 10)$	$(80, 15)$

Figure 6.24: Round 2 payoffs with the Round 3 Payoffs now inserted.

This payoff matrix has a Nash Equilibrium of (Bargain,Bargain) with a payoff of (80,15).
So substituting this back into the game matrix for Round 1 we have Figure 6.25:

		Player	
		Big	Small
Team	Big	$(20, 100)$	$(80, 15)$
	Small	$(80, 15)$	$(100, 20)$

Figure 6.25: Round 1 payoffs

Again, solving this equation to find the Nash Equilibria, we find a single equilibrium at (Small,Small) with a payoff of (100,20). Therefore the equilibrium solution strategies for the Team and the Player are:
Team: `(Small, Bargain,Low)`
Player: `(Small, Bargain,(Accept,Accept))`.

6.2.5 Exercise: Redistribution of Marks

A mark of 100% is initially distributed to the 4 members (named A, B, C, D) of a game theory group of students for success in their team-based assessment with the following proportions:

$$[A = 6, B = 46, C = 24, D = 24].$$

The lecturer allows the team to reallocate the proportions under the following rules:

(i) Any proposal is voted on by all team members and is passed only if at least 3 members agree to it.

(ii) If a team member makes a proposal that fails, that person gets a mark of zero and their marks are equally split amongst the other team members.

The team members vote one after the other in the order $[A, B, C, D]$ and each team member can see the previous team member's decisions before they make their own. (ie it is a perfect information game). Team member A makes the following proposal:
The proportions of the marks are to be redistributed according to the new proportions $[A = 40, B = 20, C = 20, D = 20]$*, and if the proposal passes, anyone who voted against it gets zero and their allocated marks equally distributed to the other team members.* Provide an argument based on backward induction of subgame perfect equilibria to decide whether or not this proposal is accepted, and who gets what mark.

Answer:

This is an extensive form game and a large game tree can be produced with various payoffs. Or we can use the following logical argument. :
We know that A will always vote "Yes" for their own proposal. We can now solve the game by backward induction. The last vote is cast by D. D can observe (a) three "Yes" votes for the proposal, (b) two "Yes" votes for the proposal, and (c) one "Yes" vote for the proposal. For (a), with three "Yes" votes, the proposal has already passed, so D prefers to get 20% rather than 0%, so will vote "Yes". For (b), with two "Yes" votes, D is the decisive voter.

By voting "Yes", then the proposal passes and D receives 20% plus 1/3 of the share of the dissenting member (which is also 20%). This amount is higher than what D receives by voting "No" (keeping their original share of 24% plus 1/3 of the original share of A). For (c), with one "Yes" vote, the proposal fails, so D's vote does not matter.
The second to last vote is cast by C. C can observe (a) two "Yes" votes for the proposal, or (b) one "Yes" vote for the proposal. Under (a), with two "Yes" votes, then C knows that D will always follow with a "Yes" vote. Thus, the proposal will pass, so it is optimal for C to vote "Yes" as she prefers to get 20% rather than 0%. Under (b), with one "Yes" vote, C is the decisive voter. By voting "Yes", then the proposal passes and C receives 20% plus 1/3 of the share of the dissenting voter (that share being the 20% share of B). This amount is higher than what C receives by voting "No" (keeping her original share of 24% plus 1/3 of the original share of A). The third to last vote is cast by B. B is in a tough spot. B observes the one "Yes" vote by A. No matter what B does, the proposal will pass (C and D will both vote "Yes"). Thus, it is optimal for B to vote "Yes" as he prefers to get 20% rather than 0%. In the end, the proposal passes with four "Yes" votes.

6.3 Perfect Information Extensive Form Games: Equilibria

As was the case for our earlier discussion on strategic games we are interested in notions of equilibrium that model the players' behaviour in steady state. We look for patterns of behaviour such that if every player knows every other players' behaviour she has no reason to change her own behaviour.

The first concept is a generalisation of the idea of a Nash equilibrium we met for strategic form games. This is a Nash equilibrium for extensive form games:

(Nash Equilibrium of extensive games with perfect information)
The strategy profile $\boldsymbol{s}^*$ in an extensive form game is a Nash Equilibrium if, for every player i and every strategy profile $\boldsymbol{r}_i$ of player i, the terminal payoff $\pi(\boldsymbol{s}^*)$ generated by $\boldsymbol{s}^*$ is at least as good according to player i's preferences as the terminal payoff $\pi(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\})$ generated by the strategy profile $\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}$ where player i chooses strategy $\boldsymbol{r}_i$ while every other player j chooses $\boldsymbol{s}^*_j$.

Equivalently, for each player i:

$$\pi_i(\boldsymbol{s}^*) \geq \pi_i(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}).$$

One way to find the Nash equilibria in extensive form games is to list each player's strategies, find the outcome of each strategy profile, and analyse this information as if it were a strategic form game.

From Definition 6.3 we see that:

...the set of Nash Equilibria of any extensive game with perfect information is the set of Nash Equilibria of the equivalent strategic form game.

Example: (Nash Equilibria of the Entry Game)
In the Entry Game we saw previously, its strategic game matrix was given in Table 6.3. We can analyse this and see that this payoff matrix has the Nash equilibria of `(In, Acquiesce)` and `(Out,Fight)`. The first equilibrium is the solution found by the method of backward induction.

In the second equilibrium the Challenger always chooses `Out`. This strategy is optimal given the Incumbents strategy to fight in the event of an Entry of the Challenger. If the Challenger chooses `Out` it makes no difference to the Incumbent whether they choose `Acquiesce` or `Fight` to the payoff, and so `Out, Fight` is an equilibrium since neither player can improve on their payoffs *given the other player's strategy.*

This raises a question of why the equilibrium `(Out, Fight)` does not occur in backward induction when analysing the extensive form game directly. This latter equilibrium point is not "stable" as we will see later.

6.4 Subgame Perfect Equilibrium

The notion of Nash equilibrium ignores the sequential nature of an extensive game; it treats strategies as choices made once and for all at the start of the game. We therefore introduce a notion of equilibrium that requires each player's strategies to be optimal not only at the start of the game, but at every possible history of the game.

For any nonterminal vertex, v the *subgame* $\Gamma(v)$ following v is the remaining game tree following the choice of v.

Because there is a subgame for every possible nonterminal vertex, there are as many subgames as there are nonterminal vertices.

Consider the extensive form game shown in figure 6.26. There are two non terminal vertices (apart from the root vertex itself) and so there are going to be two proper subgames (plus the original game stemming from the root vertex).

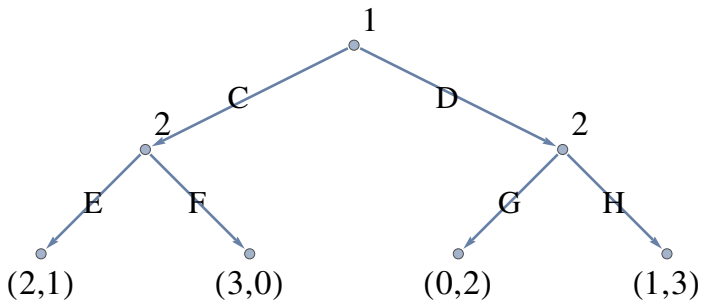

Figure 6.26: Simple Extensive Form Game.

The proper subgames are shown in Figure 6.27.

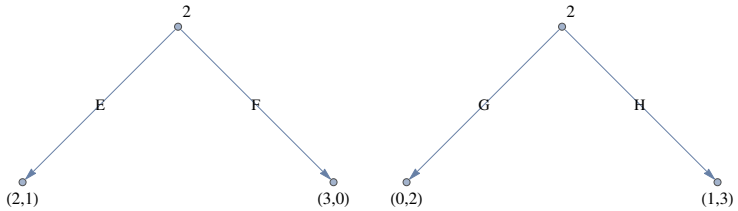

Figure 6.27: The two subgames of Figure 6.26.

The idea of subgame perfect equilibria can be informally embodied as follows:

A subgame perfect equilibrium is a strategy profile $\boldsymbol{s}^*$ with the property that in no subgame can any player i do better by choosing a strategy different from $\boldsymbol{s}^*$ given that every other player j adheres to $\boldsymbol{s}^*$.

So using this definition, we can now see that the equilibrium solution we found in the strategic form for the Entry Game, is *not* a subgame perfect equilibrium because in the subgame following `In`, the choice of `Fight` is not optimal for the incumbent, since the Incumbent is better off to choose `Acquiesce`. However the Nash equilibrium `(In, Acquiesce)` *is* a subgame perfect equilibrium, since each player's strategy is optimal, both in the whole game and in the subgame following the choice `In`.

We can be a little more precise in the following definition which clarifies the role of the subgame after each choice.

(Subgame perfect equilibrium of extensive game with perfect information)
Let π_i^v be the terminal payoffs for player i following the choice leading to vertex v. The strategy profile $\boldsymbol{s}^*$ of a perfect information extensive game is a subgame perfect equilibrium for every player i and every nonterminal vertex v after which it is player i's turn to move and every strategy $\boldsymbol{r}_i$ of player i, the terminal payoff generated by $\boldsymbol{s}^*$ after the vertex v is at least as good according to Player i's preferences as the terminal payoff $\pi^v(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\})$ in which player i chooses $\boldsymbol{r}_i$ and all other players choose $\boldsymbol{s}^*_j$.

Equivalently, for each player i and for every vertex v after which it is player i's turn to move,

$$\pi_i^v(\boldsymbol{s}^*) \geq \pi_i^v(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}) \quad \forall\ v. \tag{6.1}$$

The important point in this definition is that each player's strategy is required to be optimal for *every* nonterminal vertex after which it is that player's turn to move, and not only at the start of the game as was the case in the Nash equilibrium.

6.4.1 Subgame perfect equilibria and Nash equilibria

In a subgame perfect equilibrium, every player's strategy is optimal (according to their own individual payoffs). Thus:

Remark 15. ... every subgame perfect equilibrium is a Nash equilibrium.

In fact a subgame perfect equilibrium generates a Nash equilibrium in every subgame.

So, equivalently, we can define a subgame perfect equilibrium to be a strategy profile $\boldsymbol{s}$ in subgame Γ for which for any nonterminal vertex v the strategy profile $\boldsymbol{s}^*|v$ is a Nash equilibrium of the subgame $\Gamma(v)$. [We use the notation that $\boldsymbol{s}^*|v$ refers to the strategy profile restricted to the subgame that follows node v].

The notion of subgame perfect equilibrium eliminates Nash equilibria in which the players' threats are not credible.

For example, in the Entry game on which there are two Nash equilibria the fact that the Nash equilibrium `(Out, Fight)` is not a subgame perfect equilibrium follows from the definition 6.1 since:
for $\boldsymbol{s}^* =$ `Out, Fight`, then

$$\begin{aligned} i &= \mathit{Incumbent}, \\ r_i &= \texttt{Acquiesce}, \\ v &= \texttt{In}, \\ \pi_i^v(\boldsymbol{s}^*) &= \texttt{(In, Fight)} \equiv 0, \\ \pi_i^v(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}) &= \texttt{(In, Acquiesce)} \equiv 1. \end{aligned}$$

So for this Nash equilibrium the inequality in 6.1 is violated, and so it is not a subgame perfect equilibrium. The other Nash equilibrium `In, Acquiesce` *is* a subgame perfect equilibrium since according to the definition 6.1:

- At the start of the game the Challenger moves, so

$$\begin{aligned} i &= \mathit{Challenger}, \\ \boldsymbol{s}^* &= \texttt{In, Acquiesce}, \\ \pi_i(\boldsymbol{s}^*) &= 2. \end{aligned}$$

whereas for the other strategy of the Challenger we have

$$r_i = \texttt{Out},$$
$$\pi_i^v(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}) = \texttt{(Out, Acquiesce)} \equiv 1.$$

Therefore the inequality in Definition 6.1 is obeyed.

- The Incumbent moves after one strategy selection, namely `In`. Now we have for $i = Incumbent$,

$$\boldsymbol{s}^* = \texttt{In, Acquiesce} \equiv 1,$$
$$r_i = \texttt{Fight}, \text{ and so}$$
$$\pi_i^v(\{\boldsymbol{r}_i, \boldsymbol{s}^*_{-i}\}) = \texttt{(In, Fight)} \equiv 0.$$

So the inequality in Definition 6.1 is obeyed and so this Nash equilibrium is also a subgame perfect equliibrium.

We can therefore conclude these discussions with a formal proposition:

(Kuhn's Theorem: Existence of subgame perfect equilibria)
Every finite extensive form game with perfect information has a subgame perfect equilibrium.

Note that this does not claim there exists a single subgame perfect equilibrium, there may in fact be several.

Also, note that the restriction to finite games is important. For games in which there are infinitely many actions, it is possible that there do not exist subgame equilibria.

6.4.2 Chess

One conclusion we may draw from the result relates to chess. Under the rule that a game is a draw once a position is repeated three times, chess is finite, so that Proposition 6.4.1 implies that it has a subgame perfect equilibrium and hence also a Nash equilibrium. (Under the (official) rule that a player has the option of declaring a draw under these circumstances, chess is not finite.) Because chess is strictly competitive, we saw previously that the equilibrium payoff is unique and that any Nash equilibrium strategy of a player guarantees the player his equilibrium payoff. Thus either White has a strategy that guarantees that it wins, or Black has a strategy that guarantees that it wins, or each player has a strategy that guarantees that the outcome of the game is either a win for it or a draw. The empirical evidence suggests that Black does not have a winning strategy but this remains an unproven result.

6.5 Exercises for Extensive form Subgame Perfect Equilibria

6.5.1 Exercise: (Asymmetric Game)

Consider the extensive form game shown in Figure 6.28.

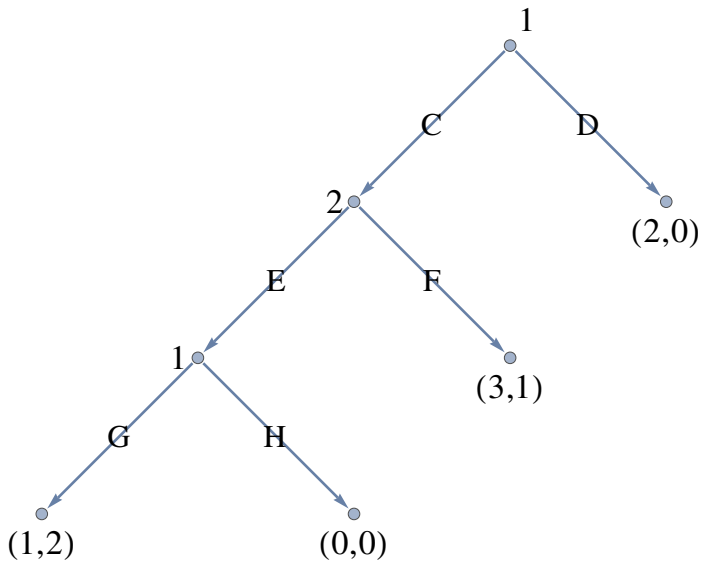

Figure 6.28: Asymmetric Extensive Form Game.

(i) Enumerate the full list of pure strategies for each player and form the equivalent strategic form game matrix.

(ii) Find the Nash Equilibria of this matrix.

(iii) Using the method of subgame perfect equilibria, find the optimum players' strategy pair.

Answer:

(i)

The pure strategies for Player 2 are simply `E` and `F`.

Player 1 potentially goes twice in this game with 2 options at each of her nodes, and so her pure strategies are :

`CG, CH, DG, DH`.

So the strategic form game matrix will be a 4×2 matrix as follows:

		"2"	
		E	F
"1"	CG	$(1, 2)$	$(3, 1)$
	CH	$(0, 0)$	$(3, 1)$
	DG	$(2, 0)$	$(2, 0)$
	DH	$(2, 0)$	$(2, 0)$

Table 6.5: Strategic form payoff matrix for the Asymmetric Game.

(ii) This is not a zero sum game, so the Nash equilibria need to be found by inspection and investigating whether a player can benefit by unilaterally changing their strategy. By inspection we find 3 standard Nash equilibria which are shown in Table 6.6.

		"2"	
		E	F
"1"	CG	$(1, 2)$	$(3, 1)$
	CH	$(0, 0)$	$(3, 1)$
	DG	$(2, 0)$	$(2, 0)$
	DH	$(2, 0)$	$(2, 0)$

Table 6.6: Strategic form payoff matrix for the Asymmetric Game showing the Nash Equilibria circled.

(iii) Apart from the root node, there are two other non terminal vertices, and so there are two other subgames, shown in Figure 6.29.

Exploring the subgame of length 1 in the figure we see that it is owned by Player 1, and so the optimum choice would be `G`. Then, exploring the subgame of length 2, which is owned by Player 2, we note that the optimum choice for Player 2, *given the optimum choice of Player 1 in the subgame of length 1*, is `E`. Why is this? Initially this might seem a risky choice since choosing option `E` opens up Player 2 to the risk of Player 1 choosing option `H` which would give Player 2 the worst possible outcome of a zero payoff. However, because we already know that *if* this subgame is entered, then Player 1 would play option `G` every time. So Player 2 can be safe in the knowledge that `E` is the

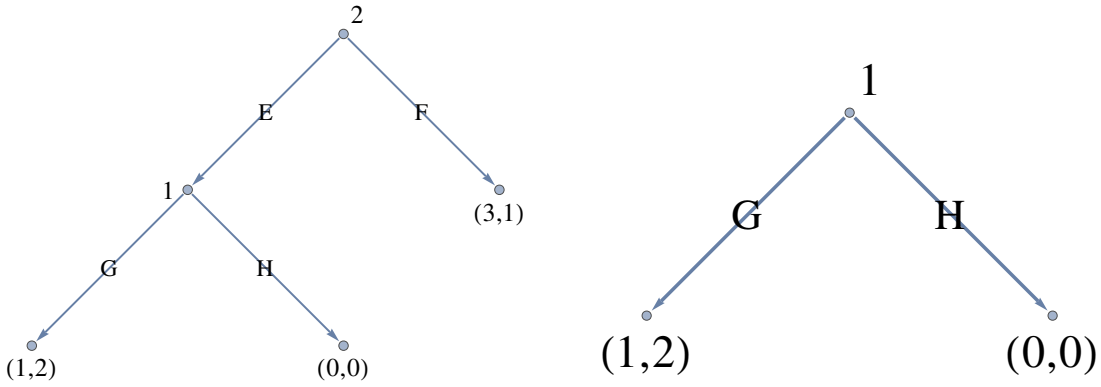

Figure 6.29: The two proper subgames of Figure 6.28.

optimum choice for the subgame of length 2 in Figure 6.29. Therefore in the last 'subgame' of length 3 (ie the full game owned by the root node of Player 1), *given* the optimum choices of the players in the other subgames, Player 1 would choose `D` as this would give her the best payoff.
So, in summary, using backward induction the optimum strategy pair of the players would be `(DG,E)` (even though player 1 never gets to play choice G note!!).

6.5.2 Exercise: (Voting by Alternating Veto)

Two people select a policy that affects them both by alternately vetoing policies that they dislike until there is just one policy left. First, person 1 vetoes a policy. If more than one policy remains, person 2 vetoes a policy. If more than one policy remains person 1 vetoes another policy, and so on until only one policy remains.

Suppose there are 3 possible policies, `X,Y,Z`, and person 1 prefers `X` to `Y` to `Z`. Person 2 prefers `Z` to `Y` to `X`.

Model this situation as an extensive form game, find its Nash equilibria, and solve it using subgame perfect equilibria to show that policy `Y` is the one that will be elected.

Answer: To model it as an extensive form game, we need to define the players (which is obvious), the players' possible options,and construct the extensive form game tree. We can then find the Players pure strategies, find the equivalent strategic form game matrix, thus find the Nash equilibria, and then solve corresponding subgames for optimum player solutions.
So, the Players, are '1' and '2', and we know Player 1 goes first.
We now need to deal with the players' preferences. The easiest way to do this is to assign numerical payoff values that represent the preferences since it is only the ordering that matters and not the absolute values.
We know that

- Player 1: `X` $\succ$ `Y` $\succ$`Z`
- Player 2: `X` $\prec$ `Y` $\prec$ `Z`

so selecting payoffs of $(X, Y, Z) = (2, 1, 0)$ for Player 1 and $(X, Y, Z) = (0, 1, 2)$ for Player 2 would satisfy these preferences.(Other values would be equally as good as long as the ordering was preserved).
We now need to construct the extensive form game tree. Player 1 can choose to veto `X`, `Y`, or `Z`. ie Player 1 has a choice of 3 options. Once Player 1 has gone, Player 2 has a choice of two options, eliminating one of the options that Player 1 has left. This would leave just one option which is therefore the selected policy and the game has ended.
So expressing this as a graphical tree we would get Figure 6.30.

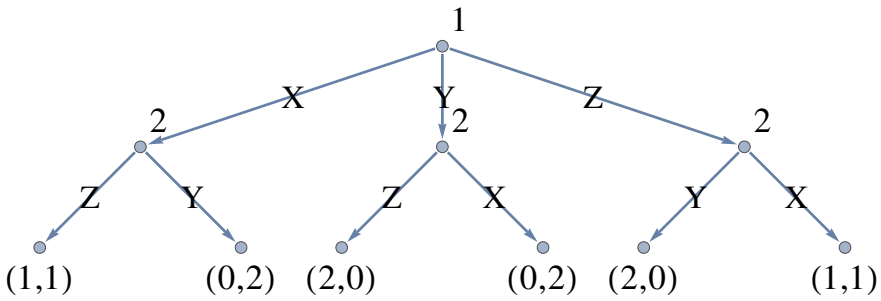

Figure 6.30: The gametree of Voting by Alternating Veto Game.

From this gametree we can now more easily see that Player 1 has pure strategies of `X,Y,Z` and Player 2 has to consider pure strategies conditioned on whichever node the game ends up at. So Player 2 has pure strategies:
`(YXX), (YXY), (YZX), (YZY), (ZXX), (ZXY), (ZZX), (ZZY)`
So the equivalent strategic form game matrix will be of size 3×8. Using the gametree we can construct the strategic form game matrix as shown in Table 6.7.

The Table also shows the Nash Equilibria circled. These have been found by observing for each payoff, whether a player can improve their payoff by unilaterally switching strategies. We note there are two Nash equilibria in this strategic form game. So we now use subgame perfect equilibria to see whether

		"2"							
		YXX	YXY	YZX	YZY	ZXX	ZXY	ZZX	ZZY
"1"	X	(0, 2)	(0, 2)	(0, 2)	(0, 2)	(1, 1)	(1, 1)	(1, 1)	(1, 1)
	Y	(0, 2)	(0, 2)	(2, 0)	(2, 0)	(0, 2)	(0, 2)	(2, 0)	(2, 0)
	Z	(1, 1)	(2, 0)	(1, 1)	(2, 0)	(1, 1)	(2, 0)	(1, 1)	(2, 0)

Table 6.7: Strategic form payoff matrix for the Voting By Alternating Veto Game showing the Nash Equilibria circled.

either of these Nash equilibria are unstable (ie are not going to appear in the extensive form game plan). To explore the subgame perfect equilibria, we note there are 3 proper non terminal vertices, and so there are 3 proper subgames, each of length 1 and shown in Figure 6.31:

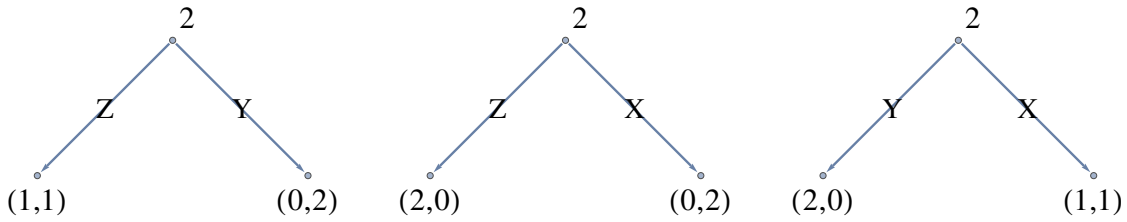

Figure 6.31: The three subgames of the Voting by Alternating Veto Game.

Each subgame is owned by Player 2. The first subgame has an optimum strategy of `Y` for Player 2. The second subgame has an optimum strategy choice of `X`. The third subgame has an optimum choice of `X`.

So the next level subgame of length two is the entire game, which has a root node owned by Player 1. What will Player 1 do? Player 1 knows the optimum choices to be made by Payer 2, given that the game progresses to each of the vertices. So the optimum choice for Player 1 is to choose strategy `Z` since this would give the maximum payoff to Player 1, *given* the strategic choices of Player 2.

Therefore the subgame perfect equilibrium is: `(Z,YXX)`. In other words, the resulting policy which remains is policy Y.

6.5.3 Exercise: (Multiple Optimal Actions)

So far we have only considered examples with just one subgame perfect equilibrium, even if there were multiple Nash equilibria. In this example we consider a case where one of the players has several optimal actions. We will use subgame perfect equilibria to show how these multiple optimal solutions can be found, and indicate that it is not guaranteed that there is ony one unique subgame perfect equilibrium.

Consider the extensive form gametree shown in Figure 6.32. Construct the Players pure strategies and use backward induction to show there are multiple optimal subgame equilibria for Player 1.

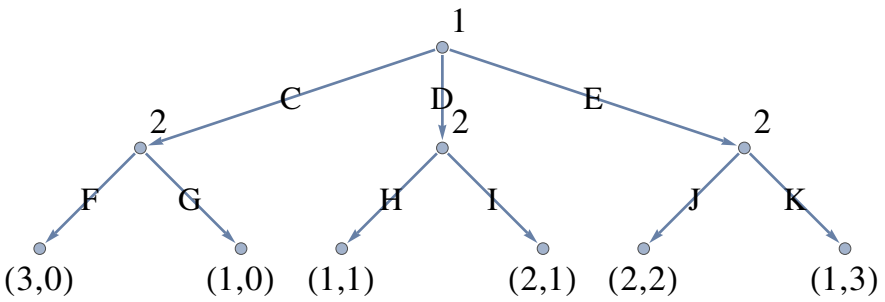

Figure 6.32: Extensive form gametree with multiple genuine equilibria.

Answer: Player 1 has pure strategies `C,D,E`. Player 2 has two choices at each of three separate vertices, which depend on the choice made by Player 1. So Player 2's pure strategies are
`(FHJ),(FHK),(FIJ),(FIK),(GHJ),(GHK),(GIJ),(GIK)`.
If we examine the subgames of length 1 in this game we see the subgames given in Figure 6.33:

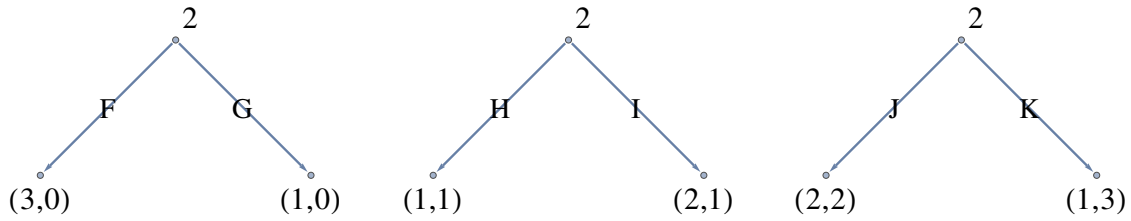

Figure 6.33: Extensive form subgames of length 1 of the extensive form gametree in Figure 6.32.

If we examine the first subgame, which is owned by Player 2, then we note that actually in this case Player 2 is *indifferent* to either of the strategies, ie either of the choices `F` or `G` could be made to give the same payoff. Similarly, examining the second subgame to choose between `H` and `J` again we note that Player 2 is indifferent to either choice. In the third subgame, Player 2 would choose strategy `K` since this gives the maximum payoff.
So Player 2 has 4 combinations of optimal strategies:

- `(FHK),(FIK),(GHK),(GIK)`.

If we now look at subgames of length 2 (ie the whole gametree owned by Player 1), Player 1's optimal move must depend on Player 2's optimal moves for *all* combinations of Player 2's strategies. By working from Player 2's optimal strategies we find that Player 1 has several optimal strategies. In particular:

- For `(FHK)` and `(FIK)` the optimal move for Player 1 is `C`.
- For `(GHK)` *ALL* `C,D,E` options are equally optimal in that they all give the same payoff to Player 1.
- For `(GIK)` the optimal choice for Player 1 is `D`.

So, this process of backward induction gives the following set of perfect subgame equilibria which are all equally valid.

- `(C,FHK),(C,FIK),(C,GHK),(D,GHK),(E,GHK),(D,GIK)`.

So this example shows that there are games with multiple equivalent perfect subgame equilibria.

6.5.4 Exercise: (Stackelberg's Model of Duopoly)

We can use the technique of subgame perfect equilibria to explore games where the payoffs are potentially continuous as well. We consider here a famous example exploring two firms in a market where one firm enters the market after another firm, and knows of the first firm's strategy. This is a development of another famous model known as Cournot's model of Duopoly in which the two firms make strategic decisions simultaneously. However in this game we consider the extensive form version.
In this game there are two firms producing some volume of goods. Firm i's cost of producing q_i goods is $C_i(q_i)$ ($i = 1, 2$ in this example). The price at which the total output Q can be sold is denoted $p(Q)$.
The terminal payoff $\pi_i(q_1, q_2)$ to each company is the *profit*, ie the cost of selling the goods minus the cost of production.

$$\pi_i(q_1, q_2) = q_i p(q_1 + q_2) - C_i(q_i)$$

So, given that Firm 1 goes first and decides to make q_1 volume of goods, Firm 2 then decides, knowing this, what volume of goods q_2 to manufacture. Thus in this example a strategy of Firm 2 is a *function* that associates an output for Firm 2 with each possible output of Firm 1.
Use the method of backward induction to find the subgame perfect equilibrium of both Firms (ie what volume of goods q_1, q_2 will they produce) in this game assuming that the cost of producing goods is linear

$$C_i(q_i) = cq_i$$

and there is a linear inverse demand function for the price of the goods such that:

$$p(Q) = \begin{cases} \alpha - Q & \text{if } Q \leq \alpha \\ 0 & \text{if } Q > \alpha \end{cases}$$

This latter constraint prevents overproduction since if the total goods produced are more that α then the price of selling them is zero.

Answer:

First, we assume Player 1 produces q_1 goods, and Player 2 therefore chooses an optimum strategy (which is a function of q_1 recall) of $q_2 = b_2(q_1)$. We then need to determine the optimum choice of q_1 *given* the optimum choice of $b_2(q_1)$.

So we first need to explore Player 2's optimum strategy.

Player 2's profit would be

$$\pi_2(q_1, q_2) = q_2[p(q_1 + q_2)] - c] = \begin{cases} q_2(\alpha - c - q_1 - q_2) & q_1 + q_2 < \alpha \\ -cq_2 & q_1 + q_2 > \alpha. \end{cases}$$

or, rearranging,

$$\pi_2(q_1, q_2) = \begin{cases} -q_2^2 - (q_1 + c - \alpha)q_2 & q_1 + q_2 \leq \alpha \\ -cq_2 & q_1 + q_2 > \alpha. \end{cases}$$

So the optimum solution corresponds to maximising this function, ie looking for the solutions of $\frac{d\pi_2}{dq_2} = 0$ gives:

$$\begin{aligned} \frac{d\pi_2}{dq_2} = 0 \Rightarrow -2q_2 &= (q_1 + c - \alpha) \text{ if } q_1 < \alpha - c \\ \Rightarrow q_2 &= \frac{1}{2}(\alpha - c - q_1) \end{aligned}$$

This is true if $q_1 < \alpha - c$ since q_2 needs to be positive. So this is the optimum strategy for Firm 2. But it depends on the strategy of Firm 1. So now working backwards, Firm 1 needs to choose an optimum strategy (ie what to choose for q_1) , now in the knowledge of what Firm 2's optimum strategy would be.

Firm 1's profit is

$$\begin{aligned} \pi_1(q_1, q_2) &= q_1 p(q_1 + b_2(q_1)) - cq_1 \\ &= q_1 \times (\alpha - q_1 - \frac{1}{2}(\alpha - c - q_1)) - cq_1 \\ &= \alpha q_1 - q_1^2 - \frac{\alpha}{2}q_1 + \frac{c}{2}q_1 + \frac{1}{2}q_1^2 - cq_1 \\ \pi_1 &= -\frac{1}{2}q_1^2 + \frac{1}{2}(\alpha - c)q_1 \end{aligned}$$

This is maximum when $d\pi_1/dq_1 = 0$.

This is satisfied when $q_1 = (\alpha - c)/2$.

So this is the optimal strategy of Firm 1. Therefore the optimal strategy of Firm 2 can

now be found. Firm 2 produces $q_2 = b_2(q_1)$ goods, ie

$$\begin{aligned} q_2 &= \frac{1}{2}(\alpha - c - q_1) \\ &= \frac{1}{2}(\alpha - c - \frac{1}{2}(\alpha - c)) \\ &= \frac{1}{4}(\alpha - c) \end{aligned}$$

So this is Firm 2's optimum strategy.
The joint optimum strategy has Firm 1 producing twice as many goods as Firm 2. Any deviation from these values will result in reduced payoff from whichever Firm chooses to deviate! From these we can work out the profit of Firms 1 and 2.

$$\pi_1 = q_1 p(q_1 + q_2) - cq_1 = \frac{1}{8}(\alpha - c)^2$$
$$\pi_2 = q_2 p(q_1 + q_2) - cq_2 = \frac{1}{16}(\alpha - c)^2$$

So this shows how backward induction can be used to find the subgame perfect equilibrium in a game in which the payoffs potentially form a continuum rather than a finite discrete set.

Chapter 7

Incomplete Information: Bayesian Games

So far we have assumed that all players had perfect information regarding the elements of a game. These are called games with complete information.

In many game theoretic situations, one agent is unsure about the preferences or intentions of others. Incomplete information introduces additional strategic interactions and also raises questions related to learning.

Examples:

- Bargaining (how much the other party is willing to pay is generally unknown to you)
- Auctions (how much should you bet for an object that you want, knowing that others will also compete against you?)
- Market competition (firms generally do not know the exact cost of their competitors)
- Signalling games (how should you infer the information of others from the signals they send)
- Social learning (how can you leverage the decisions of others in order to make better decisions)

A game with incomplete information, on the other hand, tries to model situations in which some players have private information before the game begins. The initial private information is called the type of the player. For example, types could be the privately known valuations of an object in an auction, private information on a company strength or investement plan. We consider first *strategic* games, and then *extensive form* games.

7.1 Strategic Games

7.1.1 Preliminaries

A Bayesian game is a strategic form game with incomplete information. It consists of:

- A set of players I;

- A set of actions (pure strategies) for each player i: A_i ;
- A set of types for each player i: $\theta_i \in \Theta_i$;
- A payoff function for each player i: $u_i(a_1, ..., a_I, \theta_1, ..., \theta_I)$;
- A (joint) probability distribution $p(\theta_1, ..., \theta_I)$ over types (or $P(\theta_1, ..., \theta_I)$ when types are not finite).

More generally, one could also allow for a signal for each player, so that the signal is correlated with the underlying type vector.

Importantly, throughout in Bayesian games, the strategy spaces, the payoff functions, possible types, and the prior probability distribution are assumed to be common knowledge. Very strong assumption. But very convenient, because any private information is included in the description of the type and others can form beliefs about this type and each player understands others' beliefs about his or her own type, and so on, and so on.

The function p_i summarises what player i believes about the types of the other players given her type. So, $p_i(\theta_{-i}|\theta_i)$ is the conditional probability assigned to the type profile $\theta_{-i} \in \Theta_{-i}$.

Similarly, $u_i(a|\theta_i)$ is the payoff of player i when the action profile is a and the type profile is θ_i. We call a Bayesian game finite if N, A_i, Θ_i are all finite, for all $i \in I$.

A (pure) strategy for a player prescribes an action for each possible type of player.

A pure strategy for player i in a Bayesian game is a function which maps player i's type into her action set $a_i : \Theta_i \to A_i$, so that $a_i(\theta_i)$ is the action choice of type θ_i of player i:

A mixed strategy for player i is $\alpha_i(\Theta_i) \to \Delta(A_i)$ so that $\alpha_i(a_i|\theta_i)$ is the probability assigned by α_i to action a_i by type θ_i by player i.

Recall that player types are drawn from some prior probability distribution $p(\theta_1, ..., \theta_I)$. Given $p(\theta_1, ..., \theta_I)$ we can compute the conditional distribution $p(\theta_{-i}|\theta_i)$ using Bayes' rule. Hence the label "Bayesian games".

Equivalently, when types are not finite, we can compute the conditional distribution $P(\theta_{-i}|\theta_i)$ given $P(\theta_1, ..., \theta_I)$. Player i knows her own type and evaluates her expected payoffs according to the the conditional distribution $p(\theta_{-i}|\theta_i)$, where $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$.

Since the payoff functions, possible types, and the prior probability distribution are common knowledge, we can compute expected payoffs of player i of type θ_i as

$$U(a_i', a_{-i}(.), \theta_i) = \sum_{\theta_{-i}} p(\theta_{-i}|\theta_i) u_i(a_i', a_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

when types are finite, and

$$\int u_i(a_i', a_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i}|\theta_i)$$

when types are not finite.

Suppose there are two players, player 1 and 2 and for each player there are two possible types. Player i's possible types are θ_i and $\bar{\theta}_i$. Suppose that the types are independently distributed and the probability of θ_i is q and the probability of $\bar{\theta}_i$ is $(1 - q)$. For a given pure strategy profile a^* the expected payoff of player 1 of type θ_1 is

$$q u_1(a_1^*(\theta_1), a_2^*(\theta_2)|\theta_1, \theta_2) + (1 - q) u_1(a_1^*(\theta_1), a_2^*(\bar{\theta}_2)|\theta_1, \bar{\theta}_2)$$

Similarly, for a given mixed strategy profile α^* the expected payoff of player 1 of type θ_1 is:

$$q \sum_{a \in A} \alpha_1^*(a_1|\theta_1^*)\alpha_2^*(a_2|\theta_2^*)u_1(a_1, a_2|\theta_1, \theta_2) + (1-q) \sum_{a \in A} \alpha_1^*(a_1|\theta_1^*)\alpha_2^*(a_2|\bar{\theta}_2^*)u_1(a_1, a_2|\theta_1, \bar{\theta}_2)$$

7.1.2 Bayes-Nash Equilibrium

(Bayes–Nash Equilibrium).
The strategy profile $\boldsymbol{a}$ is a (pure strategy) Bayes–Nash equilibrium if for all $i \in I$ and for all $\theta_i \in \Theta_i$, we have that

$$a_i(\theta_i) \in argmax_{a_i \in A_i} \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i)u_i(a_i', a_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

or in the non-finite case,

$$a_i(\theta_i) \in argmax_{a_i \in A_i} \int u_i(a_i', a_{-i}(\theta_{-i}), \theta_i, \theta_{-i})P(d\theta_{-i}|\theta_i)$$

Hence a Bayes–Nash equilibrium is a Nash equilibrium of the "expanded game" in which each player i's space of pure strategies is the set of maps from $\theta_i \to A_i$.

Remark 16. *Type, in general, can be any private information that is relevant to the player's decision making, such as the payoff function, player's beliefs about other players' payoff functions, her beliefs about what other players believe her beliefs are, and so on.*

Remark 17. *Notice that, in the definition of a Bayesian equilibrium we need to specify strategies for each type of a player, even if in the actual game that is played all but one of these types are non-existent. This is because, given a player's incomplete information, analysis of that player's decision problem requires us to consider what each type of the other players would do, if they were to play the game.*

7.1.2.1 Existence of Bayes–Nash Equilibria

Theorem 9. *Consider a finite incomplete information (Bayesian) game. Then a mixed strategy Bayes–Nash equilibrium exists.*

Theorem 10. *Consider a Bayesian game with continuous strategy spaces and continuous types. If strategy sets and type sets are compact, payoff functions are continuous and concave in own strategies, then a pure strategy Bayes–Nash equilibrium exists.*

The ideas underlying these theorems and proofs are identical to those for the existence of equilibria in (complete information) strategic form

Exercise: Illustration using the Battle of the Sexes Recall the Battle of the Sexes (BoS) example, which was a complete information "coordination" game. Both parties want to meet, but they have different preferences on "Cinema" and "Football". The players agree that it is better to cooperate, but they disagree about the best outcome for each of them.

This strategic game is represented in Figure 7.1.

In this game there are two pure strategy equilibria (one of them better for player 1 and the other one better for player 2), and a mixed strategy equilibrium.

	Football	Cinema
Football	$2, 1$	$0, 0$
Cinema	$0, 0$	$1, 2$

Figure 7.1: Payoff matrix for the Battle of the Sexes

Now imagine that player 1 does not know whether player 2 wishes to meet or wishes to avoid player 1. Therefore, this is a situation of incomplete information – also sometimes called asymmetric information.

We represent this by thinking of player 2 having two different types: one type that wishes to meet player 1 and the other wishes to avoid him. More explicitly, suppose that these two types have probability 1/2 each. Then the game takes the form of one of the following two possible games with probability 1/2:

	Football	Cinema		Football	Cinema
Football	$2, 1$	$0, 0$	Football	$2, 0$	$0, 2$
Cinema	$0, 0$	$1, 2$	Cinema	$0, 1$	$1, 0$

Figure 7.2: Payoff matrices for the two different types of Player 2 in the Battle of Sexes: The left hand one is the standard player, the right hand one is the player who wants to avoid Player 1

Crucially, player 2 knows which game it is (she knows the state of the world), but player 1 does not.

What are strategies in this game?

Most importantly, from player 1's point of view, player 2 has two possible types (or equivalently, the world has two possible states each with probability 1/2 and only player 2 knows the actual state). How do we reason about equilibria here?

From the previous definitions, we use the Nash Equilibrium concept in an expanded game, where each different type of player 2 has a different strategy. or equivalently, form conjectures about other player's actions in each state and act optimally given these conjectures.

Let us consider the following strategy profile `(C, (C, F ))`, which means that player 1 will play C, and while in state 1, player 2 will also play C (when she wants to meet player 1) and in state 2, player 2 will play F (when she wants to avoid player 1). Clearly, given the play of C by player 1, the strategy of player 2 is a best response.

Let us now check whether player 1 is also playing a best response to Player 2. Since both states are equally likely, the expected payoff of player 1 is

$$U_1[\texttt{C, (C, F)}] = \frac{1}{2} \times 1 + \frac{1}{2} \times 0 \equiv \frac{1}{2}.$$

If, instead, he deviates and plays `F` , his expected payoff is

$$U_1[\texttt{F , (C, F)}] = \frac{1}{2} \times 0 + \frac{1}{2} \times 2 \equiv 1.$$

Therefore, the strategy profile `(C, (C, F ))` is not a (Bayesian) Nash equilibrium.

7.2 Bayesian Extensive Form Games: (Dynamic Games with Imperfect Information)

7.2.1 Information Sets

Recall, that an extensive form game with perfect information is one in which each player knows everything that has happened before their turn and moves alone, i.e., not simultaneously. Chess is such a game.

We previously met the definition of an extensive form game in the form of a complete-information game tree reproduced as follows: (Game Tree with Complete Information and No Chance Moves).
A game tree with complete information and no chance moves is defined as the tuple $G = (T, P, S, \nu, \mu, \pi)$ such that T is a directed tree, ν is a player vertex assignment on intermediate vertices of T , μ is a move assignment on the edges of T and π is a payoff function on T.

To simplify the definition, we combine together the notion of the Tree, T, the set of possible actions S, and the move assignment function μ into a notion of a player's History on the game tree. A player can have a specific history $h_i \in H$ from the set of all possible histories, H.

We only introduce one new idea to Extensive Form Games with Perfect Information: sets of histories. We call these sets of histories "information sets" because a player regards histories within her information set as histories which could actually have occurred and all those outside her information set as those histories which did not occur. It is as if the player received information about what happened at an information set. Actually, it would be more accurate to say that a player has been told what did not happen at her information set; what is excluded in her information set. She regards any history within her information set as possible and her beliefs about the relative likelihood is governed by Bayes' Rule when that information set is reached in that equilibrium. In other words, whenever player i has a decision to make,i.e., has a turn, she is informed that some history has occurred but is not informed which one.

So in summary, an *information set* is a set that, for a particular player, establishes all the possible moves that could have taken place in the game so far, given what that player has observed. If the game has perfect information, every information set contains only one vertex, namely the point actually reached at that stage of the game. Otherwise, it is the case that some players cannot be sure exactly what has taken place so far in the game and what their position is. More specifically, in the extensive form, an information set is a set of decision nodes such that: Every node in the set belongs to one player, and, when play reaches the information set, the player with the move cannot differentiate between nodes within the information set, i.e. if the information set contains more than one node, the player to whom that set belongs does not know which node in the set has been reached. The notion of information set was introduced by John von Neuman motivated by studying the game of Poker.

To make this idea of information set unambiguous, we need some notation. Let H_i , the set of histories after which player i moves be divided up into a collection of information sets I_i such that every history where player i moves is an element of player i's information set, which is an element of Player i's "information partition".

Intuitively, a player's information partition is the set of a player's information sets. It contains all of the sets of histories which a player can discern in a game. (It is called a "partition" because every history must be contained in one and only one of those information sets.) The actions available to a player will no longer be contingent on the history which has occurred but upon the set of histories which a player

thought could have occurred at her information set.

$$A(I_i) = \{A(h) : h \in I_i\}$$

For each player $i \in P$ a partition I_i of the set of histories at which player i has a turn $\{h \in H : \mu(h) = i\}$ has the property that for histories contained within the information set, the player has the same set of available actions: If $h; h' \in I_i$ then $A(h) = A(h')$.

The reason for this is that if a player had different actions for different histories within the same partition, she should be able to tell which history had actually occurred.

The introduction of the concept of information set necessitates a change in our definition of strategy. Recall that a pure strategy in an extensive game is a function which assigns to each i's information set an action in $A(I_i)$.

So, a dynamic game of incomplete information consists of

- A set of players I;
- A sequence of histories H_t at the t-th stage of the game, each history assigned to one of the players (or to Nature);
- An information partition, which determines which of the histories assigned to a player are in the same information set.
- A set of (pure) strategies for each player i, A_i, which includes an action at each information set assigned to the player.
- A set of types for each player i: $\theta_i \in \Theta_i$;
- A payoff function for each player i: $u_i(a_1, ..., a_I, \theta_1, ..., \theta_I)$;
- A (joint) probability distribution $p(\theta_1, ..., \theta_I)$ over types.

7.2.2 Imperfect Information Extensive Form Games: Illustrations

Dynamic games are always represented in the extensive form (tree form). With imperfect information, these trees will include information sets. We will use the notation in the following figures that an information set is all nodes that are connected by a dotted line. The information set indicates that the player taking an action at that point of the game does not know which node has been reached. Consider figure 7.3 in which some nodes are in information sets and others are not.

A subgame is defined as all the branches that begin at a node not in an information set and continue until the end of the game. Thus, to count the number of subgames, one must simply count the number of nodes not in an information set. The game shown in in Figure 7.3 has 4 subgames, as shown in Figure 7.4; ie the games following nodes 8,7,2,1 (recall that the entire game is always a subgame).

The more generic definition of a subgame perfect Nash equilibrium states that it is the set of strategies for all players such that a Nash equilibrium is played in each *subgame*. If the game has perfect information, then this definition is equivalent to the one given previously (where a subgame perfect Nash equilibrium is the set of strategies such that a Nash equilibrium is played at all nodes).

We now consider finding subgame perfect Nash equilibria for extensive form games with imperfect information.

7.2.2.1 Illustration

Consider the two extensive form games shown in Figures 7.5 and 7.6, the first having only perfect information and the second having imperfect information, as indicated by the dashed line of Player 2 who is unsure which node in the game she is playing. These are two different games, with different strategies and different normal form equivalents.

For the game corresponding to Figure 7.5, Player 1 has a strategy set of `{L, R }` whilst Player 2 has strategy set of `{(LL,LL),(LL,RR),(RR,LL),(RR,RR)}`. The subgame perfect Nash equilibrium (found using backward induction) is `(R, (RR, LL))`, with equilibrium payoffs of $(2, 1)$.

For the game corresponding to Figure 7.5, we need to find the subgame perfect equilibria. By definition, a subgame perfect Nash equilibrium is such that a Nash equilibrium is played in all subgames. So, to apply backward induction, we must find the lowest nodes in our game tree that are not in an information set. The lowest node satisfying this requirement in this figure is the top node, at which Player 1 must choose between `L` and `R`. The game proceeding from this node to the end is the one and only subgame. To find the optimal strategies for this one subgame, consider that Player 2 is taking an action without knowing the action taken by Player 1, and so her strategy set is `{LL,RR }` only. This leads to an equivalent normal form game matrix shown in Figure 7.7.

It can be seen that the subgame perfect Nash equilibrium is `(L, RR)` (verify this yourselves), giving equilibrium payoffs of $(1, 2)$.

Note that when the actions of Player 1 are observed as in Figure 7.5, then Player 1 has a higher payoff (2) compared to when the actions are not observed (Figure 7.6 with payoff 1). This is the concept of a First mover advantage.

Note that in the imperfect information game as the actions are not observed, it is irrelevant whether Player 1 is written on top or Player 2 is written on top.

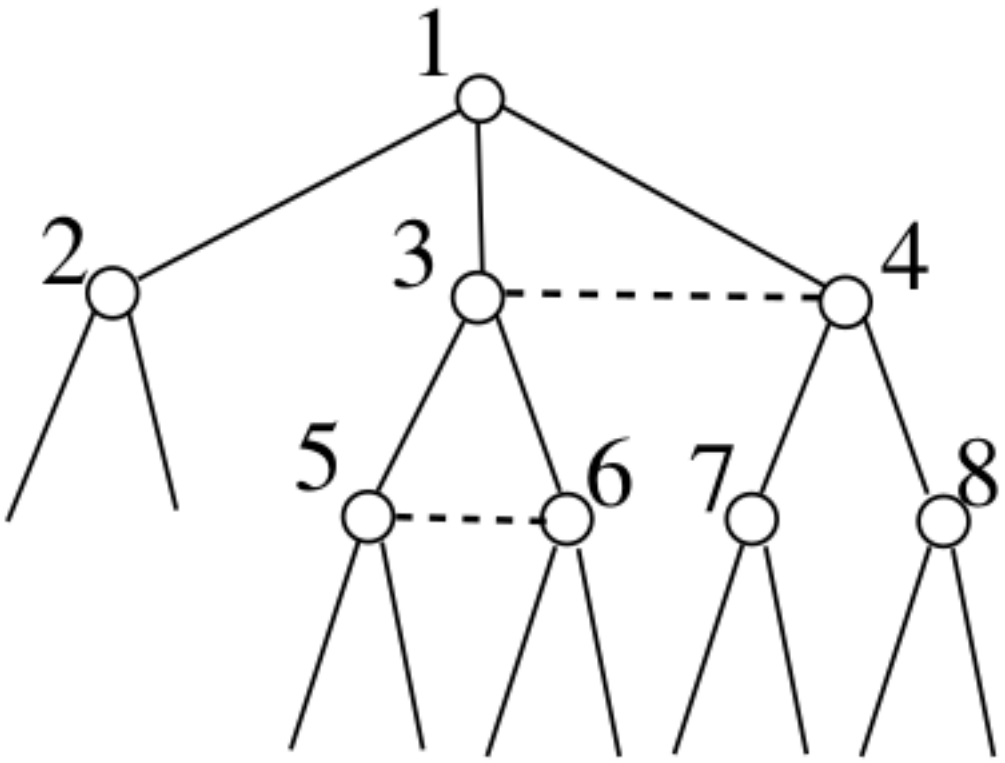

Figure 7.3: Imperfect information extensive form game tree. Information sets are indicted by a dashed line.

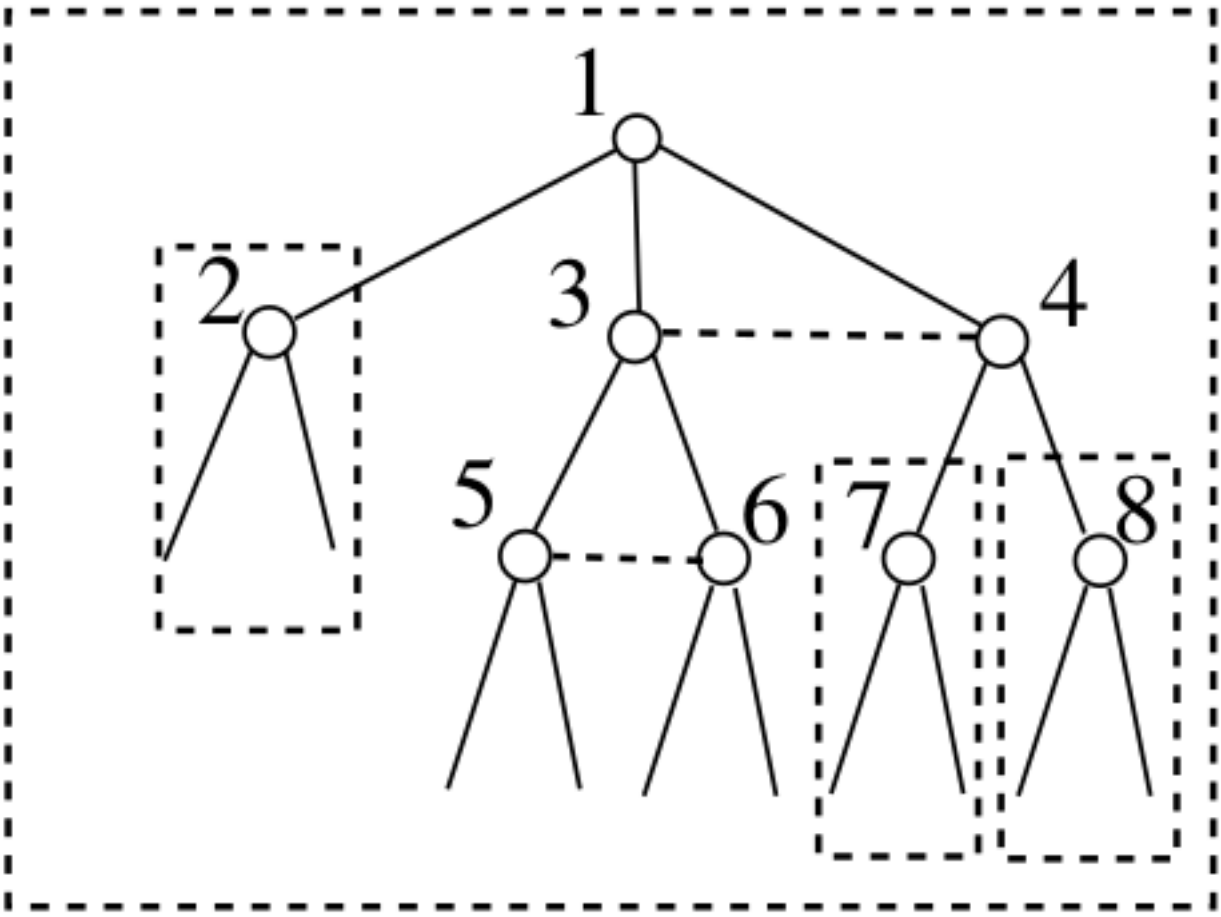

Figure 7.4: The 4 subgames of Figure 7.3 shown by dashed rectangles.

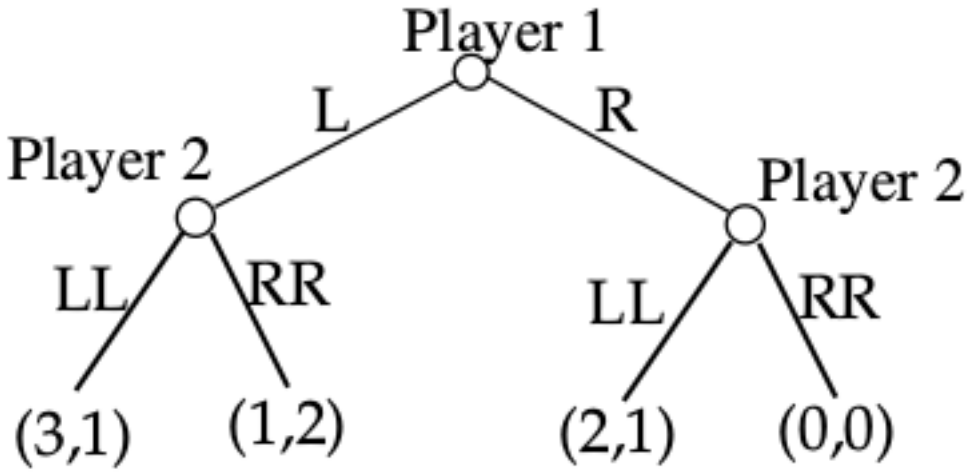

Figure 7.5: Simple 2 player extensive form game: perfect information.

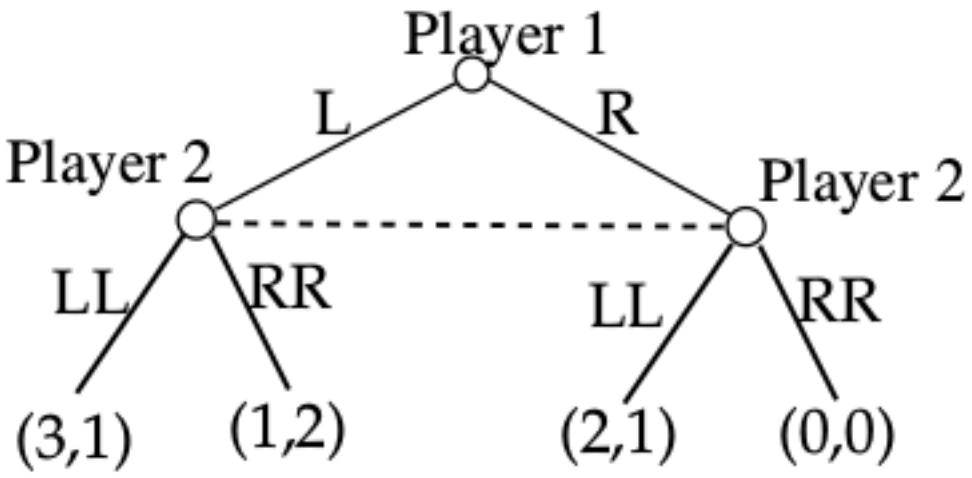

Figure 7.6: Simple 2 player extensive form game: **imperfect information**

	LL	RR
L	$3,1$	$1,2$
R	$2,1$	$0,0$

Figure 7.7: Equivalent normal form game for Figure 7.6

7.3 Exercises for Chapter VII

7.3.1 Strategic Games– Incomplete Information

Two students have been allocated to each other to work together on a joint poster assesment. Student 1 is either Lazy or Keen. Student 2 does not know which type of student, Student 1 is. Student 2 assumes that Student 1 will be a Keen type with probability p and a Lazy type with probability $(1-p)$. Each student has a choice of actions to either (i) work hard on the poster assessment (which we denote as W) or, (ii) don't work (which we denote as DW). Table 7.1 shows the payoffs for both students for the two different types of Student 1. The first number in each pair in the Table is the payoff to Student 1, the second number in the pair is the payoff for Student 2.

		Student 2				Student 2	
		W	DW			W	DW
Student 1	W	$2,2$	$0,3$	Student 1	W	$3,3$	$1,2$
[Lazy $(1-p)$]	DW	$3,0$	$1,1$	[Keen (p)]	DW	$2,1$	$0,0$

Table 7.1: Payoff matrices for the two different types of Student 1 in the poster assignment.

a) For each game in Table 7.1 determine the *pure* strategies of the students.

Answer: This is straightforwards. For the Lazy game the unique Nash equilibrium occurs at (DW,DW), since neither player can improve their payoff unilaterally by switching strategies ie neither student puts any effort in. Similarly, for the game in which one student is Keen, the Nash equilibrium is (W,W), so both students work hard. ie idleness is penalised, effort is rewarded.

b) Construct the extended normal form payoff matrix for this Bayesian game in terms of the strategies of Student 1 and Student 2.
(Hint: Each element in the extended payoff matrix should depend on the belief probability p.)

Answer: It is important to get the correct strategies of each student. Student 2 has simple strategies (W,DW). Student 1, because of the two different types of Student 1, has strategy choices of ((W,W),(W,DW),(DW,W),(DW,DW)), where the first symbol in each pair is the strategystudent 1 would follow if she were playing game 1, and the second symbol denotes the strategy she would follow if she were playing game 2. So the extended normal form payoff matrix will be a 4×2 matrix with values obtained by weighting the payoffs by the probabilities of each type of player 1 playing the W or DW action. The resulting payoff table is given in Table 7.2 and will take the student a good 5 minutes to work out:

c) Assume student 2 believes student 1 is keen with probability $2/3$. Using this value of p write down the extended payoff matrix and use this to determine the Bayes-Nash equilibrium solution of the game by considering the best response strategies of each player and type of player.

		Student 2	
		W	DW
	(W,W)	$p+2, p+2$	$p, 3-p$
Student 1	(W,DW)	$2, 2-p$	$0, 3-3p$
	(DW,W)	$3, 3p$	$1, 1+p$
	(DW,DW)	$3-p, p$	$1-p, 1-p$

Table 7.2: Payoff matrix for the extended normal form student poster assessment game.

Answer: Substituting $p = 2/3$ into Table 7.2 is straightforward and gives Table 7.3.

		Student 2	
		W	DW
	(W,W)	$8/3, \overline{8/3}$	$2/3, 7/3$
Student 1	(W,DW)	$2, \overline{4/3}$	$0, 1$
	(DW,W)	$\underline{3}, \overline{2}$	$\underline{1}, 5/3$
	(DW,DW)	$7/3, \overline{2/3}$	$1/3, 1/3$

Table 7.3: Payoff matrix for the extended normal form student poster assessment game using $p = 2/3$.

Given Table 7.3, we can consider the best response of player 1 to either choice of W and DW of player 2. These are shown as underlined payoffs in the Table Then for each of player 1's choices we can then consider which is the best response of player 2, to player 1 choosing (W,W),(W,DW), (DW,DW),(DW,DW). These are shown as overlines over the payoffs in the Table. Following this we find he Bayes-Nash equilibrium is ((DW,W),W) and is circled in the Table 7.3.

d) Using your answer in b), determine what is the threshold belief value of p, above which Student 2 will choose to Work rather than Don't Work.

Answer: For Student 2 to have an incentive to Work rather than Dont Work, the payoffs for that student need to be higher for W than DW. Looking at the elements in Table 7.2 this means that p needs to satisfy

$$\begin{aligned} p + 2 &\geq 3 - p \\ 2 - p &\geq 3 - 3p \\ 3p &\geq 1 + p \\ 2p &\geq 2 - 2p \end{aligned}$$

The consistent solution to all of these equations occurs at the unique critical value of $p = 1/2$.

e) If we tell Student 2 which type of person Student 1 is, Student 2 can change their strategies to increase their utility. Is it always the case in Bayesian games that giving one player more information (such as the state of the other player) will increase their own utility?

Answer: The counterintuitive answer is No. The challenge is to find an example in the on-line literature!!

(BMW vs Mercedes)

BMW and Mercedes are both planning to release a new high quality, comparable car. BMW are sure of their own financial strength, but they are unsure as to whether Mercedes is in a 'strong' or a 'weak' financial position following the recession. Both firms have the choice of whether to compete on price or not. The payoffs for each outcome are listed in Table 7.4.

If both firms fight then the winner (Mercedes if they are in strong position and BMW otherwise) gains a payoff of 1 and the loser receives a loss of 2 units. If neither firm fights then the payoff will be 0 to both. If one firm fights and the other does not then the payoffs will be 2 units and -1 unit respectively.

(i) Assuming BMW have a belief that Mercedes is strong with probability p and weak with probability $1 - p$ construct the normal form payoff matrices for both types of both players.

(ii) From the strategies of BMW and Mercedes, construct the extended normal form matrix, and use this to find the critical value, p^* under which BMW should choose to fight and above which they should not fight.

(iii) Find the Bayes–Nash equilibrium and the payoffs when $p = p^*$.

Answer:

		BMW	
		fight	don't fight
Mercedes	fight	$1, -2$	$2, -1$
[strong (p)]	don't fight	$-1, 2$	$0, 0$

		BMW	
		fight	don't fight
Mercedes	fight	$-2, 1$	$2, -1$
[weak $(1-p)$]	don't fight	$-1, 2$	$0, 0$

Table 7.4: Payoff matrices for the two different types of Mercedes and the one type of BMW.

(i)

(ii) The strategies of BMW are simply (fight, don't fight), abbreviated to (F,N). The strategies of Mercedes have to take into account the two different types of player they might be and so the Mercedes strategies are (fight, fight), (fight,don't fight), (don't fight, fight) and (don't fight, don't fight). Therefore the extended normal form matrix is given in Table 7.5.

		BMW	
		F	N
	(F,F)	$3p-2, 1-3p$	$2, -1$
Mercedes	(F,N)	$2p-1, 2-4p$	$2p, -p$
	(N,F)	$p-2, 1+p$	$2-2p, p-1$
	(N,N)	$-1, 2$	$0, 0$

Table 7.5: Payoff matrix for the extended normal form game.

So for BMW to have a preference to fight (F) over dont fight (N), the payoffs in the F strategy need to be at least as large as the payoffs in the don't fight (N) strategy. Hence we must have:

$$\begin{aligned} 1-3p &\geq -1 \\ 2-4p &\geq -p \\ 1+p &\geq p-1 \\ 2 &\geq 0 \end{aligned}$$

Therefore from the first 2 inequalities we find a critical value of $p = p^* = 2/3$. So for $p < 2/3$ then BMW should choose fight (F).

(iii) For $p = p^*$ the extended strategic form payoff matrix becomes Table 7.6:

From this Table we can see that BMW's N strategy is weakly dominated by F. So given this for Mercedes (F,N) strongly dominates all others. This leaves the uniques Bayes-Nash equilibrium as ((F,N),F) with payoffs of 1/3 to Mercedes and -2/3 to BMW.

		BMW	
		F	N
	(F,F)	$0, -1$	$2, -1$
Mercedes	(F,N)	$1/3, -2/3$	$4/3, -2/3$
	(N,F)	$-4/3, 5/3$	$2/3, -1/3$
	(N,N)	$-1, 2$	$0, 0$

Table 7.6: Payoff matrix for the extended normal form game for $p = p^*$.

7.3.2 (2 Generals Problem –Harsanyi)

Two Generals are in charge of two opposing armies, A and B, who are about to face each other on the battlefield. Each General knows the strength and capability of their own army, but is unsure about the state of their opponents. Assume each army can be characterised as being in one of two states: `Strong` and `Weak`. Therefore army A can be in states A_S or A_W, army B in states B_S or B_W.

Additionally, each General has a choice of 2 different actions (such as attack, or defend) (which need not be the same for each army). Let us denote the actions for army A as a_1, a_2, and the actions for army B as b_1, b_2. Depending upon which strengths of armies meet and under what strategies, the payoffs for the different encounters are all shown in Figure 7.8 (note this is a zero sum game of conflict so the gain of army A is strictly at the expense of army B).

Each General has a pure strategy for each possibly conflict, and each General must form a strategy to maximise their *expected* payoff.

	b1	b2
a1	2	5
a2	-1	20

$(A_S B_S)$

	b1	b2
a1	-24	-36
a2	0	24

$(A_S B_W)$

	b1	b2
a1	28	15
a2	40	4

$(A_W B_S)$

	b1	b2
a1	12	20
a2	2	13

$(A_W B_W)$

Figure 7.8: Payoff matrices for the two different types of the two armies. It is a zero sum game and the payoffs shown correspond to the gain of army `A`. The 4 payoff matrices show the payoffs in the 4 possible cases of the different types of each army meeting, denoted $(A_S, B_S), (A_S, B_W), (A_W, B_S), (A_W, B_W)$. Individual Nash equilibria in each game are circled.

Assume that each general has formed a probabilistic belief about the strength of their opponent's army such that the joint probabilities of the different army combinations meeting on the battlefield are given by Table 7.7. (Note, we could have worked these out by multiplying together the individual beliefs that each General has about their opponent's army strengths).

(i) Using these beliefs, use Bayes theorem to evaluate $p(A_S|B_W)$ and $p(B_W|A_S)$, the conditional probabilites of the strong version of A's army given that B's army is weak, and the probability that B's army is weak given that A's army is strong.

(ii) Find any pure strategy Nash equilibria of the 2 General's game.

	B_S	B_W
A_S	4/10	1/10
A_W	2/10	3/10

Table 7.7: Table of joint beliefs of the generals for each possible combination of army types meeting $(A_S, B_S), (A_S, B_W), (A_W, B_S), (A_W, B_W)$

Answer:

(i) Using Bayes' theorem we know that

$$p(A_S|B_W) = \frac{p(A_S, B_W)}{p(B_W} = \frac{p(A_S, B_W)}{\sum_{states} p(states, B_W)}$$
$$= \frac{1/10}{1/10 + 3/10} = \frac{1}{4}$$

and similarly,

$$p(B_W|A_S) = \frac{p(A_S, B_W)}{p(A_S} = \frac{p(A_S, B_W)}{\sum_{states} p(states, A_W)}$$
$$= \frac{1/10}{1/10 + 4/10} = \frac{1}{5}$$

(ii) The two states of each of the armies can choose from actions a_1, a_2 and b_1, b_2 *each*, so there are 16 possible expected payoffs. For example, to calculate the expected payoff if army state A_S, A_W chooses actions (a_1, a_1) and army states B_S, B_W chooses actions (b_1, b_1), then the expected payoff for this set of strategies is calculated as

$$0.4 \times 2 + 0.1 \times (-24) + 0.2 \times 28 + 0.3 \times 12 = 7.6$$

Repeating this calculation across each type of each army choosing each strategy gives the expected payoffs in Table 7.8.

	$B_S \to b_1$ $B_W \to b_1$	$B_S \to b_1$ $B_W \to b_2$	$B_S \to b_2$ $B_W \to b_1$	$B_S \to b_2$ $B_W \to b_2$
$A_S \to a_1, A_W \to a_1$	7.6	8.8	6.2	7.4
$A_S \to a_1, A_W \to a_2$	7.0	9.1	1.0	3.1
$A_S \to a_2, A_W \to a_1$	8.8	13.6	14.6	19.4
$A_S \to a_2, A_W \to a_2$	8.2	13.9	9.4	15.1

Table 7.8: Table of expected payoffs for the different possibilities of different types of each army meeting in conflict. The Bayes-Nash equilibrium is circled.

From this Table we can see that the single pure strategy Nash equilibrium in this expanded game can be found by seeing that the maximin = minimax, which occurs for the strategy choice $(A_S \to$

$a_2, A_W \to a_1), (B_S \to b_1, B_W \to b_1)$ which is circled in the Table.

The interpretation of this Bayes-Nash equilibrium is that if A is strong, take action a_2; if A is weak, take action a_1 and B takes action b_1 irrespective of the strength of B.

Each player has a pure *state-dependent* strategy.

7.3.3 Extensive form Games– Incomplete Information

7.3.3.1 Law and Order

Consider the following game of political intrigue and scandal.

A politician (P) is suspected of diverting government funds for his own personal use. Such an activity is illegal, but in order to obtain a conviction, the Chief Constable of the local police force (C) must obtain proof. The Chief Constable can choose to hire a private investigator at a cost of 150 (in arbitrary utility units). The politician suspects that the Chief Constable may be attempting to obtain proof, so he can either continue with his misdemeanors or stop. If the private investigator is hired and the politician continues with the misdemeanors, then the Chief Constable is able to obtain proof. The cost of a trial is 100 for both parties. If the Chief Constable has proof of the illegal activity, then the politician is convicted, must pay a 1000 unit fine plus the costs (trial cost and private investigator cost) of the Chief Constable. Without proof of the illegal activity, then no outcome is reached in trial and no further transfers are made. The politician receives a benefit of 10 utilities from maintaining his illegal activities. The Chief Constable places a value $V > 0$ on being able to convict the politician of the wrongdoing. What value must V take so that the politician only continues the illegal activity with probability less than 0.25?

Answer:

The game has two players: the Chief Constable and the politician. The private investigator plays a passive role. The extensive form gametree, **without any information sets for the time being**, is shown in Figure 7.9

The payoffs are taken exactly from the description of the game. For instance, if the private investigator is hired, the politician continues illegal activity, and the Chief Constable goes to trial (it is assumed that going to trial with evidence gives a guaranteed conviction), then the payoffs for C is $V > 0$ and the payoffs for P is -1340 (10 benefit -1000 fine -100 own court costs -100 C court costs -150 for C private investigator costs). The remaining payoffs can be verified similarly. Given the description of the game, where do we need to include information sets?

First, the politician must decide whether or not to continue illegal activities and must make this decision without knowing whether or not the private investigator has been hired. Thus, one information set is required for the two nodes at which the politician must take action. If the private investigator has been hired, the Chief Constable will soon know with certainty whether or not the politician is continuing with the illegal activity. But, if the private investigator has not been hired, then the Chief Constable has no way of gathering this information. Thus, we need a second information set in the lower-right subgame to account for the fact that after deciding not to hire the private investigator, the Chief Constable must choose trial or not without knowing whether the politician is continuing with the illegal activity. This leads to the final incomplete information extensive form gametree in Figurecc2.

Now that we have translated the English of the problem into the extensive form used in game theory, we can find the subgame perfect Nash equilibrium. To begin the method of backward induction, we look at the lower-left subgame nodes for the Chief Constable. As given in Figure 7.11, the Chief Constable will choose "Trial" if the politician is continuing the illegal activityl and will choose "No Trial" if the politician has stopped, since there will be no evidence and the politician will not be found guilty.

We now move by backward induction to the next node up the tree that is not contained in an information set. The next node is the top node at which the Chief Constable must decide whether or not to hire the private investigator. We know that this extensive form game with information sets is equivalent

to a normal form game, but how many strategies do each of the two players have? For the politician, the answer is easy. He has two strategies: `Continue` the illegal activity, or `Stop`. Let us make the politician the column player.

How many strategies does the Chief Constable have? One possible strategy is to hire the private investigator. In this case, the actions for the politician at the bottom of the tree are already determined as in Figure 7.11. For the case in which the Chief Constable does not hire the private investigator, then he gains no new information prior to having to make the `Trial` or `No Trial` decision. Thus, the two remaining strategies for the Chief Constable are combined strategies: (a) `Dont Hire PI` and `Trial` and (b) `Dont Hire PI` and `No Trial`. The Chief Constable will be the row player and has the three strategies discussed above.

The upper subgame is then equivalent to the 3×2 normal form game shown in Figure 7.12 where we have been careful in transcribing the payoffs from the extensive form to the normal form to ensure that the payoffs correspond to the appropriate outcomes.

We can see that the normal form game does not have any pure strategy Nash equilibria (recall that $V > 0$ by assumption). In this case we must look for mixed strategy Nash equilibria. Denote (p_1, p_2, p_3) to be the probabilities associated with each of the three strategies (in order) of the Chief Constable, where $p_1 + p_2 + p_3 = 1$. Denote $(q, 1-q)$ as the probabilities associated with each of the two strategies (in order) of the politician. First, the values for (p_1, p_2, p_3) are determined to make the politician indifferent between his two strategies:

$$\begin{aligned} -1340p_1 - 90p_2 - 10p_3 &= -100p_2 \\ 10p_2 + 10p_3 &= 1340p_1 \end{aligned}$$

This implies that $p_1 > 0$ and either $(p_2 = 134p_1, p_3 = 0)$ or $(p_2 = 0, p_3 = 134p_1)$. However, note that for the Chief Constable, strategy `(Dont Hire, Trial)` is strictly dominated by `(Dont Hire, No Trial)` and so the middle strategy would never be played, and hence we know $p_2 = 0$. Therefore if $p_2 = 0$ then $p_3 = 134p_1$ and since they add up to 1 we have

$(p_1, p_2, p_3) = (1/135, 0, 134/135)$; the CHief COnstable only hires a private investigator with very small probability.

To find the value V needs to be so that the politician continues the illegal activity with probability less than 0.25 we need to find the lower bound V so that $q = 0.25$. The probabilities $(q; 1q)$ are determined to make the Chief Constable indifferent between `Hire PI` and `(Dont Hire PI, No Trial)`:

$$\begin{aligned} qV + (1-q)(-150) &= 0 \\ ().25)V + (0.75)(-150) &= 0 \\ \Rightarrow V &= 450 \end{aligned}$$

This implies that for all $V \geq 450$ the Chief Constable can hire a private investigator with probability $p_1 = 1/135$ and keep the probability of crime capped at 0.25.

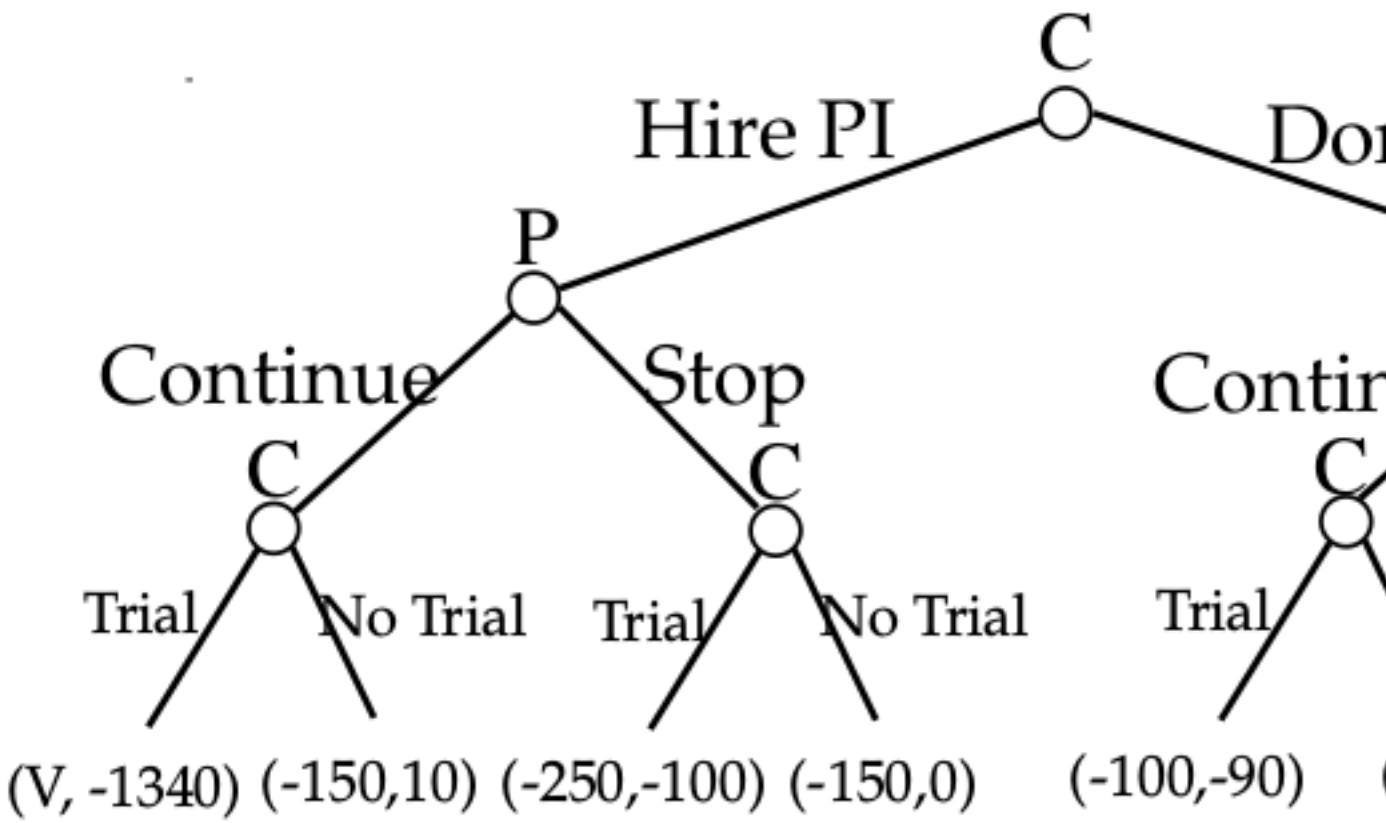

Figure 7.9: Extensive form gametree of the Law and Order game, WITHOUT the information sets for the moment.

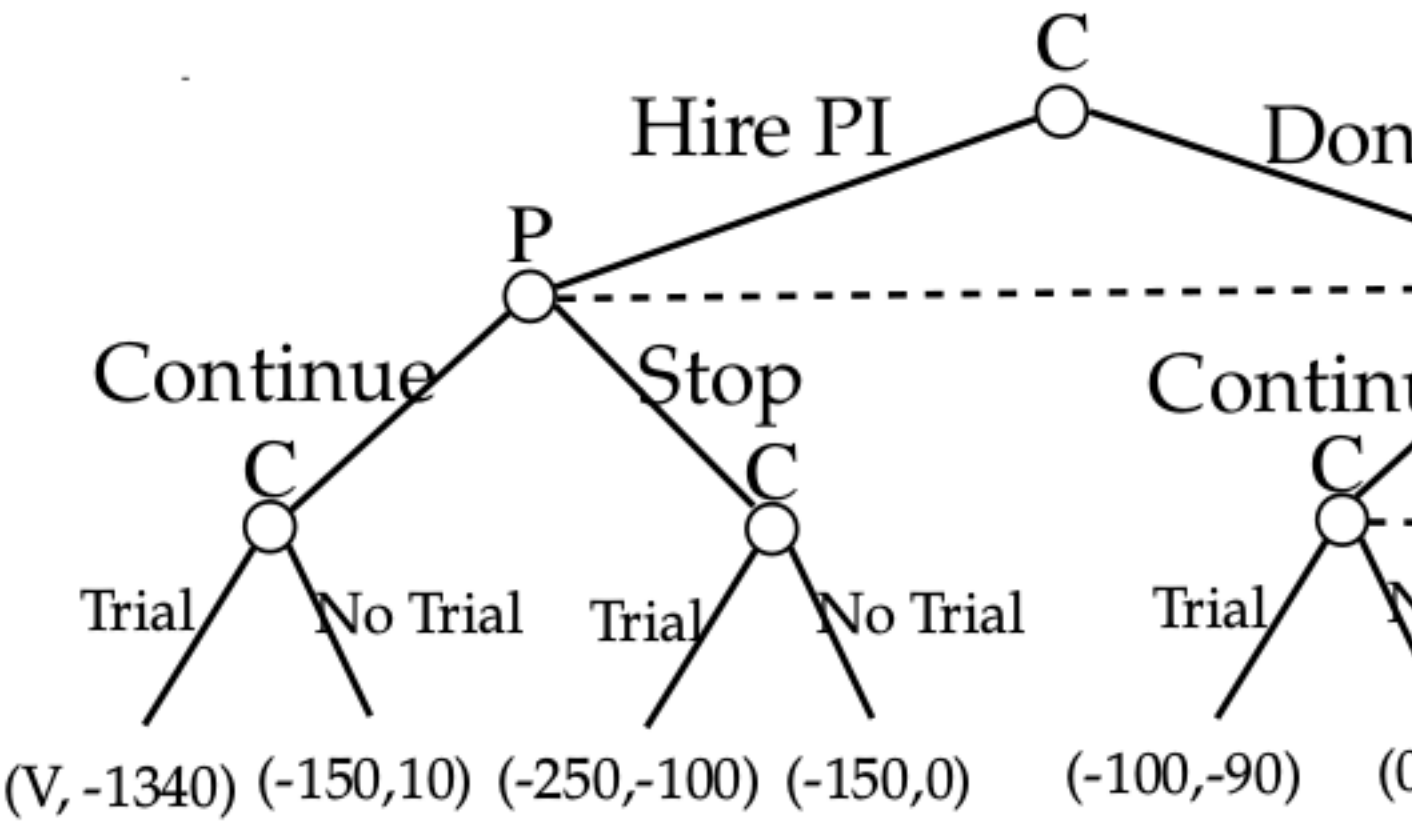

Figure 7.10: Final incomplete extensive form gametree of the Law and Order game, WITH the information sets added.

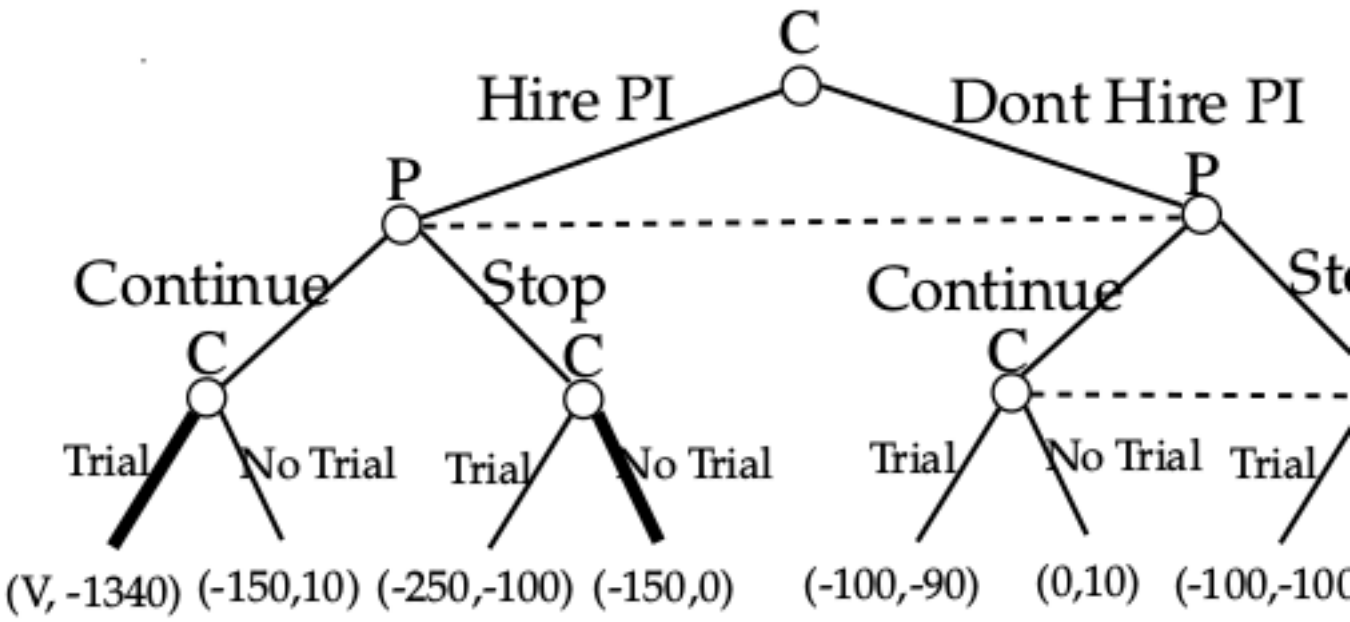

Figure 7.11: First subgames considered with the Nash equilibria shown by the THICK lines.

		Politician	
		`Continue`	`Stop`
	`Hire PI`	$V, -1340$	$-150, 0$
Chief Constable	`(Dont Hire, Trial)`	$-100, -90$	$-100, -100$
	`(Dont Hire, No Trial)`	$0, 10$	$0, 0$

Figure 7.12: Equivalent normal form game for the upper subgame in Figure 7.11

Chapter 8

Auctions

The design and analysis of auctions is one of the triumphs of game theory. Modern game theory approaches to Auction theory were pioneered by the economist William Vickrey in 1961. Its practical use became apparent in the 1990s, when auctions of the radio frequency spectrum for mobile telecommunication raised billions of dollars. Economic theorists advised governments on the design of these auctions, and companies on how to bid (see *McMillan, "Selling spectrum rights," Journal of Economic Perspectives Vol. 8, 1994, pages 145–162*). The auctions for spectrum rights are complex. However, many principles for sound bidding can be illustrated by applying game–theoretic ideas to simple examples. For further insight see *Milgrom, Auctions and bidding: a primer" (Journal of Economic Perspectives Vol. 3, 1989, pages 3–22)* for a broader view of the theory of bidding in auctions.

Auctions are examples of the application of Bayesian games. This generally corresponds to a situation of incomplete information because the *valuations* of different potential buyers are unknown. For example, if we were to announce a mechanism, which involves giving a particular good (for example a mathlete hoodie) for free to the individual with the highest valuation, this would create an incentive for all individuals to overstate their valuations. In general, auctions are designed by profit–maximising entities, which would like to sell the goods to raise the highest possible revenue.

8.1 Example: Second Price Auction (with Complete Information)

Before we consider Bayesian auctions (ie with incomplete information), it is useful to consider a slightly artificial version of an auction with *complete* information.

The *second price* auction game is specified as follows:

- There exists an object that is to be assigned to one of n players in $\{1, .., n\}$.
- Each player has her own valuation of the object. Player i's valuation of the object is denoted v_i . We further assume that $v_1 > v_2 > ... > 0$.
- Note that for now, we assume that everybody knows all the valuations $v_1, ..., v_n$, i.e., this is a *complete information* game. We will analyse the incomplete information version of this game later. The assignment process is described as follows:

- The players simultaneously submit bids, $b_1, \ldots, b_n$.
- The object is given to the player with the highest bid (or to a random player among the ones bidding the highest value).
- The winner pays the second highest bid.
- The utility function for each of the players is as follows:
  the winner receives her valuation of the object minus the price she pays, i.e., $v_i - b_i$; and everyone else receives 0.

In the second price auction, truthful bidding, i.e., $b_i = v_i$ for all i, is a Nash equilibrium.

Proof We want to show that the strategy profile $(b_1, .., b_n) = (v_1, \ldots, v_n)$ is a Nash Equilibrium – a truthful equilibrium. First note that if indeed everyone plays according to that strategy, then player 1 receives the object and pays a price v_2. This means that her payoff will be $v_1 - v_2 > 0$, and all other payoffs will be 0. Now, player 1 has no incentive to deviate, since her utility can only decrease. Likewise, for all other players $v_i = v_1$, it is the case that in order for v_i to change her payoff from 0 she needs to bid more than v_1 , in which case her payoff will be $v_i - v_1 < 0$. Thus there is no incentive to deviate from their truthful strategies for any player.

Are there other Nash Equilibria? In fact, there are also unlikely Nash equilibria in second price auctions. We show that the strategy $(v_1, 0, 0, \ldots, 0)$ is also a Nash Equilibrium.

As before, player 1 will receive the object, and will have a payoff of $v_1 - 0 = v_1$. Using the same argument as before we conclude that none of the players have an incentive to deviate, and the strategy is thus a Nash Equilibrium. It can be verified the strategy $(v_2, v_1, 0, 0, \ldots, 0)$ is also a Nash Equilibrium. Why? (Left as an exercise for the student).

Nevertheless, the truthful equilibrium (where $b_i = v_i$) is a Weakly Dominant Nash Equilibrium. In particular, truthful bidding, $b_i = v_i$, weakly dominates all other strategies.

8.2 Different types of auctions and terminology

- **Take-it or Leave-it auctions:** Retail outlets use this type of auction. A price is written on the object and prospective buyers can take it or leave it.
- **English auctions:** Sotherbys use this type of auction to sell Old Masters. An auctioneer invites oral bids for an object. The bidding continues upwards in bid value until no-one is prepared to bid any more. The object is sold to the buyer who made the last bid. It is essentially an ascending sequential bids process.
- **First price sealed bid auctions:** This is the standard format for Government tenders. Each bidder privately writes their bid on a piece of paper and seals it in an envelope. It is similar to English auctions, but in the style of a strategic form game; all players make a single simultaneous bid and the highest one obtains the object and pays their expressed bid.
- **Second price sealed bid auctions:** similar to first price auctions, except that the winner pays the second highest bid.

- **Vickrey Auction:** In a Vickrey auction the object is sold to the highest bidder but at the highest bid offered by a loser. This will typically be the second highest price unless there is a tie for first place. We always assume the bids are submitted using the sealed-bid mechanism and so a Vickrey auction is also identified as a second-price sealed-bid auction, but it need not always be so.
- **All-pay auctions:** The bidding is as in an English auction with the highest bidder gaining the object, but everyone also pays their highest bid. (Bribing corrupt officials is rather like an all-pay auction in that only one bribe can be successful, though several bribes may have been accepted).
- **Dutch auctions:** descending sequential auctions; The auctioneer announces an initially high price for the object and the price is gradually reduced until a buyer calls a halt and accepts the current price. The first buyer to do so acquires the object at the current price when he or she intervened. The auction stops when an individual announces that she wants to buy at that price. Otherwise the price is reduced sequentially until somebody stops the auction.
- **Combinatorial auctions:** when more than one item is auctioned, and agents value combinations of items.
- **Private value auctions:** valuation of each agent is independent of others' valuations;
- **Common value auctions:** the object has a potentially common value, and each individual's signal is imperfectly correlated with this common value.

These different auctions have different game theory solutions. For example, unlike the first complete-information example in which it was a Nash equilibrium to bid the true valuation, in a Dutch auction, it would be suboptimal to bid your true valuation for an object.

This is because a Dutch auction is a sequential descending auction where the first buyer to stop the auction pays the current price when he intervenes. So this is the same outcome as if all the players put their bids in sealed envelopes and the auctioneer chooses the largest bid in the envelopes. ie it is the same as a first-price, sealed-bid auction. So bidding your true valuation can only ever gain you zero utility, which is a worse expected utility than a strategy where you bid less than your valuation. In the latter case there is a possibility you might win the auction with a positive utility giving an overall positive expected utility.

8.3 Modelling Auctions

A model of an auction is a valuation structure for the bidders (i.e., private values for the case of private-value auctions), which is typically a probability distribution over the valuations available to the bidders.

Let us focus on first and second price sealed bid auctions, where bids are submitted simultaneously. Each of these two auction formats defines a static game of incomplete information (Bayesian game) among the bidders. We will determine Bayesian Nash equilibria (defined later) in these games and compare the equilibrium bidding behaviour.

More explicitly, suppose that there is a single object for sale and N potential buyers are bidding for the object. Bidder i assigns a value v_i to the object, i.e., a utility $v_i - b_i$, when he pays b_i for the object. He knows v_i. This implies that we have a private value auction (v_i is his "private information" and "private

value"). Suppose also that each v_i is independently and identically distributed on the interval $[0, \bar{v}]$ with cumulative distribution function F , with continuous density f and full support on $[0, \bar{v}]$.

Bidder i knows the realization of its value v_i (or realisation v_i of the random variable V_i , though we will not use this latter notation) and that other bidders' values are independently distributed according to F , i.e., all components of the model except the realised values are "common knowledge".

Bidders are risk neutral, i.e., they are interested in maximising their expected profits. This model defines a Bayesian game of incomplete information, where the types of the players (bidders) are their valuations, and a pure strategy for a bidder is a map

$$\beta_i : [0, \bar{v}] \to \Re^+.$$

We will characterise the symmetric equilibrium strategies in the first and second price auctions.

Once we characterise these equilibria, then we can also investigate which auction format yields a higher expected revenue to the seller at the symmetric equilibrium.

8.4 Second Price Auctions

In the auction, each bidder submits a sealed bid of b_i , and given the vector of bids $\boldsymbol{b} = (b_i, \boldsymbol{b}_{-i})$ and valuation v_i of player i, the payoff to i is

$$U_i((b_i, \boldsymbol{b}_{-i}), v_i) = \begin{cases} v_i - max_{j \neq i} b_j & \text{if } b_i > max_{j \neq i} b_j \\ 0 & \text{if } b_i < max_{j \neq i} b_j \end{cases}$$

Let us also assume that if there is a tie, i.e., $b_i = max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability.

In a second-price auction, it is a weakly dominant strategy to bid truthfully, i.e., according to $\beta(v) = v$.

Moreover, now there are no other optimal strategies and thus the (Bayesian) equilibrium will be unique, since the valuation of other players are not known. Therefore, we have established:

In the second price auction, there exists a unique Bayesian Nash equilibrium which involves

$$\beta(v) = v.$$

It is straightforward to generalize the exact same result to the situation in which bidders' values are not independently and identically distributed. As in the complete information case, bidding truthfully remains a weakly dominant strategy. The assumption of private values is important (i.e., the valuations are known at the time of the bidding).

8.4.1 Expected payment for second price auction

Let us next determine the expected payment by a bidder with value v , and for this, let us focus on the case in which valuations are independent and identically distributed

Fix bidder 1 and define the random variable y_1 as the highest value among the remaining $N - 1$ bidders, i.e., $y_1 = max\{v_2, \ldots, v_N\}$.

Let G denote the cumulative distribution function of y_1 . Clearly,

$$G(y) = F(v)^{N-1} \text{for any} v \in [0, \bar{v}]$$

In a second price auction, the expected payment by a bidder with value v is given by

$$\begin{aligned} m_{II}(v) &= Prob(v \ wins) \times E[\text{second highest bid}|v \text{ is the highest bid}] \\ &= Prob(y_1 \leq v) \times E[y_1|y_1 \leq v] \\ &= G(v) \times E[y_1|y_1 \leq v]. \end{aligned}$$

This is a general case argument. What is it if we assume a specific uniform distribution of values of the bidders?

Example: Uniform Distribution

Suppose that there are two bidders with valuations, v_1 and v_2 , distributed uniformly over $[0, 1]$. Then $G(v_1) = v_1$, and

$$E[y_1|y_1 \leq v_1] = E[v_2|v_2 \leq v_1] = \frac{v_1}{2}$$

Thus

$$m_{II}(v) = \frac{v^2}{2}$$

If, instead, there are $N - 1$ bidders with valuations independently distributed over $[0, 1]$,

$$G(v_1) = v_1^{N-1}$$

and

$$E[y_1|y_1 \leq v_1] = \frac{N-1}{N} v_1$$

and so

$$m_{II}(v) = \frac{N-1}{N} v^N$$

We will be comparing this value with the equivalent one from a first-price bid auction later.

8.5 First Price Auctions

In a first price auction, each bidder submits a sealed bid of b_i , and given these bids, the payoffs are given by

$$U_i((b_i, \boldsymbol{b}_{-i}), v_i) = \begin{cases} v_i - b_i & \text{if } b_i > max_{j \neq i} b_j \\ 0 & \text{if } b_i < max_{j \neq i} b_j \end{cases}$$

Tie-breaking is similar to before.

In a first price auction, the equilibrium behaviour is more complicated than in a second-price auction. Clearly, bidding truthfully is not optimal (why not?). So we have to work out more complicated strategies.

What about the equilibrium of a first-price auction? Clearly an agent should not bid more than their true value because it would yield negative utility whenever they win the auction. Moreover, it is also clear

that an agent should bid strictly less than their true value. This is because if they bid truthfully and win, then the utility is exactly zero. For a formal analysis of the properties of the first-price auction we need a new equilibrium concept: the Bayes–Nash equilibrium.

(Bayes–Nash equilibrium).
The strategy profile $\boldsymbol{s}^* = (s_1^*, \ldots s_n^*)$ is a Bayes-Nash equilibrium (BNE) if, for all i,

$$E[U_i(s_i^*(v_i), \boldsymbol{s}_{-i}^*(v_{-i}))] \geq E[U_i(b_i(v_i), \boldsymbol{s}_{-i}^*(v_{-i}))] \quad \forall v_i, b_i.$$

This requires that every agent i maximises its expected utility by adopting s_i^* for every value v_i and given the strategies adopted by the other agents. The expectation is taken with respect to the distribution on values of the other agents. For a BNE to make sense, both the distribution must be common knowledge and the assumption of agent rationality must be common knowledge.

Generally speaking, it can be difficult to characterize the BNE of an auction, because we need to identify equilibrium strategies that are *functions* from values to bids. To keep things simple, we focus on values which are uniformly distributed on the interval $[0, 1]$ typically. We also look for a symmetric (continuous and differentiable) equilibrium, in which every agent plays the same strategy.

We first consider the general case, then specialise to the case of a uniform distributions, showing how the bid in a first price auction is not truthful bidding.

Suppose that bidders $j \neq 1$ follow the symmetric increasing and differentiable equilibrium strategy $\beta_I = \beta$, where

$$\beta_i : [0, \bar{v}] \to \Re^+.$$

We also assume, without loss of any generality, that β is increasing. We will then allow player 1 to use strategy β_1 and then characterise β such that when all other players play β, it is a best response for player 1. Since player 1 was arbitrary, this will complete the characterisation of equilibrium.

Suppose that bidder 1's value is v_1 and he bids the amount b (i.e., $\beta(v_1) = b$).

First, note that a bidder with value 0 would never submit a positive bid, so $\beta(0) = 0$. Next, note that bidder 1 wins the auction whenever

$$max_{i \neq 1} \beta(v_i) < b.$$

Since $\beta(.)$ is increasing, we have

$$max_{i \neq 1} \beta(v_i) = \beta(max_{i \neq 1} v_i) = \beta(y_1),$$

where recall that $y_1 = max\{v_2, \ldots, v_N\}$. This implies that bidder 1 wins whenever $y_1 < \beta^{-1}(b)$.

Consequently, we can find an optimal bid of bidder 1, with valuation $v_1 = v$, as the solution to the maximization problem:

$$max_{b \geq 0} G(\beta^{-1}(b))(v - b).$$

The first-order (necessary) conditions imply

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(v - b) - G(\beta^{-1}(b)) = 0, \tag{8.1}$$

where $g = G'$ is the probability density function of the random variable y_1 . [Recall that the derivative of

$\beta^{-1}(b)$ is $1/\beta'(\beta^{-1}(b))$.]

This is a first-order differential equation, which we can in general solve.

More explicitly, for a symmetric equilibrium, we have $\beta(v) = b$, and therefore (8.1) yields:

$$G(v)\beta'(v) + g(v)\beta(v) = vg(v).$$

Equivalently, the first-order differential equation is

$$\frac{d}{dv}[G(v)\beta(v)] = vg(v),$$

with boundary condition $\beta(0) = 0$.

We can rewrite this as the following optimal bidding strategy

$$\beta(v) = \frac{1}{G(v)} \int_0^v yg(y)dy = E[y_1|y_1 < v].$$

Note, however, that we skipped one additional step in the argument: the first-order conditions are only necessary, so one needs to show sufficiency to complete the proof that the strategy $\beta(v) = E[y_1|y_1 < v]$ is optimal.

This detail notwithstanding, we have:

In the first price auction, there exists a unique symmetric equilibrium given by

$$\beta_I(v) = E[y_1|y_1 < v].$$

8.5.1 First Price Auctions: Payments and Revenues

For our general case, the expected payment of a bidder with value v in a first price auction is given by

$$m_I(v) = Prob(vwins) \times \beta(v) = G(v) \times E[y_1|y_1 < v]. \tag{8.2}$$

This can be directly compared to (8.1), which was the payment in the second price auction ($m_{II}(v) = G(v) \times E[y_1|y_1 \leq v]$). This establishes the somewhat surprising results that $m_I(v) = m_{II}(v)$, i.e., both auction formats yield the same expected revenue to the seller.

8.5.2 Specific Example: Uniform distribution

Theorem 11. *A symmetric Bayes–Nash equilibrium in a first-price sealed-bid auction for agents with independently and identically distributed values uniform on* $[0, 1]$ *is*

$$s^*(v_i) = \left(\frac{n-1}{n}\right) v_i \quad \forall \text{ agents } i \in \{1, 2, \ldots, n\}$$

Proof: We give a proof that this is a BNE for the case of $n = 2$ (ie that the optimal bid is actually $b_i = v_i/2$) and later explain how this generalises to $n > 2$. A common technique is "guess-and-verify". That is, we assume that the equilibrium strategies have a particular functional form and look to see if an

equilibrium with the functional form does indeed exist. Let us look for an equilibrium with simple linear structure

$$b_i = \alpha v_i \ \forall i \in \{1,2\} \quad 0 < \alpha < 1$$

Focus on agent 1. Fix v_1 as the realisation of agent 1's value. Let $B_2 = \alpha V_2$ denote the bid of agent 2, which is a random variable from the perspective of agent 1. For a given bid b, agent 1's expected utility is

$$(v_1 - b) \times Prob[B_2 \leq b] \equiv (v_1 - b) \times \frac{b}{\alpha} \tag{8.3}$$

where $(v_1 - b)$ is the utility conditioned on winning and $Prob[B_2 \leq b]$ is the probability of winning given bid b. For a uniform distribution on $[0,1]$ and distribution function $F(v) = v$, then

$$Prob[B_2 \leq b] = Prob[V_2 \leq b/\alpha] = b/\alpha$$

For the relevant first-order condition to find agent 1's optimal bid (and thus best-response) given value v_1, we take the derivative with respect to b, and set this to zero:

$$\begin{aligned} \frac{\partial}{\partial b}\left[(v_1 - b) \times \frac{b}{\alpha}\right] &= 0 \\ \Rightarrow \frac{v_1 - 2b}{\alpha} &= 0 \\ b &= \frac{v_1}{2} \end{aligned}$$

This holds for all $\alpha > 0$ and all v_1.

By completing the same calculations from the perspective of agent 2 rather than agent 1, we similarly get $s^*(v_2) = v_2/2 = \alpha v_2$.

For a first-price sealed-bid auction with two bidders, we have therefore verified that bidding $v_i/2$ is a BNE.

The analysis can be generalized to an auction with n bidders, by recognizing that the general form of the expected utility above (8.3) is

$$(v_1 - b)\left(\frac{b}{\alpha}\right)^{n-1}$$

since

$$Prob[max(B_2, B_3, \ldots B_n) \leq b] \equiv \left(\frac{b}{\alpha}\right)^{n-1}$$

when values are independently and identically distributed uniformly on $[0,1]$ and every agent adopts $b_i = \alpha v_i$.

8.6 Revenue Equivalence

In our analysis of auctions, we have generally considered the game from the point of view of the bidders (the buyers). In any auction, there must exist a seller on the other side of the market with payoffs. An interesting property of auctions is that the expected revenue received by a seller in a first price auction is the same as the expected revenue in a second price auction. We will demonstrate this property using the previous results in a two bidder auction in which both bidders have values drawn from the

distribution $U[0,1]$. The equivalence result holds more generally. In a *first price* auction, both bidders choose the bidding strategies $b_i = v_i/2$ for $i = 1, 2$ meaning that the price that the seller receives is equal to $\max[v_1/2, v_2/2]$. So what is the *expected* gain of the seller. For this we need to work out the average using the distribution of bids.

The values v_1 and v_2 are drawn from the uniform distribution from 0 to 1. The probability density function is equal to $f(v) = 1 \forall\ v \in [0,1]$ and the *cumulative* density function is

$$F(v) = \int_0^v f(V)\, dV \equiv v$$

Now consider the distribution of $w = max[v_1, v_2]$. The cumulative distribution function of this max function is $G(w) = F(w)^2$ [1]. Thus, the probability density function is equal to $g(w) = 2w$ (take a derivative). The formula to compute the expected payoff is given by the average $\int wg(w)\, dw$ and so the expected payoff to the seller will be

$$\begin{aligned}\frac{1}{2}\int_0^1 w \times (2w)\, dw &= \frac{1}{2}\left[\frac{2}{3}w^3\right]_0^1 \\ &= \frac{1}{2}\left[\frac{2}{3}\right] \\ &= \frac{1}{3}\end{aligned}$$

What about the revenue expected from the second price auction instead?

In the second price auction, both bidders choose the bidding strategies $b_i = v_i$ for $i = 1, 2$ and the price that the seller receives is equal to $\min[v_1; v_2]$ where v_1 and v_2 are drawn from the uniform distribution from 0 to 1. So to work out the *expected* return to the seller we need to average over the distribution corresponding to $w = \min[v_1, v_2]$.

The cumulative density function of this is $H(w) = 1 - (1 - F(w))^2$ [2].

Thus, the probability density function is equal to $h(w) = -2w + 2$ (take a derivative). The formula to compute the expected payoff is given by the average $\int wh(w)\, dw$ and so the expected payoff to the seller will be :

$$\begin{aligned}\int_0^1 w \times (-2w + 2)\, dw &= \left[-\frac{2}{3}w^3 + w^2\right]_0^1 \\ &= \left[-\frac{2}{3} + 1\right] \\ &= \frac{1}{3}\end{aligned}$$

which is the same as the expected revenue to the seller if the buyers had used a first price auction instead. ie even though the buyers paid different amounts in the first and second price auctions, the expected payoff to the seller in both auction designs is the same! This is the concept known as the Revenue Equivalence Theorem.

[1] Because the probability that $y \leq \max[x_1, x_2]$ is equal to the product of ($Prob[y \leq x_1] \times Prob[y \leq x_2]$.

[2] Because the probability that $y < \min[x_1, x_2] = 1 - Prob[v > max[v_1, v_2]]$, which is the same as $1 - Prob[y > v]^2$ since v_1 and v_2 are independent, and then use $Prob[y > v] = 1 - Prob[y < v] = v$.

We can generalise this specific example of 2 bidders to N-bidders as follows.

Example: N-player First Price Auctions: Uniform Distribution

If we assume that values are uniformly distributed over $[0,1]$ then, we can verify that for N players the optimum bidding function is

$$b(v) = \frac{N-1}{N}v$$

Moreover, since for N independent bidders the cumulative distribution function would be $G(v_1) = v_1^{N-1}$, and so we again have the expected revenue to the seller for N bidders in a first price auction will be the same as in the second price auction, ie

$$m_I(v) = \frac{N-1}{N}v^N.$$

8.6.1 General Theorem

The previous results are specific examples of a more general theorem. Consider any standard auction, in which buyers submit bids and the object is given to the bidder with the highest bid. Suppose that values are independent and identically distributed and that all bidders are risk neutral. Then, we have the following theorem:

Theorem 12. *Revenue Equivalence Any symmetric and increasing equilibria of any standard auction (such that the expected payment of a bidder with value* 0 *is* 0*) yields the same expected revenue to the seller.*

8.7 Scaling Distributions

What if the Uniform distribution describing the private information was not on $[0,1]$, but on $[A,B]$? How would it affect matters?

So, now suppose that both bidders' values are drawn from the uniform distribution between A and B, $U[A,B]$ where $B > A \geq 0$. Refer to the game corresponding to this $U[A,B]$ distribution as the "true" game and the game corresponding to $U[0,1]$ as the "updated" game. Suppose that player 1 has a value for the object in the true game equal to v_1. To get from the true game to the updated game, we have to subtract A and divide by $B-A$. Thus, the value for player 1 in the updated game is equal to

$$v_1 = (v_1 - A)/(B-A).$$

In the updated game, we know that (from the previous subsection) the optimal bid should be

$$\hat{b}_1 = \frac{1}{2}\hat{v}_1 = \frac{1}{2}\left(\frac{v_1 - A}{B-A}\right)$$

To get from the updated game back to the true game, we must multiply by $B-A$ and add A (the

reverse of the operations performed before). Thus, $b_1 = \hat{b}_1 \times (B - A) + A$,

$$\begin{aligned} b_1 &= \frac{1}{2}\left(\frac{v_1 - A}{B - A}\right) \times (B - A) + A \\ &= \frac{1}{2}v_1 + \frac{1}{2}A \end{aligned}$$

This is simply equal to the average between the lower bound A and the true value v_1.

8.8 Exercises for Chapter 8

Auctions

8.8.1 Exercise: First Price Auction with Reserve Price

Suppose that in a standard First Price Auction two bidders' valuations are private information and are randomly drawn from a uniform distribution on $[0, 1]$. However, in this version of the game, a reservation price, RP, is in place. This simply means that the auction winner must have the highest bid and that the winning bid must be greater than or equal to RP. What is the optimal bidding strategy of player 1?

Answer:

Le us consider the payoff maximisation problem for bidder 1. Bidder 2 is assumed to bid according to the linear strategy: $b_2 = a \times v_2 + c$; where v_2 is the private information value for bidder 2. The possible payoffs for bidder 1 are as follows:

$$\begin{aligned} v_1 - b_1 \text{ if } b_1 &> \max\{av_2 + c, RP\} \\ (v_1 - b_1)/2 \text{ if } b_1 &= av_2 + c \geq RP \\ 0 \text{ if } b_1 &< av_2 + c, \text{ or } b_1 < RP \end{aligned}$$

Thus, if $v_1 < RP$ then it does not matter what your bid is, because you never want to place a bid above the reservation price RP. So, the bidding strategies are only valid for $v_1 > RP$. The strategies for both bidders are then given by the linear equation:

$$b_i = av_i + c, \text{ for } v_i > RP \text{ where } i = 1, 2$$

Consider the optimal bids for player 1 under the scenario that $v_1 \geq RP$. There are two cases to consider:

1. With probability $Pr(v_2 < RP)$ player 2 will not be able to place a bid greater than or equal to RP. Thus, player 1 can win the auction with a bid equal to $b_1 = RP$. She cannot win the auction with a bid lower than RP due to the definition of a reservation price. Thus, this is the optimal bid in this scenario for player 1.

2. With probability $Pr(v_2 \geq RP)$: Players 1 and 2 will be competing in a game in which valuations are drawn from $U[RP, 1]$. As seen from the earlier subsection on scaling distributions, after scaling the $U[RP, 1]$ game down to $U[0, 1]$ the best strategy for player 1 is given by

$$b_1 = \frac{1}{2}RP + \frac{1}{2}v_1 :$$

Players in this game maximize *expected* payoffs. So, if the optimal strategy is $b_1 = RP$ for the first possibility (with probability $Pr(v_2 < RP) = RP$) and $b_1 = RP/2 + v_1/2$ for the second possibility (with probability $Pr(v_2 \geq RP) \equiv 1 - RP$), then the optimal strategy over all possibilities is given by the

expectation:

$$\begin{aligned} b_1 &= Pr(v_2 < RP) \times (RP) + Pr(v_2 \geq RP) \times (RP/2 + v_1/2) \\ &= (RP) \times (RP) + (1 - RP) \times (RP/2 + v_1/2) \end{aligned}$$

After some algebra, the optimal strategy can be expressed as:

$$\begin{aligned} b1 &= (RP)^2 + \frac{1}{2}RP + \frac{1}{2}(1 - RP)v_1 - \frac{1}{2}(RP)^2 \\ &= \frac{1}{2}(1 - RP)v_1 + \frac{1}{2}RP(1 + RP) \end{aligned}$$

whenever $v_1 > RP$. Note that this continues to be in linear form , $b_1 = av_1 + c$ where $a = \frac{1}{2}(1 - RP)$ and $c = \frac{1}{2}RP(1 + RP)$.

8.8.2 Exercise: (Nash Equilibrium in a first price sealed bid auction)

In a first price sealed bid auction, there exists a Nash equilibrium in which each player bids less than her valuation. In this example we consider 2 players/bidders whose valuations and bids come from a uniform distribution on $[0, 1]$.

Player 1 has valuation v_1 and bids b_1. Player 2 has valuation v_2 and bids $\beta_2(v_2)$.

Suppose player 2 bids $\beta_2(v_2) = v_2/2$.

(i) What is the probability that Player 1 wins if her bid is $b_1 > 1/2$?

(ii) What is the probability that Player 1 wins if her bid is $b_1 < 1/2$?

(iii) Write down the pay-off of player 1 for $b > 1/2$ and $b < 1/2$, and sketch the graph.

(iv) For what value of b_1 is player 1's payoff maximised?

Answer:

(i) If player 1 bids more than 1/2 then she wins with probability 1, since player 2's bid is distributed uniformly on $[0, 1/2]$.

(ii) If player 1 bids less than 1/2 then it depends on the bid of player 2. The probability she wins is the probability that player 2's valuation is less than $2b_1$. Since it is uniformly distributed, this probability is $2b_1$.

(iii) Therefore the payoff for player 1 is the probability she wins times her gain,$(v_1 - b_1)$, ie

$$U(b, v) = \begin{cases} 2b_1(v_1 - b_1) & if\ 0 \leq b_1 \leq \frac{1}{2} \\ (v_1 - b_1) & if\ b_1 \geq \frac{1}{2} \end{cases}$$

(iv) From the graph, or taking the derivative of the payoff and setting to zero gives $b_1 = v_1/2$.

8.8.3 Exercise: (All-pay auction)

Two bidders are bidding for the same object in an auction. Their valuations are private information: each bidder knows his own valuation, but only knows that his opponents valuation is drawn from the uniform distribution on $[0, 1]$.

Each bidder has to submit a sealed bid. The auctioneer collects both their bids and the object is allocated to the highest bidder. That is, the auction is an all-pay auction.

(i) Suppose that bidder i believes that his opponent, j's, bid is given by

$$\beta_j(v_j) = k_j \times \sqrt{v}_j$$

where v_j is his opponent's valuation and k_j is a constant.

What is i's expected payoff from bidding b_i when his valuation is v_i ?

(ii) What is player i's optimal bid as a function of v_i?

(iii) Why does this indicate player i's initial belief of player j's bidding strategy is likely to be wrong?

Answer:

(i) If bidder i wins they make a profit of $v_i - b_i$, and if they lose they make a 'profit' of $-b_i$ since they lose their stake as it is a pay-all auction. They win if their bid is higher than bidder j's bid, so the expected payoff depends in $(v_i - b_i) \times Prob(b_i > b_j) - b_i \times [1 - Prob(b_i > b_j)]$, or

$$U_i(b_i, v_i) = E[v_i \times Prob(b_i > b_j) - b_i]$$

Since bidder i assumes bidder j's bid is $b_j = k_j\sqrt{v_j}$ then

$$\begin{aligned} & Prob(b_i > b_j) \\ & \equiv Prob(b_i > k_j\sqrt{v_j}) \\ & \equiv Prob(\left(\frac{b_i}{k_j}\right)^2 > v_j) \end{aligned}$$

But since v_j is uniformly distributed on $[0, 1]$ then

$$Prob(\left(\frac{b_i}{k_j}\right)^2 > v_j) \equiv \left(\frac{b_i}{k_j}\right)^2$$

So,

$$U_i(b_i, v_j) = v_i \left(\frac{b_i}{k_j}\right)^2 - b_i$$

(ii) Maximising bidder i's expected payoff is found by taking the derivative of the payoff found in part (i) with respect to the bid b_i, and setting to zero. So we need to solve:

$$\frac{dU_i}{db_i} = \frac{d}{db_i}\left[v_i \left(\frac{b_i}{k_j}\right)^2 - b_i\right] = 0$$

or

$$2\frac{v_i b_i}{k_j^2} - 1 = 0$$

Rearranging gives:

$$b_i = \frac{1}{2}\frac{k_j^2}{v_i}.$$

(iii) This is probably wrong since we would reasonably expect a person to bid monotonically with their valuation. This expression in part (ii) yields a bidding strategy inversely proportional to their valuation. So the initial assumption that player j's bid would be proportional to their square root valuation is probably incorrect.

8.8.4 Exercise: (Exam Question 2013)

a) Explain why, in a Dutch auction, it would be suboptimal to bid your true valuation for an object.

b) A seller has a single item for auction. Two buyers each draw a "type" independently with equal probability from the set $\{£8, £12\}$. Consider a sealed-bid second-price auction in which the only possible bids are $£8$ and $£12$. Let P denote the second price in the auction. What is the expected revenue to the seller if each buyer i values the item according to their selected type t_i? (ie the winner of the auction obtains a utility of $t_i - P$ and the loser has a utility of 0, and in the case of any tied bids an unbiased coin is flipped to decide who gets the item.)

c) In a "Take-It-Or-Leave-It" offer, there are two players; a buyer B and a seller S. The buyer has a value V_B on the item for sale and the seller has a value V_S on the same object. Both V_B and V_S are drawn independently from the uniform distribution on $[0, 1]$. The buyer offers a price P to the seller. If the price exceeds the seller's valuation then the seller accepts. If not, the seller rejects the offer. By maximising the buyer's expected payoff, what price will the buyer choose?

d) Explain why the answer in c) leads to an unsatisfactory outcome compared to the case in which the seller's valuation is known at the outset.

Answer:

a) A Dutch auction is a sequential descending auction where the first buyer to stop the auction pays the current price when he intervenes. So this is the same outcome as if all the players put their bids in sealed envelopes and the auctioneer chooses the largest bid in the envelopes. ie it is the same as a first-price, sealed-bid auction. So bidding your true valuation can only ever gain you zero utility, which is a worse expected utility than a strategy where you bid less than your valuation. In the latter case there is a possibility you might win the auction with a positive utility giving an overall positive expected utility.

b) The possible bid combinations are $(8, 8), (8, 12), (12, 8), (12, 12)$. The revenue to the seller for each of these games is thus $8, 8, 8, 12$. So the expected revenue is

$$\frac{3}{4} \times 8 + \frac{1}{4} \times 12 = £9$$

c) The payoff to B is $(V_B - P)$ but only if the seller accepts. The probability the seller accepts is $Prob(V_S > P)$ and since the valuations are drawn from a uniform distribution on $[0, 1]$ then the expected payoff is

$$(V_B - P) \times Pr(P|P > V_S) = (V_B - P) \times P$$

Maximising this expected payoff wrt P gives

$$\frac{d}{dP}[(V_B - P) \times P] = 0 \Rightarrow P^* = \frac{V_B}{2}$$

d) If the seller's valuation was known in advance then whenever the buyer's value was higher the parties would transact and hence the player who values the object the most would always receive it. But when the seller's valuation is unknown the buyer chooses to shade her bid by a factor of one half. So in order for a trade to occur the buyer's valuation has to be more than twice the value of the seller's estimate, which will occur less often, and the object will therefore not always go to the person who values it the most.

8.8.5 Exercise (Random price, sealed-bid auction)

Consider an auction for a single indivisible good with two bidders, $\{1, 2\}$ Each bidder's valuation v_i is independently and uniformly distributed on the interval $[0; 1]$, and this is common knowledge among the players: A bidder observes their own valuation, but not the valuation of their opponent. Consider an auction where an object is allocated to the highest bidder. The price that this bidder pays is determined by the toss of a coin. With probability one-half, the bidder pays their own bid. With probability one-half, the bidder pays the losing bid. The person bidding lowest pays nothing. If the bids are equal, each bidder gets the object with probability one-half, and in this case, pays their bid.
Suppose that bidder 1 assumes that bidder 2 will bid a fraction, γ, of bidder 2's valuation.

(i) Write down the expected payoff to bidder 1, as a function of their own valuation v_1 and bid b_1. (Ignore ties, where both bidders bid the same amount.)

(ii) Solve for bidder 1's optimal bid, as a function of their valuation.

(iii) Use this to solve for γ and thereby, find the Bayes-Nash equilibrium of this game, where each bidder bids a constant fraction γ of their own valuation.

Answer: Suppose that each bidder bids a constant fraction γ of their valuation. Let v_1 be bidder 1's valuation, and v_2 be bidder 2's valuation. Then bidder 1's expected payoff from bidding b_1 given v_1 is

$$\begin{aligned} U(b_1, v_1) &= \int_0^{b_1/\gamma} [v_1 - \frac{1}{2}b_1 - \frac{1}{2}\gamma v_2] dv_2 \\ &= [v_1 v_2 - \frac{1}{2}b_1 v_2 - \frac{\gamma v_2^2}{4}]_0^{b/\gamma} \\ &= \frac{v_1 b_1}{\gamma} - \frac{b_1^2}{2\gamma} - \frac{b_1^2}{4\gamma} \end{aligned}$$

Differentiate this wrt b_1 and set to zero to get the optimal value of b_1 to find

$$\frac{dU}{db_1} = \frac{v_1}{\gamma} - \frac{3b_1}{2\gamma} = 0$$
$$\Rightarrow b_1 = \frac{2}{3}v_1$$

ie $\gamma = 2/3$ and hence in a Bayes Nash equilibrium, each bidder comes to the same conclusion and therefore each bidder would bid two-thirds their own valuations.