Haberdashers' School

Occasional Papers Series in the Humanities

HABERDASHERS'
BOYS' SCHOOL

Occasional Paper Number Sixty

Game Theory: An Introduction with Particular Application to Economics

Dr Ian St John

Stjohn_i@habsboys.org.uk

November 2022

Haberdashers' School Occasional Paper Number Sixty

November 2022

Game Theory: An Introduction with Particular Application to Economics

Ian St John

Abstract

This paper summarises some basic ideas of Game Theory and applies them to examples taken primarily from the study of Economics. It explains that Game Theory provides an approach to formulating strategies of action when players interact, such as the decisions of one actor affects the decisions and outcomes of action of the other, and *vice versa*. Such situations are frequently encountered in Economics, notably in the sphere of the Theory of the Firm. A range of types of Game are outlined – simultaneous, reiterated, and sequential, along with methods of analysis appropriate to each. It is shown how in most Games equilibrium outcomes can be determined – known as Nash Equilibria. Having reviewed some of the central ideas and techniques of Game Theoretic analysis, the paper applies them to some familiar problems in Economics, such as Oligopoly output, pricing, and advertising decisions, the Free Rider Problem in Public Goods, and Development Economics. We conclude that Game Theory provides a powerful tool for thinking systematically about interactive decision making scenarios, but that its conclusions are less predictively insightful than some might expect since Game Theory models make assumptions that are often unrealistic (such as rational behaviour and perfect information) and because human actors tend intuitively to follow Game Theory reasoning without being consciously aware of doing so. Game Theory enriches our understanding of human decision-taking more than it informs it.

What is Game Theory?

Game Theory is a theory of decision taking in situations where decision-makers are interdependent. A **Game** is a situation in which a limited number of decision takers (called **players**) interact which each other, such that the actions of each of the players affects the outcomes and decisions of each of the others. This means that when a player contemplates an action, they need to reflect upon how the other participants will respond to that action, since this will affect the outcomes of their action. Such actions, which take into account the possible responses of other players to that action, are called **Strategic Actions**. The outcomes for the players or decision takers in these interactive situations are called **payoffs**, and the payoffs to the players in a game can be set out in a **normal form payoff matrix**.

Game Theory analyses these interactive situations and asks what strategies or rules of action the players should follow in order to secure the best payoffs. Such strategies are called **Optimal Strategies**. The 'Theory of Games' was developed in the 1920s and 1930s. The pioneer study was by the French mathematician Emile Borel in the early 1920s, but his ideas were little known and the most important early paper was by Hungarian mathematician John von Neumann in his 1928 'On the Theory of Games of Strategy'. This dealt with the strategy of zero-sum games (where what one person wins another loses) and put forward the idea of the MaxiMin solution to such games. Von Neumann returned to the subject in the 1940s, writing, with Oskar Morgenstern, the *Theory of Games and Economic Behavior.*[1] This initiated a surge of interest in Game Theory, especially at Princeton, where Von Neumann was a member of the Institute of Advanced Study, and at the RAND corporation in California, which researched the use of Game Theory concepts in formulating military strategy. It was at Princeton that John Nash put forward, in 1950, the concept of an equilibrium in non-cooperative games, now known as a Nash Equilibrium, while it was at RAND that the well-known scenario of the Prisoner's Dilemma was developed.[2]

As is apparent, 'Games' exist whenever decision takers interact with each other and shape their behaviour in the light of that interaction, and such situations occur very frequently in social life – in sport, in politics, in negotiations, in war, in relationships. They are, in fact, a ubiquitous feature of life, and Game Theory has been applied to just about all of them. However, one of its earliest and most sustained applications has been in the sphere of economics, relating to the behaviour of firms, consumers,

---

[1] J. Von Neumann and O. Morganstern, *Theory of Games and Economic Behavior* (Third Edition, Princeton University Press, Princeton, 1953).

[2] For an introduction to the history of Game Theory, see W. Poundstone, *Prisoner's Dilemma* (Anchor Books, New York, 1992), and R. Leonard, *Von Neumann, Morgenstern, and the Creation of Game Theory* (Cambridge University Press, Cambridge, 2010).

workers, and governments. Hence, we focus primarily on economics here. But we shall draw on examples from other areas of life, and the principles outlined here apply to Games in general, not just narrowly 'economic' Game scenarios.

Some More Definitions

We have already defined the key concepts of: Game; Player; Strategic Action; Optimal Strategies; Payoffs. In economics the chief players are firms, but they can also include consumers, workers, and investors. In the case of firms, a **strategy** or rule of behaviour might be: 'If a rival firm raises its price, I will keep my price unchanged; if it reduces its price, I will reduce mine'; 'Each year I will launch a new product design'; 'If my rival advertises, I will advertise'. For consumers a strategy might be: 'If a firm raises its price up to £100 I will keep buying the product; if the price goes above £100 I will switch my demand to another firm.' For a worker a strategy might be: 'I will do an hour of overtime if my employer pays me 50% more per hour; if they pay me 100% more per hour I will do two extra hours'; 'I will ask for promotion, if I get it I will stay, if I don't I will leave'. And so on. Players formulate these strategies by estimating potential payoffs from their decisions – where the payoff is something they value and want to achieve. So the payoff for a firm may be profit; the payoff for consumers may be utility; the payoff for workers may be income or the chance of promotion. Given a range of potential strategies (rules of action) that can be pursued, *the Optimal Strategy is then the one which maximises the expected payoff.*

Classifying Games

One of the things that makes Game Theory appear rather confusing is the range of different ways that Games can be classified. These classifications relate to the assumptions underlying the Game – is it one-off or repeated; do the players have perfect information or not; do players make decisions at the same time or one after the other, etc. According to what we assume, the outcomes of the Games tend to differ. Here we outline just a few of the basic classifications of types of Game.

Cooperate or Non-Cooperative

1. A Cooperative Game

This is where players are able to, or wish to, make binding commitments to each other so that they can formulate joint strategies. For example, two firms meet up before setting prices and agree to set prices that are the same; or they may agree to each sell in different sectors of a market.

2. Non-Cooperative Games
This refers to situations where the negotiation and enforcement of binding contracts is not possible.

In economics, we are chiefly concerned with non-cooperative games. In most economies, for example, it is illegal for firms to draw up contracts to set prices or output.

Perfect and Non-Perfect Information

A key differentiator is the question of information. There are two main types of game:

1. Games of Perfect Information. These are situations where each of the players has perfect knowledge of the decision context: they each know all the possible actions of all the other players; they all know all the outcomes for all the players from each decision; and they know the preferences regarding outcomes for all the players. Since all the players are assumed to know these things they are said to be common knowledge.

2. Games of Imperfect Information. In these situations the players lack information regarding aspects of the decision taking context: they don't know the payoffs to the other players; they don't know the history of all the moves in the game; or they don't know the preferences of the other players.

Zero-Sum and Non-Zero Sum Games

A **zero-sum game** is a game where the two players have a directly competing interest. Most actual parlour games are like this – if I win at monopoly or snakes and ladders, you lose. If we play for money, then there is a fixed pot of money – if I win £10, you lose £10. They are 'zero-sum' because the outcome adds up to zero – if I win £10 (+10) then you lose £10 (-10). Similarly, in an election contest between two candidates, if one gains 5 per cent more support the other loses 5 per cent support.
These were the first Games to be studied. Von Neumann wanted to know if it is possible to predict the outcomes of such Games – assuming each player was rational and acted in the light of the possible rational actions of the other player. His answer was that it was possible to predict the outcome of two-person zero-sum games. Such outcomes that can be predicted are called **Solutions**, and they would be equilibrium outcomes, since no player would have a reason to unilaterally change their strategy once they have arrived at it. The solution to these games is called the **MaxiMin** solution.

To give an example. Consider a strategic Game involving an Army and a Guerrilla insurgency. The army has the initiative and decides whether to fight the insurgents in the jungle or at the edge of the cities. The guerrillas must then decide whether to openly attack the army or engage in skirmishes with them. The zero-sum payoff matrix is as follows.[1]

		Guerrilla rebels	
		Open Attack	**Skirmish**
Army	**Jungle Pursuit**	**-5, +5**	**-7, +7**
	Hold the Cities	**+8, -8**	**+1, -1**

Figure 1. A Payoff Matrix for a Zero-Sum Military Strategy Game

We call a table like the one above a **Normal Form Payoff Matrix**. What it does is show the payoffs to two players according to the combinations of actions (moves) that they make. Each square shows the outcomes of a particular combination of moves, with the payoff to the Row player being given first (in this case Army) and the payoff to the Column player (Guerrilla rebels here) being given second. For example, the top left hand square shows that if the army pursues the guerrillas into the jungle the army does badly and its payoff is -5. By contrast, the guerrillas do well fighting in the jungle and their payoff is +5. Because we are assuming a zero-sum game, a loss of 5 to the Army is matched by a gain of 5 to the Guerrillas.

The aim of Game Theory is to predict what will be the outcome of this Game situation given the payoffs of different strategies for the players and assuming they are rational. What Von Neumann demonstrated in his 1928 paper is that such Games have a predictable outcome generated by what he called the **MiniMax** theorem. This states that:

> *The determinate outcome of a two-person Zero-Sum Game occurs when one player Maximises his Minimum payoff (**MaxiMin**), while the other player attempts to Minimise his opponent's Maximum payoff (**MiniMax**). When each player does this, neither player can change their behaviour and do better, so this is the stable equilibrium of the Game.*

Let us apply this principle to the above Game. We start with the Army and assume that the Army wants to Maximise its Minimum payoff from the situation it finds itself in. Should it enter the jungle or stay in the cities? If it enters the jungle its worse payoff is -7, when it hits a skirmish with the Guerrillas. This is its worst outcome in the entire

[1] This example is taken from M. Shubik (ed.), *Game Theory and Related Approaches to Social Behavior* (John Wiley and Sons, New York, 1964), pp. 14-18.

Game. By contrast, the worst outcome the Army can encounter if it remains based in the cities is +1. So strategy 'Hold the Cities' is its MaxiMin. Given this, what will the Guerrillas do? Their strategy should be to Minimise the Maximum payoff to the Army given their choice of the second row (Hold the Cities). This occurs when the Guerrillas engage in the column two strategy, skirmish. Then they have limited the army to only +1, whereas if they had taken the Army on in open combat the Army would have scored 8 and the Guerrillas would have lost 8. Hence, we can predict that the outcome of this Game is the combination of strategies 'Army Holds the Cities' and Guerrillas engage in 'Skirmishes'.

To see that this result is quite logical, take each decision taker in turn. If we start with the Guerrillas, we can see, looking at the column scores for the Guerrillas, whatever the Army does, the Guerrillas will always do best if they skirmish. So, if the Army enters the jungle, the Guerrillas score +5 if they fight them in open battle and +7 if they skirmish. Hence, they skirmish. And if the Army stays in the cities, the Guerrillas score -8 in open battle and -1 if they skirmish. Thus, the Guerrillas always skirmish. Since the Guerrillas will never willingly engage in Open Attack we call Open Attack a **Dominated Strategy** – one they will never use. Equally, the Army always does better in the cities than in the jungle whatever the Guerrillas do – if they enter the jungle they make losses of -5 or -7, whereas if they stay in the cities their gains are either +8 or +1. Clearly, the Army must always stick to the cities and will never enter the jungle (so jungle is a **Dominated Strategy** for the Army). And if the Army is in the cities, and the Guerrillas always skirmish, we end up in the bottom right payoff box, which is the MiniMax predicted outcome.

Such zero-sum games were the first to be studied. However, they are not the main focus of this paper since zero-sum games are not very common in economics. In economic games the outcomes are not fixed, and while players can do better or worse than the other, a win for one does not directly lead to an equal loss for the other. Both players can often enjoy positive outcomes (for example two firms both setting a high price for a good). The games are therefore **Positive Sum** with a range of payoffs.
A more important division between the games we consider in economics have to do with the timing of what strategy to use.

The Timing of Decisions

Here there are three main possibilities.

1. **Simultaneous Games**. This is where each participant or player chooses their optimal strategy **at the same time**, so that they don't know what strategies the other players are actually choosing. In deciding what to do they have to think about what the other players *might* do – but they won't know for sure until the game is played. So,

a firm choosing whether to set a price high or low will need to reflect upon what prices other firms might set; but it won't actually know the prices other firms set until they all announce their prices simultaneously.

2. **Repeated or Reiterated Games**. This is where the simultaneous game is replayed numerous times. In each case the players won't know for sure what the other players will do before they make their decision. But they will have played the simultaneous game on multiple occasions. For example: imagine firms which launch new clothing brands each spring. Every year they have to decide whether to set high or low prices, and each year they won't know what the other firms have decided. However, they will know what those firms have done in previous years.

3. **Sequential Games**. These are situations where players take decisions consecutively. One player makes a decision and then the second player takes a decision, knowing what the first player has done. Thus, in our clothing example: firm one announces the price of its spring collection, and firm two knows this decision when it decides how to price its product and this knowledge can inform its decision as to which strategy to pursue.

We now consider these simultaneous, reiterated, and sequential Games in more detail.

Simultaneous One-Off Games with Perfect Common Information

What we are trying to do when confronted with a one-off interactive game is predict what is the best strategy for a player to pursue, assuming that they are rational and wish to maximise their expected pay-off from the interaction. Remember that the agent is having to decide what to do given that they don't know for sure what the other agent will do – although both *do* know the payoffs for themselves and others of each possible combinations of actions. Both players reveal their actions at the same time, and they do so only once. The question is: what should a player do in this situation? A key step to answering this question is the concept of a Dominant Strategy.

> **A Dominant Strategy is a rule of action that will yield the optimal or best outcome for a player irrespective of what the other player decides to do.**

This strategy is always the right thing for a player to do whatever the other player decides to do. By contrast, a **Dominated Strategy** is a strategy which is NEVER the right thing to do whatever the other player does.

Let us see what a Dominant Strategy looks like.[1] Imagine two companies, A and B, both of whom are releasing a film on the same day. The question is: should each firm advertise their new film? To analyse this question we create a payoff matrix. The figures in the below table refer to the profits the firms make millions of dollars if they do or do not decide to advertise.

		Firm B	
		Advertise	**Don't Advertise**
Firm A	**Advertise**	**100, 50**	**150, 0**
	Don't Advertise	**60, 80**	**100, 20**

Figure 2. An Advertising Game with Two Dominant Strategies

In this matrix, each box contains the different payoffs to each firm. The first number in the box refers to the payoff to the Row player on the left; the second number refers to the Column player. For example, taking the top left-hand box, 100 is the payoff to Firm A and 50 is the payoff to firm B. This payoff matrix summarises the outcomes for the two firms according to whether each firm decides to advertise or not. Thus, if firm A advertises and Firm B advertises as well, the payoffs are $100m for firm A and $50m for Firm B (top left-hand box). If Firm A does not advertise but Firm B does, then the outcome is $60m profit for Firm A and $80m profit for firm B (bottom left-hand box). And so on. The question is: given these four possible sets of outcomes, what strategy should each firm choose given it does NOT know what the other firm will do? Should each firm advertise or not?

Start with Firm A. If it advertises, it will get either 100 profit if B also advertises, or 150 profit if B does not advertise. Its outcomes are therefore 100 or 150. But if Firm A does not advertise, it will get 60 if Firm B advertises and 100 if Firm B also does not advertise. Its outcomes for not advertising are therefore 60 or 100. *In this case Firm A should definitely advertise.* If Firm B advertises, then Firm A will earn 100 if it also advertises and only 60 if it does not. So it earns more profit by advertising. And if Firm B does not advertise, then Firm A will make 150 profit if it does advertise, and only 100 if it does not. In this case, whatever Firm B does, whether it chooses to advertise its new film release or not, Firm A will in both cases make more profit if it DOES advertise its new film. *Since to advertise is the optimal strategy for Firm A whatever Firm B does, we say that to advertise is the Dominant Strategy for Firm A.*

[1] This example is based on R.S. Pindyck and D.L. Rubinfeld, *Microeconomics* (Pearson Educational International, New Jersey, Seventh Edition, 2009), pp. 482-483.

What about Firm B? Should firm B advertise or not? Again, it should advertise. If Firm A advertises and Firm B also advertises, Firm B gets $50m profit; but if Firm B does not advertise its film it gets no profit at all (top right hand box). Hence, if Firm A advertises, Firm B is better off advertising too. And if Firm A does not advertise, then Firm B will make $80m profit if it advertises, and only $20m if it does not. Again, therefore, Firm B does better if it advertises. Thus, Firm B does better when it advertises whatever Firm A does. *For Firm B, too, advertising is the Dominant Strategy.*

When every player has a Dominant Strategy whatever the other player does, we say there is an **Equilibrium in Dominant Strategies.** There is a clear determinate outcome. In this case: the equilibrium outcome is both firms advertise their new films and the payoffs will be $100m profit to Firm A and $50m profit to Firm B.

Equilibrium without Dominant Strategy

Situations where BOTH players have a Dominant Strategy are the simplest to model. Each player then just follows its dominant strategy which, by definition, doesn't depend on what the other player does. But clearly one can imagine many situations where what is the best strategy for one player to pursue depends on what the other player does. We can show this my amending the payoffs in our advertising example. Let us imagine the payoffs to Firms A and B from advertising and not advertising are as follows:

		Firm B	
		Advertise	**Don't Advertise**
Firm A	**Advertise**	**100, 50**	**150, 0**
	Don't Advertise	**60, 80**	**200, 20**

Figure 3. An Advertising Game with One Dominant Strategy

This payoff matrix is the same as the first but with one difference – shown in the bottom right-hand corner. We now assume that if Firm A does not advertise and Firm B also does not advertise, then Firm A will make a profit of $200m. What should Firm A now do?

If Firm B advertises, then Firm A should advertise, since if it also advertises it will make 100, and if it does not it will make only 60. But if Firm B does NOT advertise, then Firm A should also NOT advertise – for if it advertises it makes a profit of 150, and if it does not advertise it makes a profit of 200. The point here is that Firm A **no longer has a dominant strategy**. It can no longer say that it should do one thing whatever Firm B does. Its optimal strategy now depends on what Firm B does: if Firm B advertises, then Firm A should advertise; but if Firm B does not advertise, then Firm A should not advertise either. *Firm A's optimal strategy depends on what Firm B chooses to do.*

What should Firm A do? To answer this, it needs to try and figure out what Firm B is most likely to do. It needs to know what Firm B's payoff outcomes are and see what Firm B's most likely strategy will be. *In this case, Firm B does have a dominant optimal strategy, and that strategy is to advertise its film*. If Firm A advertises, then Firm B will make 50 if it also advertises and zero profits if it does not. So it should advertise. And if Firm A does not advertise, Firm B will make 80 profits if it advertises and only 20 profits if it does not. Again, Firm B is better off advertising. Given this, Firm A can conclude that Firm B WILL advertise. And if Firm B advertises, then the best strategy for Firm A will be to advertise as well, for it then earns 100 if it advertises and only 60 if it does not. Thus, again, the equilibrium outcome of this game scenario is for both firms to advertise.

But note that the two equilibriums are different: in the first the equilibrium emerged from both firms following their Dominant Strategy. In the second example, only Firm B followed a dominant strategy, and Firm A's strategy was a response to that. Both these are examples of Nash Equilibria, but the second is more general than the first.

The Nash Equilibrium

When seeking stable or equilibrium strategies in games, we have seen that when both players have a Dominant Strategy then this can yield a stable outcome. But stable outcomes are also possible when one or both of the players are not pursuing Dominant Strategies. To model these stable equilibria more generally we use the idea of a Nash Equilibrium, put forward by John Nash in 1950.

> **A Nash Equilibrium seeks to locate the best strategy for Player A to pursue in the light of what Player B is likely to pursue, while Player B is pursuing their own best strategy given the likely behaviour of Player A. When both players arrive at this mutually consistent best strategy position they will have no tendency to change these strategies.**

The outcome means that A does their best for themselves given what B does, while B does the best for themselves given what A does – at the same time. When this outcome is actually realised it follows that neither of the players would regret what they in fact did and neither would want to change their strategy even if they subsequently could. The central insight of the Nash Equilibrium is that each player's best strategy depends on what the other player does and *vice versa*. It is thus based on the key feature of Game scenarios – namely the interdependence of players in formulating strategies. Formally, this is expressed as:

> If bc(r) =c* is the best response for column player given what r does, and br(c) = r* is the best response for player row given what player column does, then a **Nash Equilibrium** is the strategy profile (r*, c*) such that:
>
> $c^* = bc(r^*)$ and $r^* = br(c^*)$

A Nash Equilibrium is the logical outcome of many Games and emerges by each player thinking through, in advance, what the other player will do and formulating their strategy in response – but then the other player will revise their strategy in the light of the likely response of the other and so on. This process of second-guessing the moves of each player will tend towards a point where, once reached, neither will want to revise their strategy any longer. Hence the outcome is stable and for this reason is the equilibrium solution of the game.

We have seen equilibrium solutions before, but these emerged when one or both players had a Dominant Strategy. Here there may be no dominant strategy for either player, but still an equilibrium predictable outcome of the game can be arrived at – provided, of course, that both players are rational and have perfect knowledge regarding the payoffs to both parties.

To see how a Nash Equilibrium can emerge, consider the following example from Avinash Dixit and Barry Nalebuff's book, *The Art of Strategy*.[1] The case they consider is two clothing companies, Rainbow's End and B.B. Lean, that sell by mail-order. They each set their prices for shirts at the same time and neither knows what price the other is setting. The prices range from $38 to $42, and the profits for each combination of prices are set out in the payoff matrix below.

[1] A.K. Dixit and B.J. Nalebuff, *The Art of Strategy* (W.W. Norton and Company, New York, 2008).

		B. B. Lean's price				
		42	41	40	39	38
Rainbow's End's price	42	43,120 43,120	***43,260*** 41,360	43,200 39,600	42,940 37,840	42,480 36,080
	41	41,360 ***43,260***	41,580 41,580	***41,600*** 39,900	41,420 38,220	41,040 36,540
	40	39,600 43,200	39,900 ***41,600***	***40,000*** ***40,000***	39,900 ***38,400***	39,600 36,800
	39	37,840 42,940	38,220 41,420	***38,400*** 39,900	38,380 38,380	38,160 ***36,860***
	38	36,080 42,480	36,540 41,040	36,800 39,600	***36,860*** 38,160	36,720 36,720

Figure 4. Payoff Matrix for Two Firms Setting the Price of Shirts

This is a one-off simultaneous game and each player has access to all the information in this table. What price will the companies charge for shirts?

Let's start with B.B. Lean. Looking at all the possible payoffs to B.B. Lean, it can be seen that the highest possible payoff B.B. Lean can get is 43,260, which occurs when it prices at $41 per shirt. So B.B. Lean might be tempted to charge $41. But, B.B. Lean knows that it is not guaranteed 43,260, because the actual payoff it will get at $41 depends what Rainbow's End will charge. Now suppose Rainbow's End suspects B.B. Lean will be tempted by the 43,260 payoff at $41. Well, if Lean does charge $41, the best price for Rainbow's End to charge would then be 40, undercutting Lean by $1. For that will yield 41,600 – the highest payoff Rainbow's End can get given that B.B. Lean goes for $41. Thus, if Rainbow's End expects B.B. Lean to charge $41 then it will charge $40. Yet of course, B.B. Lean knows this too! It knows if it sets $41 then Rainbow's End is likely to set $40. Now if Rainbow's End sets $40, then the highest payoff B.B. Lean can get is if it sets a price of $40 too. If it does this, both companies get a payoff of $40,000. Rainbow's End will expect B.B. Lean to do this because it's the rational thing for Lean to do given Rainbow's End sets a price of $40. But note that if B.B. Lean does set $40 then there is no reason for Rainbow's End to revise its strategy. Given Lean does set a price of $40, there is no better payoff for Rainbow's End than the $40,000 it gets from charging $40. Hence it will stick at $40. And if Rainbow's End sticks at $40, so will B.B. Lean, since it can't do better than charge $40 if Rainbow's End charges $40. So, by thinking through what each player is likely to do in the light of what the other player is likely to do, the outcome of this game will be for both firms to set shirts at $40 in their catalogue. This, then, is the Nash Equilibrium in this game.

Indeed, once the catalogues are printed and each party sees that the other charged $40 for their shirts, neither would want to change their decision in the light of this information. Rainbow's End sees that, given that B.B. Lean charged $40, then they did the right thing by charging $40 – any other price they could have chosen ($38, $39 etc) would have yielded them a lower payoff. They did the right thing to set $40. And the same applies for B.B. Lean: once seeing that Rainbow's End selected set a price of $40, Lean could not have done better than any price other than $40. *Thus, neither would want to revise their decision in the future other things being equal.* Again, then, the Nash Equilibrium is a stable outcome for this game even if it was played again.

To see that this outcome is the logical one for this game, look at it now from the point of view of B.B. Lean. Lean tries to think what Rainbow's End will do. Well, the highest payoff for Rainbow's End is 43,260 and occurs when they charge $41. Lean sees this and therefore decides that if Rainbow's End charges $41, then they will do best if they charge $40. Rainbow's End realises B.B. Lean will conclude this, and sees that if B.B. Lean charges $40 they will also do best to charge $40. And so we end up at $40 as the mutual Nash Equilibrium again.

Once a Nash Equilibrium is apparent to both players, the game will tend towards that outcome. If B.B. Lean expects Rainbow's End to select $40, then B.B. Lean will also charge $40 since that will yield the best payoff that is then possible and *vice versa*. What is more, although we are dealing with simultaneous one-off games here, a Nash Equilibrium is a strategy that no player would want to depart from *once they know the decision of the other*. So, once a Nash point is arrived at, it will persist if there are subsequent plays of the same game since neither player will have an incentive to depart from it.

A Nash Equilibrium exists when each player is doing the best it can given the actions of the other player. So, if the two players are Firms, then Firm A is doing the best it can given what Firm B is doing, and Firm B is doing the best it can given what Firm A is doing. This outcome is stable in that neither Firm will have an incentive to deviate from what it is doing. Our earlier example of two firms choosing whether or not to advertise (Figures 2 and 3) eventuated in a Nash Equilibrium: we saw that for both Firms, A and B, the decision to advertise led to their best possible profits, so both Firms will advertise and will not revise that decision in the light of the other firm's decision to advertise.

A Dominant Strategy equilibrium is a special case of a Nash Equilibrium. However, the existence of a single Nash Equilibrium is not guaranteed. A game may have more than one Nash Equilibrium, or it may have no Nash Equilibrium at all.

How to Locate Nash Equilibria

The following procedure explains how to go about determining which payoffs are Nash Equilibria – i.e. outcomes which neither firm will want to deviate from once they have happened.[1]

Imagine two players, Row and Column. Each player has two possible strategies, A and B. The payoff matrix is as follows:

		Column	
		A	**B**
Row	**A**	**√100, 50**	**√125, 100√**
	B	**50, 100√**	**75, 75**

Figure 5. Locating Nash Equilibria through Ticks

Remember the payoffs to Row are always given first in each box, the payoffs to Column are given second.

To analyse this situation, take each player in turn. Start with Row. We need to work out what is the best strategy for Row to follow given what Column may do. So first, if Column A pursues strategy A, then Row is best off pursuing A also, since if Row does A it will get a payoff of 100, but if it does B then its payoff will be 50. Consequently, if Column does A, Row should do A. We show this by putting a tick next to Row's payoff of 100 in that case. Now imagine that Column does strategy B. In this case again Row is best off doing A, for then its payoff will be 125, whereas if it does strategy B when Column does B, its payoff is 75. Hence, we put a tick next to 125 in the top right-hand box. This completes our analysis of Rows actions.

Now do the same thing from the point of view of Column. Assume that Row does strategy A. In this case Column will be best off pursuing strategy B, since its payoff then will be 100 compared to 50 if it does A. So put a tick next to 100 for B in the top right box. If, however, Row does strategy B then Column will get a better payoff if it does A rather than B. We therefore put a tick next to 100 in the bottom left-hand box. Now look at the payoff table. Quite simply, *any box with two ticks is a Nash Equilibrium.* In our example this is the top right-hand box. This is a Nash Equilibrium since this is the combination of strategies to which the game will tend, and once this outcome is revealed, neither player would want to revise their decision.

[1] C.f. A. Goolsbee, S. Levitt, and C. Syverson, *Microeconomics* (Macmillan, New York, 2016), p. 471.

In practice this Nash Equilibrium might emerge as follows. Row, for example, needs to predict what Column will do. Given the payoffs available to Column (50 and 100 if they choose A, 100 or 75 if they choose B), then Column will do best if they choose B. Row expects Column to choose B, in which case Row will do better to choose A over B, since if Row does A it will get 125, and if it does B it will get only 75. It will choose A. And, of course, Column can predict this too. Column knows Row will choose A, and if Row chooses A then Column does best with B, making 100 compared to only 50 if Column chooses A. Thus, if Row thinks Column will choose B, Row will choose A; and if Column thinks Row will choose A, Column will choose B. And if Row thinks Column will choose B, Row will choose A, in which case Column will choose B, and so we end up in the top right-hand box which is the Nash Equilibrium – as predicted. Each is doing the best they can given what the other does and will have no incentive to change. Thus, if Row did strategy A and Column also did A, this would not be a Nash Equilibrium for Column, since then Column's payoff would be 50, whereas if it shifted to B its payoff would be 100. Column would NOT be doing the best it could given Row's strategy of A, and if the game were played again Column would not stick with A but shift to B. But if, in this one-off game, Row chose strategy A and Column chose strategy B, then when both observed the results neither would, after the event, want to change their action. And this is a Nash Equilibrium, as Hal Varian summarises:

> **A Nash equilibrium can be interpreted as a pair of expectations about each person's choice such that, when the other person's choice is revealed, neither individual wants to change his behaviour**.[1]

In this example, the combination of Row A and Column B is the only Nash Equilibrium. Are there any Dominant Strategies here? Yes. For Row, strategy A is the Dominant Strategy. Row will do A whatever B does. Strategy A is always the right strategy for Row whether Column does A or B. By contrast Column does **not** have a Dominant Strategy: for Column, A is its best strategy if Row chooses B, but B is its best strategy if Row chooses A. So, Column's strategy depends what Row does, but Row's best strategy does not depend on what Column does.

To get the hang of the method of analysing game payoffs, let us consider a more complex example. In this case there are two grocery stores in a town, Spa and Co-Op. Each is faced with the same question: should it leave its store as it is; remodel the existing store; or knock the store down and build an entirely new one? Again, we assume that the game is played once only, and that each store announces its decision at the same time. The payoff matrix is as follows:

[1] H. Varian, *Intermediate Microeconomics* (W.W. Norton and Company, New York, Ninth Edition, 2014), p. 543.

		Spa		
		Build New Store	Remodel Existing Store	Leave store as it is
CO-OP	Build New Store	200, 200	300, 400 √	√400, 150
	Remodel Existing Store	√400, 300	√450, 450 √	300, 175
	Leave Store as it is	150, 300	175, 350 √	350, 300

Figure 6. Locating a Nash Equilibrium in Store Design

To make sense of these complex outcomes, apply the tick-box method we used above.

We start with Spa. Imagine first that CO-OP builds a new store. What is the best thing Spa can do in this situation? Reading across the outcomes for Spa we see that its highest payoff is 400 when it remodels its existing store. So, we tick that outcome. If CO-OP remodels its existing store, we see that the best thing Spa can do is remodel its store, which yields a pay-off of 450. Thus, we tick this. Now imagine that CO-OP leaves its store as it is. Then the best strategy for Spa is to re-model its existing store, which carries a payoff of 350. Again, then, we tick this. It can be seen here that, for Spa, the best thing it can do whatever CO-OP does is re-model its existing store. This always yields its best outcome whatever CO-OP does. Hence for Spa, remodelling the store is a **Dominant Strategy**.

Now let's see what CO-OP should do. If Spa builds a new store, then looking down the column we see that the best thing that CO-OP can do is remodel its existing store, since this gives a payoff of 400. So, we put a tick next to 400 for CO-OP. What about if Spa remodels its existing store? Then we see that CO-OP's best option is also to remodel its existing store – so give the number 450 a tick. Lastly, if Spa decides to leave its store as it is, the best thing CO-OP can do is build a new store, since this yields a payoff of 400. We therefore tick this outcome.

Note that, unlike Spa, CO-OP has no Dominant Strategy. While, if Spa builds a new store or remodels its existing store, CO-OP should remodel its store, if Spa decides to leave its store as it is, then CO-OP should build a new store. What CO-OP does will consequently be influenced by what it thinks Spa will do. There is no one optimal strategy it should pursue whatever Spa does. However, we can see that under no circumstances is it optimal for CO-OP to leave its store as it is. This is never a best payoff whatever Spa does. We call a strategy that a player will never wish to pursue a **Dominated Strategy**.

Is there a Nash Equilibrium in this game? Yes there is. Remember, a Nash Equilibrium occurs when both players are doing the best they can given the actions of the other. It is located where a box contains two ticks – and this is the outcome where Spa chooses to remodel its store, and CO-OP chooses to do the same thing and remodel its store. This Nash Equilibrium is the determinate outcome to which this game will converge providing the players know the various payoffs in advance. As we have seen, for Spa remodelling their store will be their Dominant Strategy – they will choose to do this whatever CO-OP does. Given Spa does, then, remodel its store, then for the CO-OP remodelling its own store is the best thing is can do. We can therefore predict that, in this game and with this payoff matrix, both firms will remodel their stores.

Finally, to revert to our original example in this section, that of Rainbow's End and B.B. Lean. In the table, the best strategies for each player given what the other does is indicated, not by a tick, but by bold italics. So, for example, if B.B. Lean chooses a price of $42, then Rainbow's End is best off choosing a price of $41 with a payoff of 43,260, and this is indicated by the bold italic number in the second square from the top on the left. If Rainbow's End selects a price of 39, then the best B.B. Lean can do is charge a price of $40, yielding a payoff to Lean of 38,400. And so on. Only one box has best-payoffs for **both** players and this is the 40,000 payoff box in the centre, which is therefore the Nash Equilibrium for this game – as we have already seen.
This example allows us to consider a technique for simplifying the search for a Nash Equilibrium – the technique of elimination of strategies.

The Successive Elimination of Strategies

It sometimes happens that a player, when they look at the payoff matrix in a game, will see that there are strategies that they should NEVER play, whatever the other player does. As we have noted, these are called **Dominated Strategies.** Look again at the payoff matrix for the Rainbow's End and B.B. Lean price setting game.

	B. B. Lean's price				
Rainbow's End's price	**42**	**41**	**40**	**39**	**38**
42	43,120 43,120	***43,260*** 41,360	43,200 39,600	42,940 37,840	42,480 36,080
41	41,360 ***43,260***	41,580 41,580	***41,600*** 39,900	41,420 38,220	41,040 36,540
40	39,600 43,200	39,900 ***41,600***	***40,000*** ***40,000***	39,900 ***38,400***	39,600 36,800
39	37,840 42,940	38,220 41,420	***38,400*** 39,900	38,380 38,380	38,160 ***36,860***
38	36,080 42,480	36,540 41,040	36,800 39,600	***36,860*** 38,160	36,720 36,720

Figure 7. Payoff Matrix for Two Firms Setting the Price of Shirts

Does B.B. Lean, for example, have any Dominated Strategies it will never want to deploy? Yes. One such strategy is price 42. If you look down the column of payoffs to Lean if it charges 42, you will see that, for whatever price Rainbow's End charges, Lean will do better if it charges 41 or 40 or 39 than if it charges 42. In every case the payoff to $42 is lower than these three other prices. We can therefore say that price $42 is a dominated strategy that B.B. Lean will never use. We can thus delete this column from the table – it simply isn't relevant. The same is true for the price $38. If you look down the column of payoffs at price $38 you will see that the payoffs at this price can be beaten by some other price. For example, price $39 **always** beats the payoff to $38. So B.B. Lean will **never** choose price $38 and we can delete this column too.

What about Rainbow's End? This company knows that Lean will never charge $42 or $38. This means Lean will charge $41, $40, or $39. Now whichever one Lean chooses, Rainbow's End can always do better if it does not charge $42, since it can do better by charging some lower price. For example, if Lean chooses $41, Rainbow will do better than $42 if it charges $41, $40, or $39. So, Rainbow's End will **never** choose $42, and so we can delete this row. The same applies to Rainbow's End if it charges $38, and we can delete this row also. Hence, by deleting two columns and two rows, we end up with the following simplified game matrix.

		B. B. Lean's price		
		41	40	39
Rainbow's End's price	41	41,580 41,580	***41,600*** 39,900	41,420 38,220
	40	39,900 ***41,600***	***40,000*** ***40,000***	39,900 ***38,400***
	39	38,220 41,420	***38,400*** 39,900	38,380 38,380

Figure 8. Simplified Payoff Matrix after Deleting Dominated Strategies

With this simplified payoff matrix, locating the Nash Equilibrium is straightforward. Thus, considering Rainbow's End's strategy, if B.B. Lean charges $41, the best Rainbow's End can do is charge $40, and so this is figure is highlighted in bold italics. If Lean charges $40, the best price for Rainbow is $40, and this is in bold italics. And if Lean charges $39, Rainbow is best off charging $40. Doing the same for B.B. Lean shows that, whatever Rainbow's End does, Lean is best off charging $40. The only box where both best strategies coincide is the box for $40-$40, and this is the Nash Equilibrium, as we have already discovered.

Best Response Curves

The results we have been representing through payoff matrices can also be exhibited in diagrammatic form through **Best Response Curves**. The below diagram shows the best response curves for Rainbow's End and B.B. Lean for our above game.

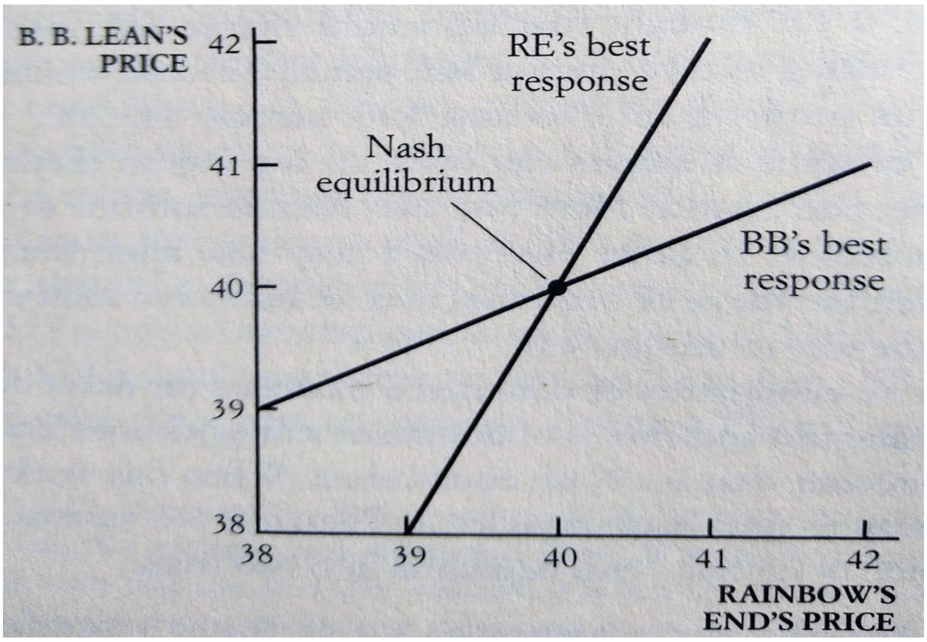

Figure 9. Best Response Curves for the Shirt Price Setting Game

How do these curves work? Each curve shows the best strategy each player can use given what the other player does. For example, taking B.B. Lean's Best Response curve. If Rainbow's End chooses the price \$38, the best price B.B. Lean can charge is \$39. If Rainbow's End selects \$40, then B.B. Lean should charge \$40, and if Rainbow's End charges \$42, then Lean should charge \$41. Rainbow End's Best Response curve shows its best price for any given price set by B.B. Lean.[1]

The virtue of these curves is that they show best response prices for a continuous series of alternative prices, whereas the payoff matrix we considered earlier only showed particular discrete prices varying by whole dollars. Again, the Nash Equilibrium is apparent at \$40 each, where the two Best Response curves intersect. At this point each firm is at its best response point given the pricing strategy of the other, and there will be no tendency for either firm to diverge from this point. If the price combinations for the two companies' shirts are either side of the crossing point then they will tend to converge towards the equilibrium point where they coincide. For example, if Rainbow's End's price is \$41, then B.B. Lean's best response is around \$40.50; but if B.B. Lean's price is \$40.50, Rainbow's End's best price is about \$40.25, and so we gradually converge on \$40 apiece.

[1] These two linear lines are generated by the equations: Lean's Best Response Price = 24 + (0.4 x Rainbow's Price) and Rainbow's BR Price = 24 + (0.4 x Lean's Price). See Dixit and Nalebuff, *The Art of Strategy*, p. 125.

Issues with the Nash Equilibrium

The Nash Equilibrium concept allows us to locate outcomes to which Game situations will tend to gravitate, since each player is doing the best they can given the action of the other player, so neither will have an incentive to shift their strategy away from that position. Yet this kind of simple Nash Equilibrium is far from resolving all the issues even of quite simple game scenarios. Let us look at a few problems which can emerge.

Multiple Nash Equilibria

One problem is that a Game might have more than one Nash Equilibrium. Consider the following example.[1] There are two cereal producers, Kellogg's and Nestle. Two new cereal types have just become available – one is called 'Sweet' and the other is called 'Crispy'. Each firm can launch one only. Should it launch Crispy or Sweet? The payoffs are as follows.

		Nestle	
		Crispy	**Sweet**
Kellogg	**Crispy**	**-5, -5**	**√10, 10 √**
	Sweet	**√10, 10√**	**-5, -5**

Figure 10. Competing Nash Equilibria

First, we locate the Nash Equilibrium. Start with Nestle. If Kellogg's opt for Crispy cereal, then Nestle will make a loss of 5 if it also makes Crispy, but will make profit of 10 if it makes Sweet. So Sweet is its best strategy and we tick this outcome. If Kellogg goes for Sweet, then Crispy is the best choice for Nestle, so we tick that outcome. Now consider Kellogg. If Nestle produce Crispy, Kellogg will do best to choose Sweet, since then it makes a profit of 10 (bottom left-hand corner). We hence tick 10. But if Nestle opt for Sweet, then Kellogg is best off going for Crispy, so we tick this.

As can be seen, there are two Nash Equilibriums in this Game. The top right-hand box, with Kellogg Crispy and Nestle Sweet, is one Nash Equilibrium with two ticks; and the bottom left-hand box, with Kellogg Sweet and Nestle Crispy, is another Nash Equilibrium. Both are stable equilibria, but we don't know which will occur. Indeed, even if the firms knew these respective payoffs for each other, we don't know what outcome will occur since each firm won't know what the other firm will do since the

[1] For this example, see Pindyck and Rubinfeld, *Microeconomics*, pp. 484-485.

potential profits in each case are the same. The actual outcome depends on what the other firm does and neither firm knows this before the game is played.

For example, if Nestle decides on Crispy then if Kellogg also goes for Crispy the two firms will clash and each will make a loss of -5. But if Kellogg goes for Sweet then Nestle will make 10 profit. In a one-off game anything is possible here. Of course, if one firm could get a tip-off about what the other is planning, then it can adjust its behaviour accordingly and a Nash Equilibrium can result. Thus, if Kellogg hears that Nestle is planning to go with Crispy, it will go with Sweet and *vice versa*. Then a stable Nash Equilibrium will emerge. We can see here why the firms would have an incentive to collude. But we are assuming competitive games without collusion. Similarly, if this game were played several times, a Nash Equilibrium might emerge as one firm shifts to what the other is not producing. But we are assuming a one-off game, so this too is not possible.

So in this game, even though there is more than one Nash Equilibrium, we cannot know whether either will actually occur in a one-off play.

Non-Optimal Nash Equilibria – The Prisoner's Dilemma

A Nash Equilibrium can sometimes yield outcomes that are not Pareto efficient. This means, given the outcome of the Game, it would be possible to move to another set of payoffs that left one player better off without making another player worse off. The classic example of this is the game called the **Prisoner's Dilemma** – the most famous of all Games.

This Game originated at the RAND institute in 1950 when two Game Theorists, Merrill Flood and Melvin Dresher, constructed a game in which two players, unable to see or communicate with each other, were asked to pick between two strategies – to either cooperate with the other player or defect against them. What made the Game interesting was that, if each player followed what was clearly their dominant strategy, namely to defect, they would collectively do worse than if they did not follow their dominant strategies and informally cooperated instead. The first two people to play the Game were the economist Armen Alchian and the mathematician J.D. Williams. In this early form the Game was played for small amounts of money, but shortly afterwards, when the Princeton mathematician Albert Tucker was asked to give a talk on Game Theory to the Stanford psychology department, he changed the scenario to one in which two prisoners were confronted with the option of confessing to a crime or staying silent – and this is the form of the Game which is well known today.

This Game models a situation where two persons who committed a crime together have been arrested by the police. The problem the authorities have is that they have very little evidence against the suspects. They therefore need a confession. Their method for securing a confession is as follows: they hold the prisoners in separate cells and make sure they cannot communicate with each other. They then offer each prisoner a deal: if you confess, and the other suspect does not, then you will be set free, but the other suspect will get six years. If both prisoners confess then they will both get three years in jail. But if neither confesses they will get one year each. The point here is that the best outcome for *both players together is not to confess* – then they get one year each. However, the logic of the Nash Equilibrium for this Game pushes them to confess, with the result that they both end up doing three years – ending up, in other words, in a position worse than if they both stayed silent. Why this happens can be seen by studying the payoff matrix for this Game.

		Suspect B	
		Confess	**Not Confess**
Suspect A	**Confess**	**√-3, -3√**	**√0, -6**
	Not Confess	**-6, 0√**	**-1, -1**

Figure 11. The Prisoner's Dilemma

We calculate the Nash Equilibrium as follows. Take suspect A. If suspect B confesses, then the best thing suspect A can do is confess, since then their payoff is a negative 3, representing three years; if B confesses and A does Not Confess, then suspect A is looking at 6 years in jail. So if B confesses, the best thing A can do is confess also. Hence the tick by the payoff -3 in the top left-hand box. What about if suspect B stays silent? If B does Not Confess, the best thing A can do is Confess. Remember if B does not confess and A does, then A gets away scot-free, hence the payoff of 0. If B does Not Confess and A also chooses Not Confess, then A will get a year in jail (bottom right-hand corner). Hence, if B does Not Confess, A should still confess, and we place a tick next to 0 in the top right hand box. Thus, whether suspect B confesses or not, suspect A should still Confess. Hence Confess is A's **Dominant Strategy** – A should confess whatever B does.

The same reasoning applies to suspect B. If A confesses, then B should confess since then B faces 3 years in jail, compared to 6 years if B does not also confess. Thus, we place a tick next to B's -3 payoff in the top left-hand box. If A does Not Confess, then B will do better to Confess, since then he will go free, whereas if he stays silent, too, he will get a year in jail. Accordingly, we place a tick next to zero in the bottom left box. As can be seen, it is the **Dominant Strategy** of B, just as it is for A, to Confess, which means B will Confess whatever A does.

Thus, the police tactic has worked: both the criminals, against whom there was hardly any evidence, have both confessed to the crime with the result that they have each ended up in jail for three years. This is a sub-optimal outcome. If they had both held their tongues they would each be doing just one year in jail. Their mutual best strategy was to stay silent; but they have ended up at their mutual worst outcome – three years in jail each. Yet this suboptimal outcome is the Nash Equilibrium for this Game. This means, even when each player finds themselves doing three years in jail instead of a possible one, they will *still* confess as they did. Given what the other suspect did, each suspect did the right thing and would not change their decision even if they could.

For example, given that B **did** confess, then A would do best to confess, since then A gets 3 years in prison – if they had not confessed then they would be looking at 6 years. And given that A did in fact confess, B would be glad they confessed also, since if they had not then they would be facing 6 years instead of 3. To reiterate: it was the rational Dominant Strategy of BOTH players to Confess, and yet the result of this rational action has been to produce a result that is not efficient. If they had both Not Confessed they would BOTH be better off than they are now, since they would both be doing 1 year in prison compared to both doing 3.

What drives this outcome is the attempt by each player to minimise the worst possible scenario they face. Take suspect A. If A keeps silent and B also keeps silent, then the worse that can happen is that A will have a year in jail. But if A keeps silent and B in fact confesses, then A will get 6 years in jail. By confessing A has ensured that the worst thing that can happen is three years in jail, and if B does not confess then A might even walk free. Confessing limits the worst outcome for A to three years. And the same is true for B. If he doesn't confess then he will get either six years in prison or one. If he confesses he will either go free or do three years. In either case B limits the potential damage by confessing. This strategy of acting to maximise the minimum gain (or in this case limiting the potential loss) is a **MaxiMin Strategy**. Each player is rationally acting to limit the damage from the situation – they are making the best of a bad job, even though by doing so they together ensure that the collective result is worse than they could gave got by staying silent. But staying silent is not the MaxiMin strategy for either since it opens them up to the possibility of a MiniMin outcome, where they do six years in jail.

Such MaxiMin sub-optimal Nash Equilibria outcomes are common in life. Their main characteristic is that it appears as rational for an individual to do some action which it would be irrational if everyone did it. For example, consider two rival countries considering whether to invest in a new, highly destructive (but also very expensive) generation of weapons like the H-bomb. Considered on its own, each country will want to invest in the weapon since, either the other country does not do so, in which case it has a huge miliary advantage, or the other country also invests, in which case the initial country will be relieved it did so also. Of course, the other country reasons similarly, so both countries build the weapons system. This is the Nash Equilibrium to this game. But, again, the outcome is not in the best interest of both countries: they have spent billions of dollars developing weapons that in effect only cancel each other out. Neither is any safer or stronger as a result. They would have been better off agreeing to NOT develop the weapons in the first place since the defence aspects would have been unchanged but they would have saved huge amounts of money. The problem remains, however, that each country will see that it will benefit if it reneges on the deal and builds the weapon if the other country does not. For this reason arms control agreements have a tendency to break down (precisely because they are not Nash Equilibria), and hence have to be monitored very carefully.

Such reasoning is also at the heart of the **Free Rider Problem**, where rational individuals conclude that it is not in their interest to pay for a product or service that would have been provided anyway. It is, for example, always in an individual's interest to ride the subway for free as long as they can get away with it. Whether *you* pay for your ticket or not, it is always in *my* interest not to pay for my ticket. This is my Dominant Strategy. Yet if this is my dominant strategy it is yours also, and everyone else's, so without the threat of enforcement Game Theory would predict that everyone will try to travel for free on the subway – even if this means that at the end of the month the entire system goes bankrupt and all will be forced to walk through the rain instead – a clearly less good outcome for all the travellers.

Prisoner's Dilemma situations can help us explain the strategic behaviour of oligopolistic firms. Consider two airlines, Ryan Air and EasyJet, who have to decide what price to set for tickets on the same route when they do not know what price their rivals will set. The fares are to be announced on the same day.
A possible payoff matrix showing varying profits is as follows.

		Ryan Air	
		High Fare	**Low Fare**
EasyJet	**High Fare**	**50, 50**	**20, 75√**
	Low Fare	**√75, 20**	**√30, 30√**

Figure 12. A Prisoner's Dilemma in Oligopolistic Pricing

This payoff matrix shows the profits each airline will get for differing price policies. For example, if both firms set a Low price for plane tickets, the profits are 30 each (bottom right-hand corner). But if both set a High price, each get profits of 50. What price will the firms set?

Consider EasyJet. If Ryan Air sets a High Fare then EasyJet is best off going with Low Fare, undercutting Ryan Air and gaining market share, since it then makes 75 profit. We indicate this by a tick in the bottom left-hand box. If Ryan Air instead sets Low Fare, then EasyJet is also best off setting Low Fare since then its profits will be 30 (compared to 20 if it were to set a High fare, when it would lose market share to Ryan Air). So we put a tick next to 30. As can be seen, for EasyJet the Dominant Strategy will be to charge a Low Fare whatever Ryan Air does.

What will Ryan Air do? If EasyJet sets a High fare, then Ryan Air will make 75 profits by setting a Low Fare. Hence, we put a tick next to 75 in the top right hand corner. And if EasyJet sets a Low Fare, Ryan Air is best off matching this and setting a Low Fare too. We place a tick next to 30 in the bottom right-hand box. So again, for Ryan Air, setting a Low Fare is the Dominant strategy – whatever EasyJet does, Ryan Air will do best by setting a Low Fare.

Since both airlines have the same Dominant Strategy, to set Low Fares, the result will be that both indeed set Low Fares. And this is the Nash Equilibrium: given what their rival did (set a Low Fare) neither will regret their decision to set a Low Fare, for to have set a High Fare would have left them exposed and struggling for passengers. There will be no tendency to revise their Low Fare decision.

Yet, as in the Prisoner's Dilemma example, the outcome is not an optimal one for either player. If they had both set High fares, then each would earn more profit than they do by setting Low Fares, and the combined payoff to them both would be 100 as against 60.

Thus, Game Theory predicts that, in a non-cooperative situation, an oligopoly industry will tend towards low prices and low profits as no firm can risk charging high prices and being out competed by its rivals.

Of course, this outcome, while good for consumers, is not good for the firms themselves. The best outcome for **both** firms will be high prices, for then they will both make 50. One way to get this is by collusion (e.g. if the directors met before announcing fares and agree to both go High). But this is illegal in practice – and also violates our assumption of a competitive non-cooperative game.

Irrational Behaviour and MaxiMin Strategies

MaxiMin strategies are especially relevant when the other player in a Game scenario behaves irrationally, in ways that do not always conduce to their self-interest. In these situations the behaviour of the other player can be unpredictable and then a strategy of maximising minimum gains can make good sense.

Take, as an example, two software firms, Firm 1 and Firm 2. Firm 1 is the market leader and has the bigger share of the market. Each firm is considering whether to invest in some new software. The payoffs are as follows.

		Firm 2	
		Don't Invest	**Invest**
Firm 1	**Don't Invest**	**√0, 0**	**-10, 10√**
	Invest	**-100, 0**	**√20, 10√**

Figure 13. A MaxiMin Strategy when confronted by a risk of a large loss

Start with Firm 1. If Firm 2 does not invest, then the best thing Firm 1 can do is not invest either. Then its profits don't change, whereas if it were then to invest it would make 100 in losses. So put a tick next to zero in the top left box. If Firm 2 does invest, then Firm 1 will do best to invest as well (since if it doesn't it will make 10 in losses). So we place a tick next to the 20 in the bottom right box.

Now look at Firm 2. If Firm 1 doesn't invest, then Firm 2 will do best to invest. So put a tick by 10 in the top right-hand box. And if Firm 1 invests in new software, then Firm 2 will also do best if it invests in new software also and we place a tick by the 10 in the bottom right box.

In one sense the outcome of this Game is clear. It is the Dominant Strategy of Firm 2 to invest in new software, since whatever Firm 1 does, Firm 2 will do better if it invests. If Firm 2 indeed invests, Firm 1 will do best to invest as well since then it will make 20 as opposed to -10 if it does not invest. Thus, the outcome will be both will invest and this is the Nash Equilibrium.

But suppose Firm 2 doesn't understand this. Suppose it doesn't follow its Dominant Strategy to invest and in fact decides not to invest? Then in this case Firm 2's profits fall by just 10 (compared to if it had invested) and it still breaks even. *But the outcome for Firm 1 is terrible*: if Firm 1 goes ahead and invests large sums in a new software when Firm 2 hasn't invested at all, then it makes a loss of 100! If you are Firm 1 this is worrying and if you think that Firm 2 does behave erratically and you can't be sure they will invest, then you might well have to revise your preferred strategy. The strategy yielding the maximum payoff is to invest; we call this the MaxiMax strategy – Maximising the Maximum payoff (in this case profits of 20). But it carries with it the risk of a very large loss. Hence, a risk averse strategy is not to seek the maximum gain, but to limit the size of the potential loss – a MaxiMin strategy. In this case the MaxiMin strategy for Firm 1 is Don't Invest. While the best that can happen to Firm 1 is now that it just breaks even with profit of zero, the worst that can happen (if Firm 2 actually invests) is Firm 1 makes a loss of 10. This is not great, but it is much less than the potential 100 loss if Firm 2 does not invest. Firm 1 pays a possible price for this conservatism: if it had assumed that Firm 2 would follow its Dominant Strategy and Invest then Firm 1 would have invested too and the result would have been a profit of 20 for Firm 1. So, by following a MaxiMin strategy Firm 1 is not profit maximising. Rather it is loss minimising.

Firm 1's strategy here has been rather cautious. Yes, there was a possibility that Firm 2 would not invest. But how likely was this, given that Invest was Firm 2's clear Dominant Strategy? One way for Firm 1 to refine its decision taking in this case would be to assign probabilities to Firm 2's conduct.

Maximising the Expected Pay-Off

When selecting Game strategies, the players will often not be sure what the other player will do even when the pay-offs are known and they lead to clear Nash Equilibria. This uncertainty can be factored into decision taking by *forming estimates of the probabilities of each course of action by the other player*. Once having done this, a player like Firm 1 can act so as to maximise its *expected* payoff from a Game.

In our example, we have seen that for Firm 2 the Dominant Strategy is to Invest since then its profits will be 10 whatever Firm 1 does, whereas if it does not invest it will make no extra profit whatever Firm 1 does. So, if Firm 2 knows these payoffs and is rational, it really ought to Invest. Suppose, therefore, that Firm 1 considers it 90% likely that Firm 1 will Invest. There is then a 10% probability that Firm 2 will go against its best interests and Not Invest. Given this, Firm 1's expected payoffs for Investing are a 90% chance of making 20 and 10% chance of making -100. Adding these together gives:

$(0.1)(-100) + (0.9)(20) = 8$

Thus, 8 is the expected payoff of Investing given the likelihoods Firm 1 has estimated. If it does Not Invest the expected payoffs are:

(0.1)(0) + (0.9)(-10) = -9

Clearly, Firm 1 should Invest. If, on the other hand, Firm 1 thinks there is a 30% chance that Firm 2 will go against its own interests and Not Invest, the expected payoffs are:

Invest: (0.3)(-100) + (0.7)(20) = -16
Not Invest: (0.3)(0) + (0.7)(-10) = -7

With a 30% chance of Firm 2 choosing Not to Invest, Firm 1 is best off not investing as its likely losses will then be much less.

Thus, what Firm 1 will do in this kind of case depends on the probabilities it assigns to Firm 2's actions, and in a one-off Game these cannot be certain and the outcome will always have a degree of uncertainty.

Mixed versus Pure Strategies

So far we have assumed that there is, for any given one-off Game, one strategy that maximises each player's payoff. These optimal strategies are called **Pure Strategies**. However, there are some situations where a Pure Strategy is not an optimal strategy, and the player should, in fact, vary their strategy in a random way – deciding what strategy to select by, say, tossing a coin. Such strategies are called **Mixed Strategies**. A classic example of a Mixed Strategy being the best strategy is that of taking of a penalty kick. In this situation a penalty taker is faced with two competing Pure Strategies – to shoot the ball to the right of the goal-keeper or to the left (we simplify by excluding a shot down the middle). Equally, the goal-keeper must decide, before the penalty is taken, whether to dive to the right or the left. To model this, imagine standing behind the penalty taker looking on to the goal beyond. The possibilities are that the penalty taker shoots to the left, and the goal-keeper also dives to the left – and saves the shot; the penalty taker goes left and the goal keeper goes right – then there is a goal; the penalty shooter goes right as does the goal keeper – the shot is saved; and the penalty taker goes right and the goal keeper goes left, and there is a goal. The payoff matrix is as follows:

		Goal-keeper	
		Left	**Right**
Shooter	**Left**	**-1, 1**	**1, -1**
	Right	**1, -1**	**-1, 1**

Figure 14. A Penalty Shot Zero Sum Game with no Pure Strategy Nash Equilibrium

So, if the shooter and goal-keeper both go left, the payoff to the shooter is -1 (no goal) and the payoff to the goal keeper is 1. But if the shooter goes left and the goal-keeper goes right, the payoffs are 1 to the shooter (goal) and -1 to the goalkeeper (no save).

Are there any Dominant Strategies or Nash Equilibria in this Game? No. For example, suppose the shooter shoots Left and the goal-keeper jumps right. A goal for the shooter! But this doesn't mean the shooter should always shoot left. If the goal-keeper had also gone left then there would have been a save and no goal for the shooter. And of course, if the shooter always shoots Left then the goalkeeper would soon wise up and jump left in future and the shooter would never score again. Put another way: remember a Nash Equilibrium strategy is one that a player would not change *even once they know what the other player actually does*. So jumping to the right is not a Nash Equilibrium for the goal-keeper. Looking back on his decision he will regret it: I should have gone left, he will say. If in the next penalty the shooter goes left again and this time the goal-keeper also goes left, the penalty will be saved. Now the goalkeeper will be happy – but not the shooter. The shooter will now wish they had gone right. And so on. The simple fact is that in a zero-sum Game of this type there is NO equilibrium pure strategy since there is no outcome satisfactory to *both* players at the same time.

There is, however, an optimal strategy to pursue for both players in this penalty Game. This is a **Mixed** strategy with randomised variation. One can mix strategies in different ratios: so a player could shoot right 70% of the time and left the other 30%. But if the goalkeeper knows that the shooter more often goes right then they will jump right too, so this is not a good mixed strategy. *In the penalty-kick Game the best strategy is to randomise between left and right in a ratio of 50/50*. A shooter could toss a coin between each penalty, going left or right according to what comes up. And the same is true of the goal-keeper. Assuming all other things equal (strength of shot, aim, ability of the goal-keeper) then the goalie will jump the right way 50% of the time and save 50% of the penalties and *vice versa*. And, in fact, studies of penalties in top-flight football do show that shooters and goalkeepers mix going left or right or staying in the centre in equal proportions, suggesting that they do follow the suggested randomised optimal strategy.[1]

Repeated Simultaneous Games

Thus far we have been considering one-off simultaneous games, where each player has to decide a strategy for a once-and-for-all Game. So, for example, EasyJet and Ryan Air had to formulate their pricing behaviour for one particular holiday season. In reality, such scenarios are often repeated. While airlines must announce their prices for this year, next year they will do so again, and the year after, and so on. We call these **Repeated Games**. In these cases, actions are taken and payoffs received over and over again. What is the effect of repetition on optimal strategies in Game situations?

The basic effect of repetition in Games is to promote cooperative strategies. This is because, with repeated plays, the actors realise that by pursuing their immediate perceived gain by defecting, they end up in an outcome that is less good than they could have arrived at by cooperating. Thus, in our Prisoner's Dilemma case, if the same two prisoners were picked up again and again and subjected to the same set of alternatives to confess or don't confess, they will realise that, while to confess seemed their Dominant Strategy, it caused them to end up doing three years in prison – whereas if they both did NOT confess they would only do one year.

But how to get to this cooperative outcome given that the two prisoners cannot communicate? The answer is – they can communicate by the strategies they select. How, in a repeated game, can Player 1 signal to Player 2 that they are willing to cooperate? Quite simply: by not defecting. After the Game is played, Player 2, who maybe did defect, sees that Player 1 did not. What matters now is how Player 2 responds to this information. Of course, Player 2 might think – I'm up against a mug here, and I will defect again. In which case, Player 1 might be forced to prove they are not a fool and defect next time also – so they are both back in the poor mutual defection outcome. But as soon as one player realises that the other is signalling that they wish to cooperate and reciprocates, then the result can be the emergence of cooperation where neither player defects and their joint outcomes are systematically better than when they each tried to cheat on the other.

Does this kind of move from defection to cooperation actually happen? Yes, it seems that it does. The most famous study of this process was undertaken by Robert Axelrod who set up a competition between computer programmers to design strategies to

[1] C.f. Goolsbee *et. al*., *Microeconomics*, pp. 476-7.

follow when confronted by a simple 'Prisoner's Dilemma' type Game.[1] Let's reframe the Game into an economic pricing problem. Each year Firm 1 and Firm 2 have to announce their prices. They can either price High or Low, and the payoffs are as follows.

		Firm 2	
		Low Price	High Price
Firm 1	Low Price	10, 10	100, -50
	High Price	-50, 100	50, 50

Figure 15. Learning to Cooperate in a Reiterated Simultaneous Game

We are familiar with this kind of payoff matrix. As we have seen, the Nash Equilibrium is for both firms to charge a Low Price. But this is only if the Game is played once. If it is played multiple times it is in the long-run mutual interest of the firms to set a High Price – which becomes, in effect, the equilibrium. What Axelrod found is that, if this Game is played over and over again, the most successful computer programme was that which led to the mutual High Price solution and that the optimal strategy that yields this outcome is **Tit-For-Tat (TFT)**. This strategy was submitted by the Game Theorist and mathematician Anatol Rapoport. Tit-For-Tat works on the following simple rule: always begin by cooperating, then in each subsequent play do what the other player did in their last play.[2] So, if Firm 1 goes Low in the first move, then Firm 2 will go Low in the second play. But if Firm 1 goes High, then Firm 2 will go High next move and so on.

You might think they will go back and forth from High to Low indefinitely. But both players know that if they do this they will tend towards the poor-performing mutual defection outcome we studied earlier. What tends to happen is Firm 1 charges a high price. Firm 2 maybe defects and goes low. Firm 1 immediately responds by cutting its price, effectively punishing Firm 2 for going low. Firm 2 sees this and goes High next round. Firm 1 decides to give Firm 2 another chance to cooperate (this is called *being nice*) and go High. So now both have gone High. Gradually this will settle down to be the most common outcome. Whenever either Firm defects and goes Low, the other immediately responds and punishes by going Low also next round. In the long-run, as the Game is played over and over, both players tend towards the cooperate High Price strategy. Thus, Axelrod found that the computer programme Tit-For-Tat had the highest overall payoff at the end of the tournament. It beat all other strategies.

[1] R. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984)
[2] *Ibid.*, p. 13.

This tendency towards cooperation in reiterated games is useful for the understanding of oligopoly behaviour. In oligopolistic market structures we can never be sure how firms will behave in response to their expectations of the behaviour of other firms. In a one-off game such as the price-setting problem of Firms 1 and 2 we would predict that the firms would compete and seek to undercut each other by charging low prices. But if the firms interact on a continuous basis and value future as well as current returns, then we would predict a tendency to tacit collusion with each charging high rather than low prices going forward. The contrasting outcomes of these two strategies are shown below.

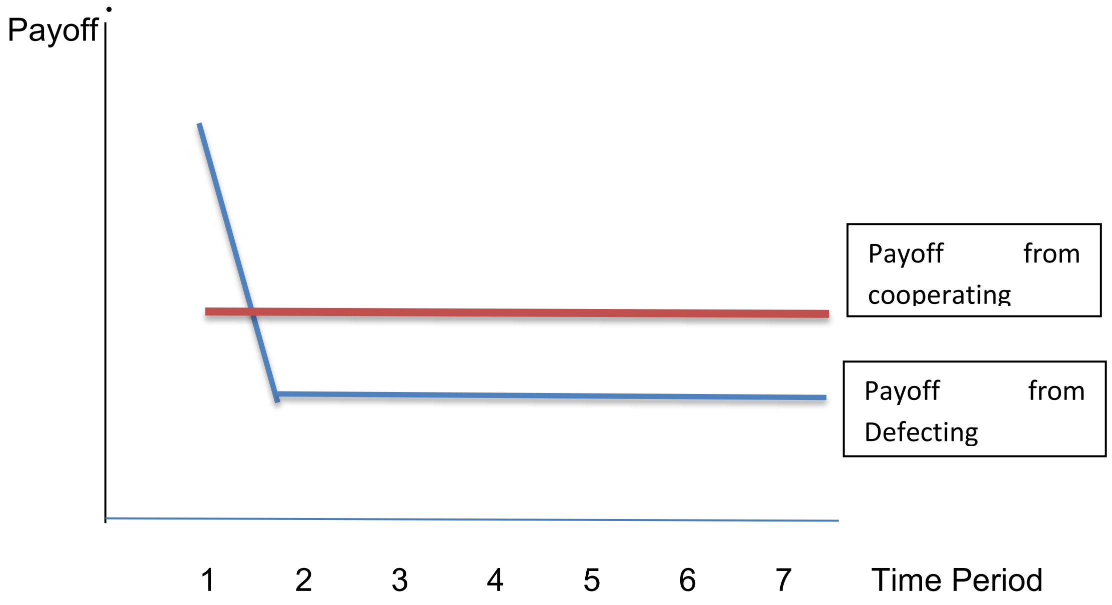

Figure 16. The Higher long-run payoffs to Cooperation

In the Time Period 1 a given firm, like Firm 1, will gain a higher payoff by defecting than it would by cooperating. But if it continually defects it will receive only the lower payoff line. If it cooperates in setting higher prices it will make better returns in the long run. Yet this result may always break down: at any given moment in time a firm can make more by defecting – but suddenly cutting price and initiating a price war. But unless, by so doing, it is able to either remove a competitor or significantly expand its share of the market, it is likely to lock all firms in a cycle of low prices to the mutual loss of all firms. On balance, Game Theory predicts that oligopolistic firms will cooperate rather than compete with each other since this will yield higher payoffs in the long run to all parties.

Tit-For-Tat is a widespread strategy in society. An example, studied by Tony Ashworth, was the behaviour of front-line soldiers in World War One. Soldiers of opposing forces confronted each other for prolonged periods across fixed trench systems. The uncooperative-defect strategy would have seen them blast and fire at one another non-stop in a race to the bottom right-hand corner of a typical Prisoner's Dilemma game. But this, they soon discovered, was in the interest of neither. Instead, the troops learned to informally cooperate. Thus, when sending out patrols at night they took care not to encounter rival patrols from the enemy. When conducting artillery fire they would desist during meal-times or avoid hitting hospital tents or supply roads – since as one British soldier commented at the time: 'It would be child's play to shell the road behind the enemy's trenches … but on the whole there is silence. After all, if you prevent your enemy drawing his rations, his remedy is simple: he will prevent you from drawing yours'.[1] Rifle fire had a habit of going over the heads of the enemy. Of course, this tacit collusion infuriated the generals, and one solution was to continually re-deploy troops along the front so as to break the habits of reiterated interaction. So just when one artillery battalion had learned the timings of lunch, another would arrive and bombard during this time, breaking the informal cooperation and sparking a new wave of Tit-For-Tat.[2]

Sequential Games

Thus far we have been dealing with Simultaneous Games, where each player announces their move at the same time. But it is quite possible that one player makes its move only after seeing what the other player has done – so that Easy Jet only fixes its prices after having seen what prices Ryan Air has set. Such Games, where one player moves first and other players observe their move before formulating their move, are known as **Sequential Games**.

To model such Games we no longer use payoff matrices: we use, instead, a **Decision Tree** or **Extensive Form**. What these allow us to do is plot unfolding actions over time. So, now, we can see how the actions of player A impact on the options of player B and so how player B is likely to respond to the actions of A, and this in turn will influence how A behaves in the first place. We can put this more strongly: when deciding what action to take, player A must consider how player B is likely to act according to the various possibilities that arise when A takes various courses of action. Player B's responses will determine the payoffs to each of Player A's initial moves. Thus, once player A has worked out what B is likely to do, Player A will then act in the light of this information. This is called **Backward Induction.** In essence, what a player must do is sketch out different scenarios as to how the Game will play out; work out which outcome yields them the highest payoff; and then reason back to decide which

---

[1] Quoted *ibid*., p. 79.
[2] T. Ashworth, *Trench Warfare: The Live and Let Live System* (Holmes and Meier, New York, 1980); R.H. Frank, *Microeconomics and Behaviour* (McGraw-Hill, New York, Eighth Edition, 2010), pp. 420-421; Axelrod, *Evolution of Cooperation*, Chapter 4.

action they should initially take. Only once having played out the Game moves theoretically can an initiating player know what move to actually take.

Dixit and Nalebuff introduce the idea with a simple example. Many will be familiar with the *Peanuts* cartoon strip where Lucy offers to hold an American football for Charlie Brown to kick. Charlie Brown must decide whether to accept Lucy's offer or not.

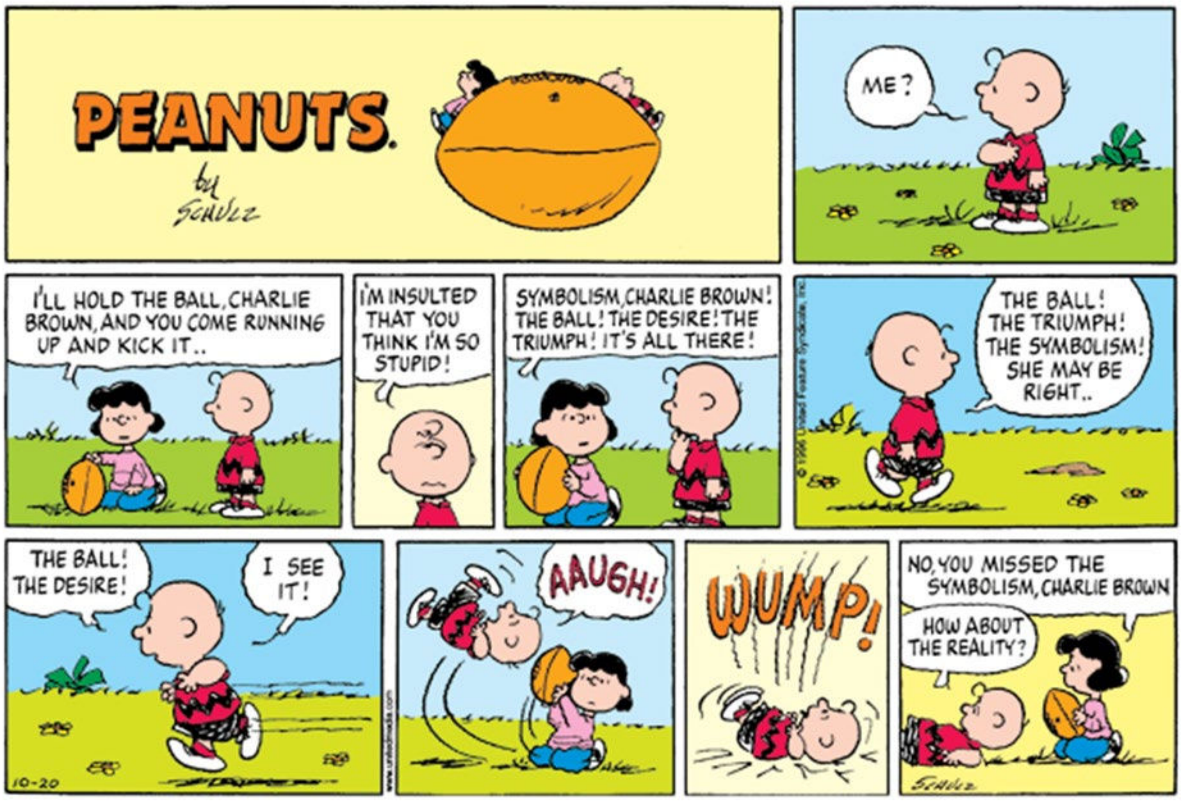

Of course, Charlie Brown accepts the offer every time, and every time he ends up on his back. That he suffers this fate arises out of his failure to carry out Backward Induction from the future likely actions of Lucy. What Charlie Brown needs to do is ask himself the following question: if I opt to take up the offer of the kick, what will Lucy do? Lucy then has two options: either she can hold the ball steady and let me kick it; or she can move the ball away at the last moment and watch me fall flat on my back. Which is she likely to do? Everything suggests that she will swipe the ball away. Given *this expectation of Lucy's behaviour*, what should Charlie Brown do? Obviously, he should not take up the offer and walk away. Alas, Charlie Brown doesn't reason back in this way from likely outcomes and the result is he takes up the offer – and of course, he ends up on his back! This sequence of events can be illustrated through a Decision Tree as follows.

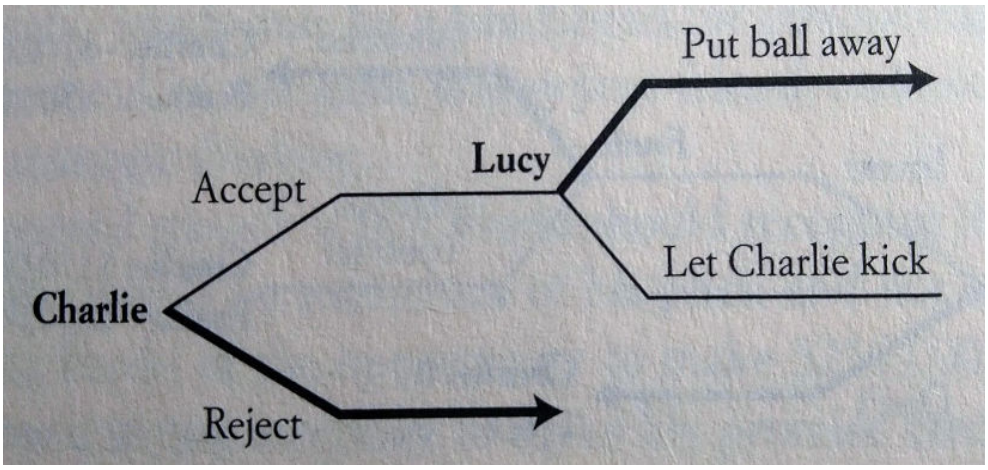

Figure 17. The Extensive Form of Charlie Brown's Kick Dilemma

As can be seen, Charlie Brown can either Accept or Reject Lucy's offer. If he Accepts, then Lucy has the choice of moving the ball away or letting Charlie kick it. Reflecting on her choice, Charlie ought to conclude that it is most likely that she will take the ball away – which is why this is shown with a bold line. In which case, Charlie's choice, reasoning backwards, is between: have the ball taken away and end up on his back; or Reject the offer and walk away. While 'Let Charlie kick' was a possible outcome, it was never likely to happen and so Charlie Brown should have discounted it. In the story that is the option he is tempted by – but as the decision tree shows, this is irrational.

Let us consider a more elaborate example.[1] Imagine two cinema companies, Warner Brothers and Disney, each of which has a new big film to release in a given year. When should each release their big new film? We assume there are three possible dates – March, May, and December. We first present the Game situation as a normal form payoff matrix.

		Disney's Opening Date		
		May	**December**	**March**
Warner Bros Opening Date	**May**	**50, 50**	**√300, 200√**	**√300, 100**
	December	**√200, 300√**	**0, 0**	**200, 100**
	March	**100, 300√**	**100, 200**	**-50, -50**

Figure 18. Normal Form Payoff Matrix for Simultaneous Film Decision

[1] This example is taken from Goolsbee *et. al*, *Microeconomics*, pp. 485-488.

The payoffs here are profits in millions of dollars.

Consider this Game first as a Simultaneous one. Is there a Nash Equilibrium? As before we use our tick-check technique. So, begin with Warner Bros. If Disney opens its film in May, then Warner Bros payoffs are 50 in May, 200 in December, and 100 in March. Since the December payoff is highest, Warner Bros should open in December if Disney chooses May, so we put a tick next to 200 in December. If Disney opens in December, then Warner Bros is best off opening in May, so we put a tick by 300, and so on. Thus apportioning our ticks, we can see that there are two Nash Equilibria in this Game – with Warner Bros opening their film in May and Disney in December; or with Warner Bros opening theirs in December and Disney in May. We cannot be sure which equilibrium will emerge, but the outcome makes a difference: both Warner Bros and Disney do best if they open in May, but not if they both open in May!
We now plot the same figures in the form of a decision tree.

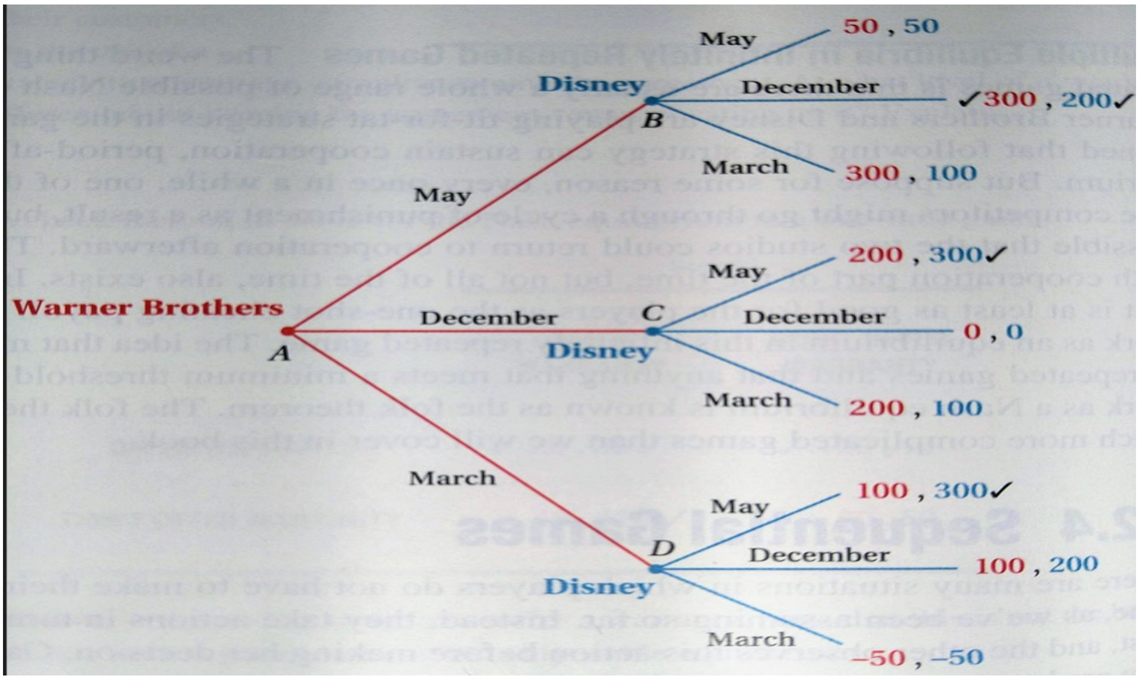

Figure 19. Extensive Form of the Film Opening Game

In this example, we assume that Warner Brothers gets to choose its release date first. Before it chooses the date, it runs through the different scenarios. Remember we assume that both players know all the data in this decision tree: they both know all the possible payoffs guiding the various decisions. The payoffs to Warner Bros are in red, and the payoffs to Disney are in blue.

First Warner Bros asks: if I choose May for the film, what will Disney do? Disney are now deciding at node B between May, December, and March. If Warner Bros release in May, Disney do best ($200 million) to release in December. Hence we tick this December choice for Disney. So: Warner Bros knows that if it releases in May, Disney will release in December, and in this case, Warner Bros will make $300 million. Next Warner Bros asks: what happens if I open my film in December? Then Disney will choose May and we put a tick next to this outcome for Disney ($300m). So if Warner Bros go with December, Disney will go with May. Finally, if Warner Bros open in March, then Disney, choosing at node D, will opt for May, yielding them $300, which we tick. Thus, if Warner Bros opens in March, Disney will open in May.

What, then, will Warner Bros actually choose? In effect, it faces three possible combinations with Disney: May-December; December-May; March-May. Looking at each of these likely outcomes, which is best for Warner Bros? Clearly, its best outcome is to go with the May-December combination, for then its payoff is $300m, compared to $200m and $100m. Reasoning back from these outcomes, it is apparent that May is Warner Bros best choice. Thus, Warner Bros opens in May and Disney delays its film to December. This is the outcome of the Game.

There are three things to note about this outcome:

1. It is a Nash Equilibrium. It is the only outcome to be ticked as a best decision for both players. And neither player would revise their decision once the game is played. Warner Bros would not revise their decision from a May launch, and Disney, given Warner Bros May launch, would not change their release date from December. Each player is doing the best that they can.
2. Where the Simultaneous version of the Game had two possible outcomes, the Sequential version has yielded just one.
3. Of the two possible Nash Equilibria, the one that has emerged is more favourable to Warner Bros. They earn $300m by launching in May, compared to $200m if they had launched in December. This result is the consequence of a **First Mover Advantage.** By choosing strategies first, Warner Bros is able to shift the outcome of the game to their advantage. This is commonly true – though not a universal rule.

A useful way to simplify such decision trees is to prune away branches that simply won't be relevant. Looking at node B, we can see that Disney will never choose May or March, so we can eliminate those branches. At node C, Disney will not select December or March, and at node D, Disney will not select December or March – so these options can be deleted. We thus end up with a simplified decision tree which is easier to navigate.

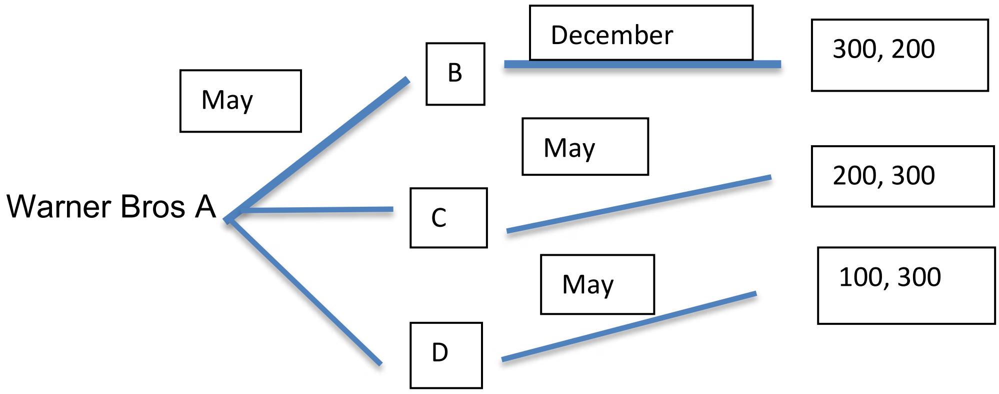

Figure 20. Simplified Decision Tree after 'Pruning' of Dominated Strategies

By simplifying the decision tree in this way, and beginning with Disney's decisions, we see that Warner Bros just has three options to consider – and by looking down the possible payoffs for Warner Bros, it is apparent that the first 300-200 payoff combination is the best for them and so they will choose May.

Strategic Moves and Credible Commitment

Thus far we have assumed that the payoffs in any Game are given to the players and they formulate their actions in the light of them. However, it is possible for players to make **strategic moves** to shift the payoffs in ways that further their interests.[1] One way is to make **side payments**. For example, suppose one player benefits much more from one payoff combination, while for the other player there is little to choose between two outcomes. Then the first player, who is set to gain much more, might make a side payment (in effect a bribe) to the other player to shift the payoffs towards the combination of strategies that best suits them. Provided the size of the bribe necessary is less than the differential advantage they derive from the preferred outcome, both players will be better-off and a new equilibrium is likely.

Another strategic move a player can make is to shape the behaviour of the other player in ways that favours the interests of the initial player. One way to do this is to make a **Credible Commitment** to a certain course of action. An example of this is in the **Game of Chicken**. This game, apparently once popular with American youths, involves two cars driving head on towards each other. The first person to swerve out of the way of the incoming car is a 'chicken' and loses the game, the other player

[1] The concept of Strategic Moves was developed by Thomas Schelling in his *The Strategy of Conflict* (Oxford University Press, New York, 1963).

gaining the credibility of being braver. Of course, if neither player swerves then both die. The payoffs to this Game can be represented as follows.

		Driver 2	
		Swerve	**Straight**
Driver 1	**Swerve**	**2, 2**	**√1, 3√**
	Straight	**√3, 1√**	**0, 0**

Figure 21. The Game of Chicken

In this Game, the highest single payoff to each driver is 3, which occurs when they keep driving straight and the other driver swerves first. The second highest payoff is 2, and occurs when both swerve simultaneously – then it's an honourable draw and neither is the chicken loser. The third highest payoff is 1, when a player swerves first: yes this player is 'chicken', but at least they live to tell the tale. The worst outcome for both players is 0, which happens when neither swerves and both die.

There are two Nash Equilibria in this Game – the bottom left-hand box and the upper right-hand box. Consider Driver 1. If Driver 2 swerves, then Driver 1 will be best off driving straight as then they will win the game. And if Driver 2 goes straight, Driver 1 should swerve as then they at least live. And if Driver 1 swerves, Driver 2 is best off going straight; while if Driver 1 goes straight, Driver 2 ought to swerve.

The problem here is we can't predict what will actually happen. If Driver 2 thinks that Driver 1 will swerve, then Driver 2 is best off driving straight. This is the top right Nash Equilibrium. Unfortunately, Driver 1 might think that Driver 2 will swerve, so they will go straight, potentially yielding the bottom left Nash Equilibrium. But if *both* do this then they both go straight and both end up dead.

How ought a driver play this Game and win? A useful strategic move is to *make a credible commitment to driving straight.* This is sometimes called the 'madman theory'. Thus, take Driver 1. To win the game he needs to convince Driver 2 that he is totally committed to driving straight. If Driver 1 can make Driver 2 believe that, whatever happens, he will drive straight even at risk of death, then Driver 2 will be much more likely to swerve. He might do this by arriving at the road track speaking with bravado about his preparedness to die. He might pretend drunkenness. He might lift his hands from the steering wheel, and so forth. In this way, by showing an irrational willingness to risk death for a game, Driver 2 might be so scared that he will swerve in order not to be destroyed by such a mad rival. In which case, Driver 1 secures the result he wanted – even if it was all a bluff.

It is easy to see how such aggressive posturing to ensure that an opponent believes that a player is prepared to play a Game with reckless disregard for negative outcomes might be a useful tactic in war: when two countries possess nuclear weapons, for example, the country that is able to convince the other that it will actually use them might well achieve its demands by forcing its adversary to back down. The Cuban Missiles crisis of 1962 has often been seen in this light, with both America and Russia seemingly being prepared to launch nuclear war over the stationing of nuclear weapons in Cuba. In this case it might be argued that the American threat was more credible, since they had more at stake over the issue, and Russia blinked first and swerved to avoid Armageddon – though this does not mean that the USA would actually have launched nuclear war over the Russian bases in Cuba. It was enough that the Soviets thought they *might*.

Credible Commitment and the First Mover Question

This use of Credible Commitment can solve a problem that we left unresolved earlier when discussing the outcome of the decision tree regarding the release dates of new films by Warner Brothers and Disney. We found there that there was a *first mover advantage* to Warner Brothers: because Warner Brothers announced the date of their film first they were able to secure the more lucrative May slot, leaving Disney with the less lucrative December date. But if Disney had gone first and opened in May, then Warner Brothers would have had the less lucrative pay off. What decides who in fact goes first and gets the advantage?

Credible Commitment can answer this question. Suppose Warner Brothers is resolved to open in May since this yields the highest payoff – namely $300 million, with Disney opening in December and making $200 million. The below matrix provides a simplified set of payoffs to this Game.

		Disney's Opening Date	
		May	**December**
Warner Bros Opening Date	**May**	**50, 50**	**√300, 200√**
	December	**√200, 300√**	**0, 0**

Figure 22. A Simplified Matrix of the Film opening Game

Figure 22 is a revised version of the payoff matrix we considered in Figure 18. All we have done is remove the March option since, in fact, neither player wanted to open in March. March was a Dominated Strategy and so can be removed. In this revised form, it can be seen that each cinema would prefer to open in May. The May-December combination characterises both Nash Equilibria. If Disney opens in May, then Warner Brothers will be pushed to December and Disney makes $300m and Warner Brothers make $200 million. Alternatively, if Warner Brothers opens in May they get $300m, and Disney will be pushed to December, earning $200m. The best outcome for each is to lead with May. One can see that this is similar to the Game of Chicken. Who will blink first? If neither of them blink and they *both* open in May, then their profits will be only $50. A bad outcome for both.

Just as the scenario is similar to Chicken, so is the best strategy. If Warner Brothers want to get the May slot then they must convince Disney *that they are committed and determined to open in May whatever Disney does.* If they can convince Disney that they are in earnest about opening in May, then Disney will back down, preferring a $200 million payoff to a $50 million. How can Warner Brothers do this? One method is to spend money in ways that commit it to a May opening and which will be wasted if it opens in December. It might take out a big advertising campaign announcing its new film will release in May. It can buy in promotional merchandise linked to a May opening. It can place orders for extra popcorn which will go off if not eaten in May and so on.

How much must Warner Brothers stand to lose if the film opens in December to make its threat credible? The answer is – it must potentially lose a minimum of $151 million. Why? The reason is simple. If Warner Brothers engages in spending and actions that sum to $151 million if the film *doesn't* open in May and in fact opens in December, then it will have changed its own payoffs from the Game. It will reduce by $151 million its payoffs from opening in December. The payoff matrix then becomes as follows.

		Disney's Opening Date	
		May	**December**
Warner Brothers Opening Date	**May**	**50, 50**	**300, 200**
	December	**(200-151) = 49, 300**	**(0-151) = -151, 0**

Figure 23. Revised Payoffs caused by Warner Bros Committing spending to May

The top row of the matrix, corresponding to a Warner Brothers opening in May, stays the same. What has changed is the lower December row for Warner Brothers. As can be seen, the old payoffs to Warner Brothers to opening in December have been reduced from $200 and $0 to £49 and -$151. This is because, if Warner Brothers open in December, the $151 million they spend advertising and promoting their film for a May release will have been lost. Quite simply, Warner Brothers have tied their own hands. They have made it clear to Disney that, for them, a December opening is a Dominated Strategy – they will NEVER open their film in December whatever Disney does. So: while Disney would do best if they opened in May and Warner Brothers opened in December, Warner Brothers have made it clear that they are opening in May whatever Disney does, so Disney, to avoid the disaster of going head to head with Warner Brothers by opening in May and earning only $50 million profit, will 'swerve first' and opt for December, when they will make $200 million, while Warner Brothers will make $300 million – minus any costs it incurs in its Credible Commitment strategy. Provided those costs are less than $100 m (which is quite possible – for example the extra popcorn they bought up will now hopefully be sold) then Warner Brothers will make a net gain by intimidating Disney and securing first-mover advantage.

Credible Commitment and Barriers to Entry

The concept of Credible Commitment can be used in economics to help understand the phenomena of **Barriers to Entry** in Monopoly or Oligopolistic markets. Imagine a monopoly firm which wants to dissuade a rival from breaking into its market. One way it can do this is to threaten a price war if a new firm tries to enter – thereby hopefully dissuading the new firm from even trying. The problem here is that a price-war can easily damage the incumbent firm, so the potential new entrant asks: will the monopoly firm *actually* embark upon a price war, or is it just *bluffing*? If the threat of the price war is not credible, the new firm might well decide to go ahead and enter, the monopoly then facing a situation where it will lose its monopoly position and be forced to accept long-term reduced profits. As in our Warner Brothers example, the monopoly needs to make its threat to launch a price war credible so the rival is scared off, *and it can do this by shifting its own payoffs in the game so that it will be worse off if it doesn't initiate a price war.*

Goolsbee *et. al*. show how this might work. They take a situation where Apple is the only maker of a tablet computer. Samsung are considering whether to try and break into this market. The initial payoffs are represented in the following decision tree.

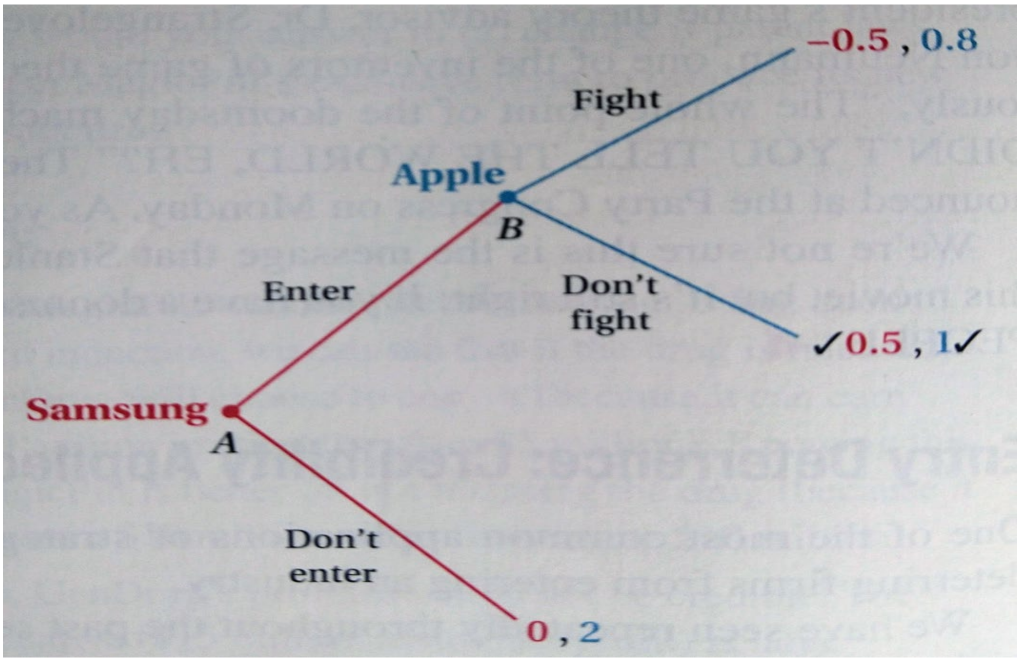

Figure 24. Extensive Form for Entrance Decision Game

Figure 24 depicts the choice facing Samsung: should it enter the tablet market or not? The figures are billions of dollars. If Samsung does not enter the market then it follows the Don't Enter branch and its profits are zero and Apple's profits remain at $2 billion per annum. Now suppose Samsung does enter. We move to node B, and Apple must now choose: should it Fight and initiate a price war or Don't Fight and accept Samsung in the market? If it Fights, then Samsung will make a loss of $ 0.5 billion. This should dissuade Samsung from entering. But, in Fighting, Apple has reduced its own profits to $0.8 billion – compared to the $2 billion it was making when it was a Monopoly. More important, if Apple Fights its profits of $0.8 billion are *less* than the $1 billion it would make if it just accepted Samsung in the market. So, while Apple does best if Samsung does not enter, if Samsung does enter then Apple will do best if it does NOT Fight. And if Apple does not fight, then Samsung will make $0.5 billion profit. *Reasoning back from these outcomes, Samsung will deduce that Apple won't actually fight and therefore the best strategy for Samsung is to enter the market.* The result – Samsung enters and Apple does not fight – is the Nash Equilibrium for this Game.

The problem Apple has in this scenario is convincing Samsung that its threat to initiate a price war is credible, for if Samsung is convinced Apple will fight then Samsung is best not entering at all. As in our Warner Brothers example, *Apple needs to shift the terms of its payoffs in this Game so that Fight becomes its Dominant Strategy if Samsung enters.* One way Apple can do this is by building up excess capacity by investing in extra factories. If Apple does this, and carries extra capacity and produces more output, then its price and profits are likely to fall. As can be seen in the below revised payoff diagram, if Samsung does not enter and Apple remains a monopoly it will make $1.2 billion instead of the previous $2 billion. More significantly, it cuts the profits it will receive if it does Not Fight Samsung from $1 billion to $0.6 billion. The reason this is so important is that Apple will now make lower profits if it accepts Samsung in the market ($0.6 billion) than it would if it Fights Samsung ($0.8 billion). Hence, at decision point B, Apple can be expected to Fight and Samsung now knows this.

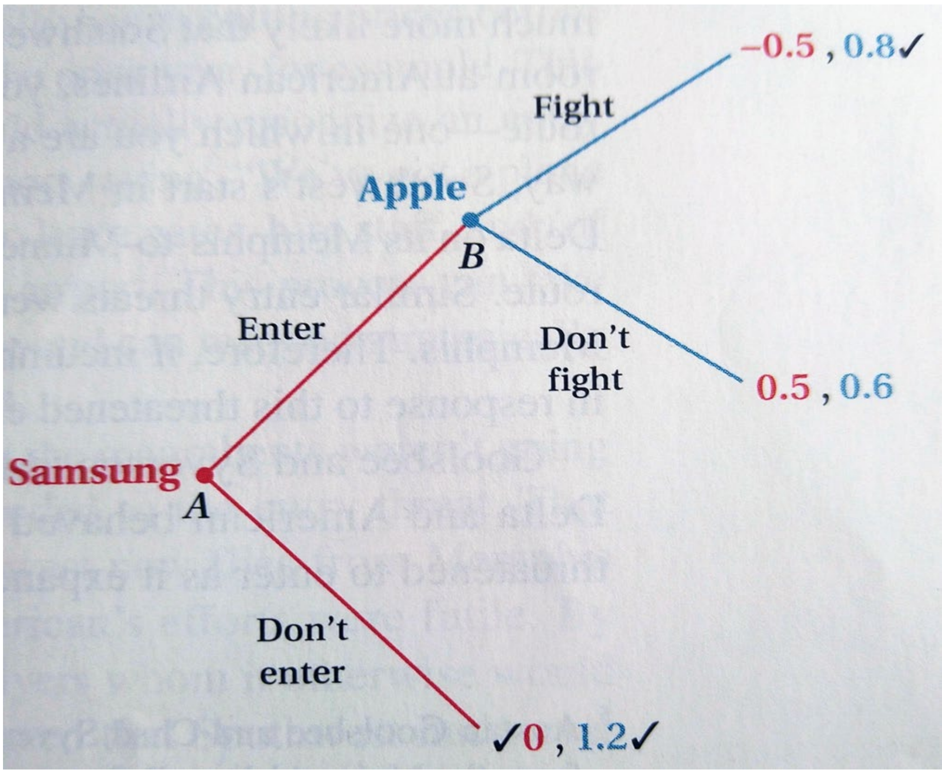

Figure 25. Revised Payoffs from the Samsung/Apple Game

We are assuming that the profit to Apple in the event of a price war remains at $0.8 billion – it hasn't fallen because, if there is a price war, then Apple will need the extra capacity to manufacture greater amounts at a lower price. The effect, therefore, of investing in more plant has been to signal to Samsung that Apple is in earnest about responding to an incursion with a price war. If Samsung enters the industry Apple will fight, and if Apple fights, Samsung will make a loss – and so might well be dissuaded from entering. Working back along the decision tree, Samsung will decide Don't Enter and we end up at the lower branch outcome of zero profits for Samsung and $1.2 billion for Apple. This is the Nash Equilibrium: neither player will revise their strategy in the light of this outcome.

Of course, this outcome is not ideal for Apple. Where they used to make a $2 billion profit as a monopoly without challengers, they are now making only £1.2 billion due to the cost of extra capacity. In effect, they have had to take a profit reduction of $0.8 billion for the privilege of remaining a monopolist in the marketing of tablets. But this trade-off is rational: if they had not been able to stop Samsung entering then their profit would have been just $1 billion, as shown in the previous decision tree. So, even despite the cost of the extra capacity, they are $0.2 billion better off than they would have been. Their strategy has been rational. This is a Game Theory statement of the Limit Price strategy for dissuading entrants many students will be familiar with.

Some Economic Applications of Game Theory

Game Theory is a very large subject and we have only sketched, here, some of the basic ideas. To conclude, we shall consider some applications of Game Theory to economics. Of all the academic disciplines, economics has made the most use of Game Theory, the reason being that the basic assumption of Game Theory – that strategic goal-seeking decision takers formulate strategies in the light of what other goal-seeking actors are doing or expected to do – applies to many situations in economic behaviour. Indeed, the first major study of Game Theory by Von Neumann and Morgenstern was explicitly formulated in the context of economics. In fact, economists had already developed some of the ideas of Game Theory, including the Nash Equilibrium, the basic concept having been present in economics since at least the 1830s. It therefore makes sense to consider some economic applications of Game Theory.

Game Theory's most direct application to economics occurs in the theory of **Oligopoly**. This is to be expected. Oligopoly theory studies the behaviour of a small number of firms (ranging from two to about ten) when those firms respond to the expected behaviour of the other rival firms when setting price, output, product design, advertising campaigns, and so on. Since the firms are **interdependent**, with the reactions of each firm to the decisions of another shaping the payoffs to all, it is apparent that this is a classic Game Theory scenario. And, of course, many of the examples of Games we have considered have basically been oligopolistic ones – such as decisions about pricing and advertising, or launch dates for films or commitments to undertake price wars and so on.

Some Applications of Game Theory to Oligopoly

The Cournot Model of Duopoly

In 1838 the French mathematician, Augustin Cournot, presented a theory of the output decision of two rival firms that produce the whole output of an industry. This situation, where two firms dominant an entire market, is a special case of Oligopoly called Duopoly. It turns out that Cournot's solution was the first discovery of the concept of a Nash Equilibrium, so his model remains important to this day.

The model assumes two firms producing identical products with equal and constant Marginal Costs of production (Cournot in fact considered two firms producing bottled mineral water). The case Cournot considered is that of a **Simultaneous Game** of the kind we considered above: namely, each firm must decide how much output to produce given that each firm makes the decision at the same time and does not know for sure how much the other will produce. Clearly, what each firm decides will impact on the other firm, and their combined choices will impact on each other as their output decisions will help to determine the total output (and hence price) of the industry.

Cournot assumed that, when each firm came to set its output, *it took the output of the other firm as given and fixed at some level.* In other words, Firm 1 will decide how much output to produce in the light of what it *assumes* Firm 2 will produce. This means that of the total market demand for the product Q, Firm 2 will produce Q2, leaving Q1 to be produced by Firm 1, where Q1 + Q2 = Q, and Q1 = Q – Q2. The effect of this is shown in the below diagram.

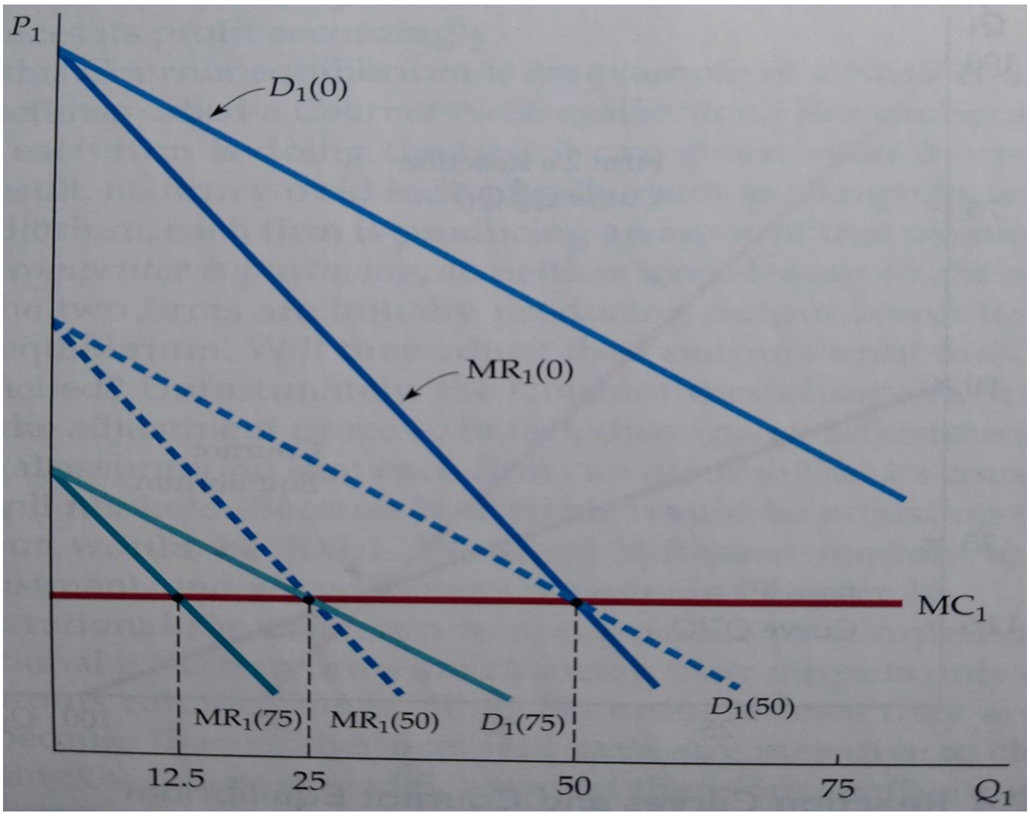

Figure 26. The Cournot Model of Duopoly[1]

In this diagram the total market demand for the good (e.g. mineral water) is shown by the upper blue line D1(0). This is the Total Demand for the product and slopes down from left to right since more is bought the lower the price. The Marginal Revenue line corresponding to this market Demand (Average Revenue) line is MR1(0). The Marginal Cost line shows MC constant at MC1. If this industry were a monopoly, with just one firm, it would produce where MC=MR at output 50.

Now suppose Firm 1 assumes that Firm 2 will produce zero output. In this case the demand curve to Firm 1 will be the demand curve of the industry and Firm 1 will be, in effect, a monopoly producing 50 units where MC=MR. So, if Firm 2 is expected to produce zero, Firm 1 will produce 50 units.

Assume, next, that Firm 1 expects Firm 2 to produce 50 units. The effect of this is to reduce the demand for the product available to Firm 1 by 50 units at all prices. The demand curve facing Firm 1 therefore shifts to the left to D1(50), where the horizontal distance between the total market demand curve D1(0) and the new demand curve to Firm 1, D1(50), is 50 units. The Marginal Revenue curve associated with this new Demand=AR curve is MR1(50). Since Firm 1 is profit maximising, it sets output at

[1] This diagram is taken from Pindyck and Rubinfeld, *Microeconomics*, p. 451.

MC=MR, which occurs at output 25. Note that, with Firm 1 producing 25 and Firm 2 producing 50, the total output of the two firms (Q1 + Q2) is 25 + 50 = 75. The industry is now producing 75 units, which is greater than the profit maximising output 50, and price will be correspondingly lower.

Now suppose Firm 1 thinks Firm 2 will produce 75 units. In this case Firm 1's demand curve is the market demand curve D1(0) shifted horizontally to the left by 75. The result is the Demand line D1(75), with the MR line MR1(75), leading to a profit maximising output for Firm 1 of 12.5.

Finally, if Firm 1 thinks that Firm 2 will produce the 100 units then Firm 1s demand line will be right across to the vertical line and there will be no demand for Firm 1's product at all and it will produce zero. Firm 2 will then supply the entire market demand.

What we have seen here is that Firm 1's output decision is a decreasing function of how much it thinks Firm 2 will make. We summarise the figures as follows:

Firm 2 assumed output	Firm 1 actual output
0	50
50	25
75	12.5
100	0

Figure 27. Firm 1's Output choices in the light of Firm 2's Assumed Output

We can plot these figures as a linear curve, relating Firm 1's output to its expectations of Firm 2's output. This line is called a **Reaction Curve** – and is in fact the Best Response curve we encountered earlier.

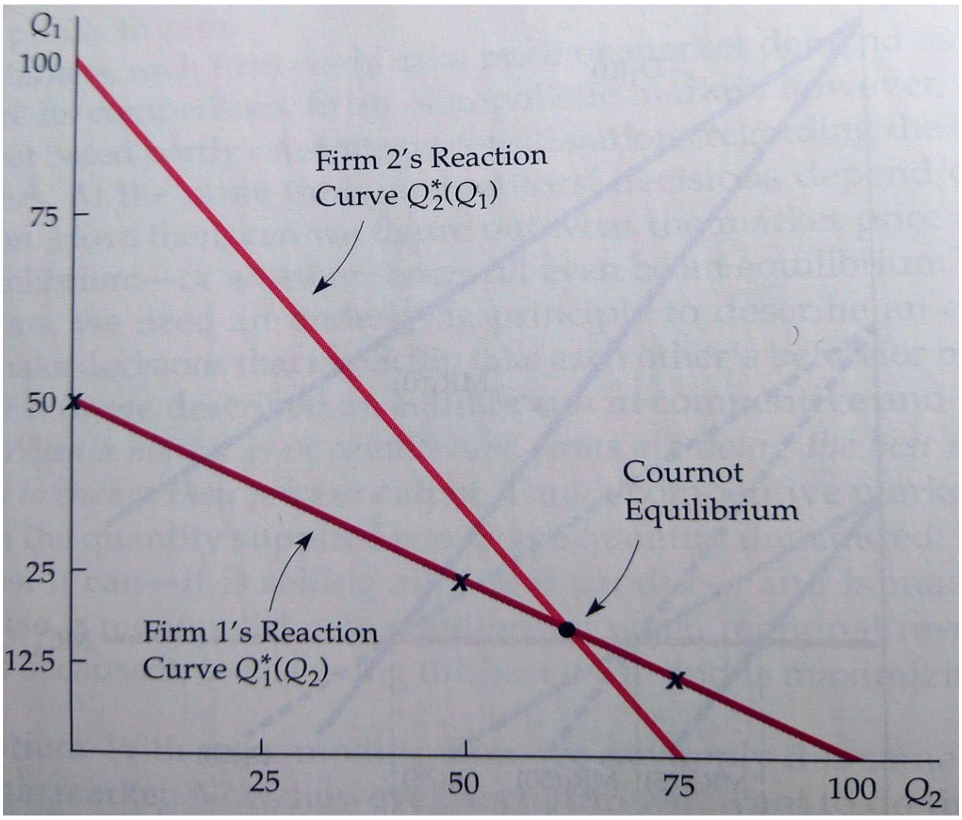

Figure 28. Reaction Curves for Firms' 1 and 2

Thus, Firm 1's Reaction Curve shows *how much it will want to produce for any given output of Firm 2*, ranging from an output of 50 when Firm 2 makes zero, down to an output of zero when Firm 2 makes 100. As we generated a Reaction Curve for Firm 1 in response to the possible outputs of Firm 2, so can we produce a Reaction Curve for Firm 2 in response to the possible outputs of Firm 1. If we assumed identical firms producing identical products with equal MC and demand lines, then Firm 2's Reaction Curve will be the mirror image of Firm 1's and the market demand will be split equally. In our example, Firm 2's Reaction Curve is more elastic to changes in Firm 1's output (its slope being 0.75 compared to 0.5 for Firm 1).

The equilibrium output for the two firms in this market is determined where the two Reaction Curves cross. This is the **Cournot Equilibrium** and is, as we saw above, an example of a Nash Equilibrium. Where the two curves intersect then both firms are producing what they would wish to produce in the light of what the other is producing – and so neither would change their output decision given the output of the other firm. How is the equilibrium achieved? Suppose at first the firms are not at their intersection point – for example, imagine that Firm 1 is producing Q1o in the below diagram.

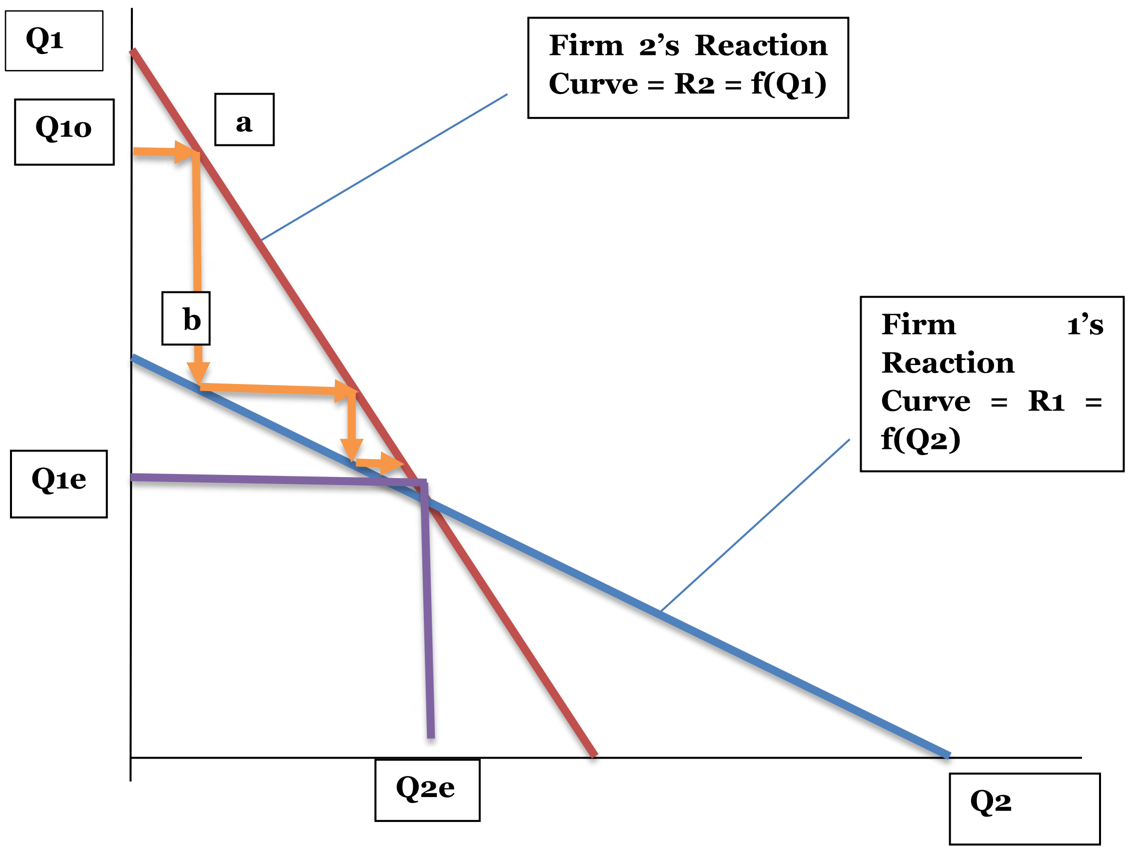

Figure 29. Process for Establishment of Cournot Equilibrium[1]

At this output, Firm 2 would want to produce at the point **a** on its Reaction Curve. When Firm 2 produces at the output corresponding to **a** then Firm 1 will want to produce (given its Reaction Curve) at the lower output corresponding to point **b**, which in turn causes Firm 2 to increase output along its Reaction Curve, which causes Firm 1 to cut output again, leading to a further but smaller increase in the output of Firm 2, and so on. As is apparent, the arrow lines converge on Q1e and Q2e and this is the Cournot-Nash equilibrium, where neither firm will have any reason to revise their output decisions.

The Bertrand Model

Reading Cournot's work, another French economist, Joseph Bertrand, was inspired to come up with an alternative model of Duopoly behaviour. Where Cournot assumed that firms responded to each other's *output* decisions, Bertrand argued that they were

[1] This diagram is based in Frank, *Microeconomics and Behavior*, p. 428.

more likely to respond to one another's *pricing* decisions. He assumed again (like Cournot) that each firm was identical and produced at constant Marginal Cost. If Firm 1 sets a price of P1, then Firm 2 has three choices: either charge a price greater than P1, in which case it would sell zero; charge a price the same as P1, in which case the two firms would split the market between them; or charge a price below P1, in which case Firm 2 would gain all the market and Firm 1 would sell nothing.

The option of setting a price below P1 will always yield more profit so Firm 2 will undercut Firm 1. In Game Theory terms, a price cut is always Firm 2's **Dominant Strategy**. But if price cutting is the Dominant Strategy for Firm 2 it is also the Dominant Strategy for Firm 1. So Firm 1 will then undercut Firm 2. And so on, with both Firms cutting price until eventually P=MC. At this point neither firm will want to cut price any further since to do so will make a loss. So, the outcome of the Game is that for each firm P=MC and the market is shared equally between the two firms. This is the Nash Equilibrium. It is also an example of a Prisoner's Dilemma outcome, in the sense that both firms would have been better off splitting the market at the initial price P1 rather than the lower price P=MC. This is apparent if we set up this game in terms of a payoff matrix.

		Firm 2	
		Price P1	**Lower Price P2**
Firm 1	**Price P1**	**100, 100**	**-50, 200√**
	Lower Price P2	**√200, -50**	**√0, 0√**

Figure 30. Payoff Matrix for the Bertrand Model

Each firm faces a choice of setting a price P1 or undercutting it with a lower price P2. The figures shown are Super Normal Profits in £millions, and the negative figures are losses corresponding to the fixed costs of production (£50m). We establish the Nash Equilibrium point by the usual technique and find that it is at the Lower Price P2. Quite simply, for both Firms Lower Price is the Dominant Strategy.

So, at the higher price P1, Firm 1 will make profits of 200 if it charges Lower Price P2 when Firm 2 charges P1, while Firm 2 makes 200 profits at P2 when Firm 1 charges P1. If Firm 1 charges a price of P1 and Firm 2 undercuts it with P2, then Firm 1 won't sell any output at all and its revenue will be zero, leaving it with a loss of £50m as it still has to pay its fixed costs. The same will happen to Firm 2 if it is undercut by Firm 1. *Each firm will do better charging a lower price whatever the other firm does.*

The upshot is that both firms charge the low Price P2 and this is the Nash Equilibrium since neither would then want to raise their price given the others P2 strategy. Supernormal profits have been competed away to zero. Of course, as is typical in such cases, both firms would have done better to price at P1 and make £100m profits each. But in the absence of binding collusion (which is illegal) the logic of this game is for each firm to undercut P1 and initiate a price war driving price down to P2.

The Stackelberg Oligopoly Model

In the Cournot duopoly model both firms set their output decisions simultaneously and assumed that *their own output decision would not affect the output decision of their rival*: remember that the model worked by Firm 1 formulating its output based on its expectation of the output of Firm 2. For Firm 1, Firm 2's output was fixed, and Firm 1 was, in effect, filling the gap in the market left by Firm 2. But clearly this is wrong since the very existence of Reaction Curves for BOTH firms means that just as Firm 1 formed its output decisions based on the expected output of Firm 2, so Firm 2 formed its output decisions based on expectations of Firm 1. The decisions of both firms are simultaneously self-determining.

Heinrich Von Stackelberg, a German economist who published his theory of duopoly behaviour in 1934, developed a Leadership Model that relaxed the assumptions of Cournot's work in two ways.[1] First, the Game is not simultaneous but sequential: one firm sets its output and the other firm reacts to that *actual* output – not the *expected* output as in the Cournot model. Second, the firm that sets its output first knows that its output decision will influence the output of the other firm. For example, where in the Cournot model Firm 1 sets its output in the expectation that its decision will not affect Firm 2 (which has already selected its output), in the Stackelberg model Firm 1 knows that the output it chooses **will** affect the output of Firm 2 because it knows Firm 2's Reaction Curve: it knows that Firm 2 will adjust its output in response to the output Firm 1 produces. The firm that sets its output first is the **Leader**, while the firm that sets its output second in the light of the Leader's output is the **Follower**. This is, then, a **Sequential Game** and the basic point to emerge is that there is a **First Mover Advantage**: the firm that sets its output first achieves a better payoff in profits than the firm that sets its output second.

To understand the economics of this, refer again to the Cournot duopoly diagram.

[1] Stackelberg joined the Nazi Party in 1931 and became a member of the SS. He died in Madrid in 1946.

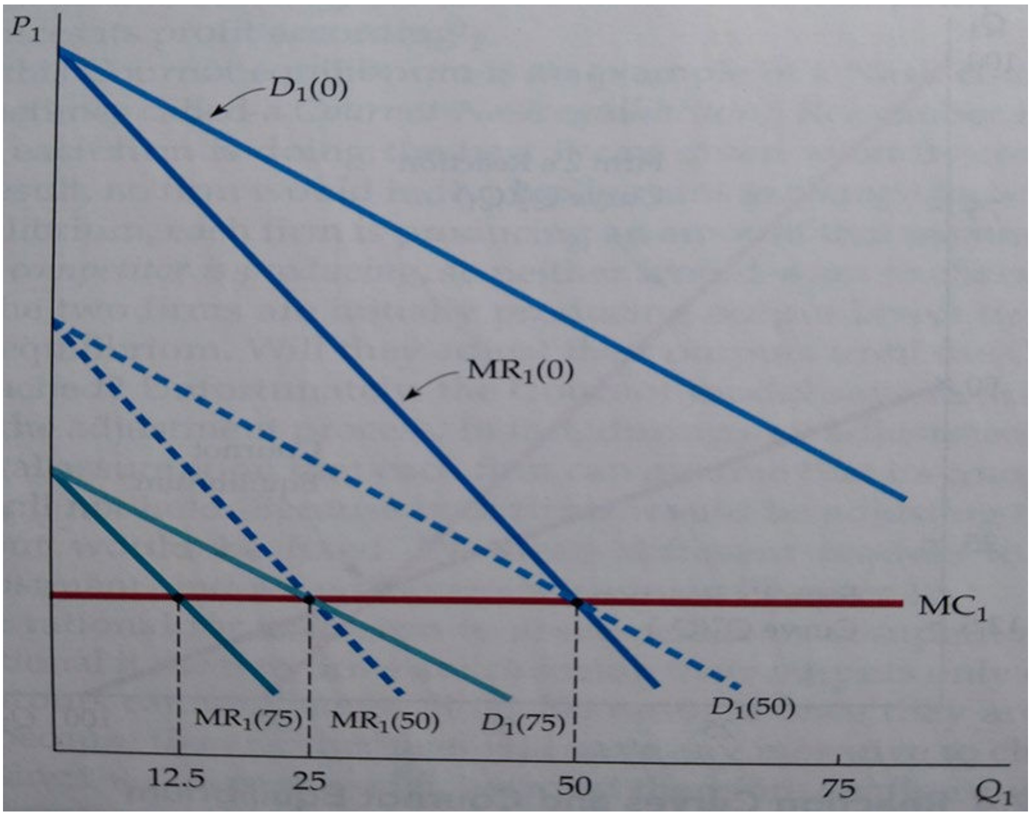

Figure 26. The Cournot Model of Duopoly

It will be remembered that D1(O) is the total market demand curve for the product and D1(50) is the demand curve facing Firm 1 when Firm 2 is assumed to be producing 50 units, to which there corresponds an MR line MR1(50). The point to note is that Firm 1's demand line D1(50) is drawn *on the assumption that the output of Firm 2 is fixed*. However, we have seen, this is not true, since Firm 2's output varies with Firm 1's output according to its own Reaction Curve.

Quite simply, the more Firm 1 produces, the less Firm 2 will produce. What this means is that the slope of the Firm 1 demand curve as shown, D1(50), is too steep (i.e. too inelastic). Yes, as Firm 1 increases its output it will have to lower its price. *But it won't have to lower its price at the same rate as the market demand curve D1(0) suggests since as it increases its output, Firm 2 will lower its output, so the net increase in output will be smaller than the diagram suggests.*

In terms of the diagram, Firm 1's demand curves will pivot outwards and have a lower gradient, and as the demand curves pivot out so do the MR curves, and this in turn means that at, say, the equilibrium point corresponding to D1(50), where MR1(50) = MC1, MR1(50) will be higher than shown, and so will exceed MC1, and Firm 1 will want to increase output until they are equal again. *Thus, the effect of assuming that an increase in the output of Firm 1 will cause the output of Firm 2 to fall is that Firm 1* *will increase its output and take a larger share of the market at the expense of Firm 2*. And this means that Firm 1 will make more profit than Firm 2. This is the prediction of the Stackelberg model. Of course, the same thing would happen if Firm 2 increased its output first. The general point is that: in this Sequential Game, where the action of one firm impacts on the action of the second firm, the first firm to act will enjoy a first-mover advantage.

Earlier we explained how first-mover advantage could be modelled through decision trees. Let us see how this works in this case.

There are two firms, Leader and Follower. Leader faces a choice between two outputs, 6 and 9, while Follower must choose between 7 and 10. We first depict their payoffs in a normal form matrix.

		Follower	
		Output 7	**Output 10**
Leader	**Output 6**	**66, 77**	**√48, 80√**
	Output 9	**√72, 56√**	**45, 50**

Figure 31. Payoffs to Leader and Follower in Stackelberg Model

There are two Nash Equilibria in this Game. Either Leader produces 9 and Follower 7, or Leader produces 6 and Follower 10. The outcomes of these two equilibria are not the same: when Leader produces more than Follower, Leader makes the most profit; when Follower makes more than Leader, Follower makes the most profit. This is the implication of the Stackelberg model. So which equilibrium will prevail? It all depends on who sets their output first. The process can be tracked in the decision tree expression of this Game.[1]

[1] This decision tree is taken from the EconPort website: http://www.econport.org/econport/request?page=web_experiments_modules_stackelberg

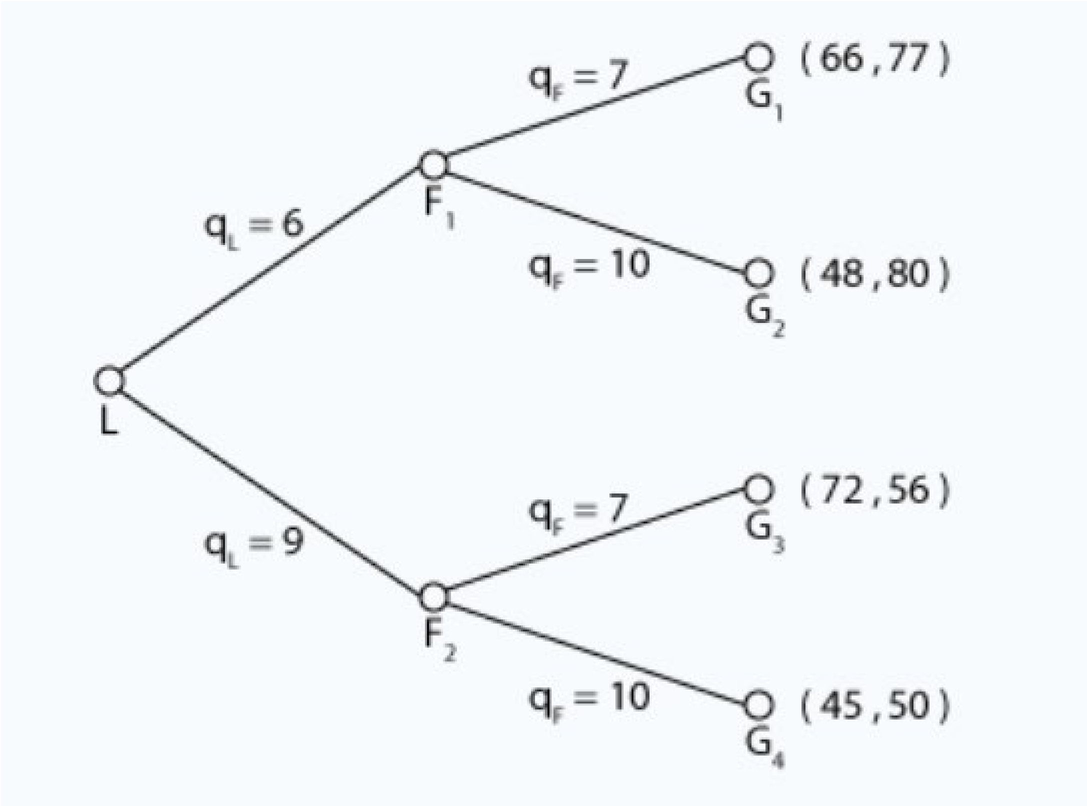

Figure 32. Decision Tree for Stackelberg Output Decision Game

The Leader has a choice of two outputs, 6 or 9. It first reflects what will happen if it chooses 6. Follower, at F1, will then choose output 10, since this will give it the higher payoff of 80 compared to 77 if it chooses 7. So Leader assumes that if it chooses the output of 6, then Follower will choose 10, leaving a payoff to Leader of 48. Suppose, now, that Leader produces 9. In this case, Follower will choose 7 since this maximises its possible payoff at 56. This is good news for Leader, which now makes a payoff of 72. Clearly, given that Leader's payoff is higher under output 9 than output 6 (given Follower's responses), Leader will opt for output 9, and the Game will end up at point G3. This is a Nash Equilibrium: given the action of the other player, neither firm would change their output choices at this point. Whoever selects their optimum output first makes higher profits.

Oligopoly and Advertising

As I write this, the TV schedules are full of Christmas adverts by the major food retailers – Marks and Spencer, Tesco's, Aldi, Asda and so forth. Why do these food retailers advertise so much? Its slightly odd, since, really, few people need to be reminded to buy food for Christmas, and the costs of these adverts is very high. It is quite likely that the various retailers would be better off not advertising at all and just enjoy the pure benefit of the increased Christmas food spending. This, of course, is just what does *not* happen. And Game Theory helps to explain why.

Why do firms advertise at all? In most cases its not to increase total demand for the product they make; it is, rather, to increase their share of the total market. Food adverts at Christmas will, probably, make each of us buy a bit more food. But the reason Asda or Sainsburys advertise is so that we buy our Christmas food from *them* and not some rival. This is why their advertise. The problem is, they are essentially locked in a Prisoner's Dilemma.

Consider two rival retailers, Tesco and Sainsbury's. Their payoffs from advertising are as follows (in millions of pounds). In each box the payoff to the Row player is given first (Tesco's) and the payoff to the second player (Sainsbury's) is given second.

		Sainsbury's	
		Advertise	**Don't Advertise**
Tesco	**Advertise**	**√30, 30√**	**√70, 20**
	Don't Advertise	**20, 70√**	**50, 50**

Figure 33. Payoffs to Tesco and Sainsbury's from Advertising

The Nash Equilibrium for this Game is for both firms to advertise, even though they would both be better off if neither advertised (so the outcome is similar to a Prisoner's Dilemma, when both prisoners would be better off being silent but both confess). So, starting with Tesco, if Sainsbury's Advertises, Tesco will need to advertise too, since if it does not its share of the Christmas market will be squeezed and its profits will be only £20m, whereas if it advertises its profits will be £30m. So, we put a tick next to Tesco's £30m payoff in the top left-hand box. But if Sainsbury's does NOT advertise, Tesco's will benefit even more by advertising as it can gain market share from the silent Sainsbury's, so hence we put a tick next to the payoff to Tesco's of £70m in the top right-hand box. Obviously, Advertise is Tesco's **Dominant Strategy** – it will run a Christmas advert whatever Sainsbury's does.

The same reasoning applies for Sainsbury's. If Tesco advertise, Sainsbury's will do better to advertise too, since their payoff then will be £30m, compared to only £20m if they don't advertise. Hence, we place a tick next to £30m for Sainsbury's payoff in the top left-hand box. And if Tesco do NOT advertise, Sainsbury's will do even better and earn £70m if they advertise and thereby take market share off Tesco's. So we place a tick next to £70m for Sainsbury's in the bottom left hand box.

As can be seen, the only box with two ticks is the top left-hand box – where both shops run Christmas adverts and earn £30m in profits each. This is the Nash Equilibrium – which means that, in the light of what the other retailer does, neither regrets its decision to run a Christmas advert campaign. By contrast, suppose Tesco's had decided not to run a campaign and now it finds Sainsbury's has – well then Tesco's would regret that decision and would (if it could) reverse it, since by not advertising it makes only £20m profit to Sainsbury's £70m. Advertise-Advertise is the only outcome that neither retailer would change even if it could.

The irony is, of course, that both firms would be better off if they abandoned Christmas adverts altogether, for then each would earn £50m in profits compared to £30m. The problem is neither firm will occupy that position for, as we have seen, it is always in the individual retailer's interest to advertise irrespective of what the other retailer does. Any retailer seeking not to advertise, and thereby make the bottom right-hand box possible, will be badly hit by the other firm that DOES advertise. So no one takes that risk and both firms advertise – and as a result our TV screens are saturated with food adverts by oligopolistic retailers. Whether this is an optimal outcome for consumers is another question!

The Free-Rider Problem and Public Goods

A familiar problem discussed in economics is that of the under-provision of Public Goods. Public Goods are ones which, if they are provided at all, everyone benefits from them – such as a general fire service or a flood protection barrier. The problem is that, if everyone benefits from them whether they have paid for them or not, everyone will have an incentive to avoid paying – with the result that the good is not provided at all.

This outcome can be modelled through Game Theory. For example, imagine two electronics stores operating in a shopping mall.[1] There is no security guard, and each firm loses £300 a week due to theft. Assume that a security guard patrolling the mall would cut these thefts to zero, but the guard costs £500 a week to hire. No one store would wish to hire a guard since the reduction in theft (£300) is less than the guard's salary (£500). Would it be in the interests of the two firms to join together and hire a guard? Yes. The cost per store would be £250, and the savings per store would be £300. But Game Theory predicts they won't hire a guard at all and will both be worse off as a result. The reason is shown by the below payoff matrix.

[1] This example is taken from J.M. Perloff, *Microeconomics with Calculus* (Pearson, Harlow, Second Edition, 2011).

		Store 2	
		Hire Guard	Don't Hire Guard
Store 1	Hire Guard	-200, -200	-200, 300√
	Don't Hire Guard	√300, -200	√0, 0√

Figure 34. Paying for a Guard and the Free Rider Problem

The above table shows the change in profits if the stores do or do not hire a guard. If each firm, by itself, hires a guard then each firm will lose £200. This is because the guard costs £500 a week but only saves £300 in thefts, so each firm, left to its own devices, will not hire a guard. This is shown by the -200 entries in the matrix. But what if one firm hires a guard and they split the costs? Both would be better off. However, this won't happen. Consider Store 1. If Store 2 hires a guard and then asks Store 1 to contribute £250 towards the costs of paying them, Store 1 will refuse. If Store 2 hires a guard then Store 1 gets the £300 benefit of reduced theft for free. So it is better off not contributing and 'free-riding' on Store 2's guard. Thus, we tick 300 in the bottom left-hand box as the optimum strategy for Store 1 if Store 2 hires a guard. And if Store 1 hires a guard then Store 2 is better off not contributing and getting £300 in reduced thefts for free.

For both players the Dominant Strategy is not to hire a guard: whatever the other store does, do not hire a guard! As a result, the firms end up in the bottom right-hand corner. The two ticks in this box show that this is the Nash Equilibrium. Thus, in such cases where one player cannot be excluded from the benefits the other pays for, the outcome will tend to be that no mutually beneficial good will be provided. This is why Public Goods usually have to be provided by the state, which can force everyone to pay. In the shopping mall example, maybe the owner of the mall could hire the guard and force both stores to provide £250 to the cost. They will both be better off as a result compared to their Nash Equilibrium.

Economic Development

As a final example of the use of Game Theory in economics, we consider a problem in economic development.[1] Imagine a developing country with two firms – a Steel Firm and a Ship Builder. The firms are interdependent. If the steel firm is to invest in new plant, it needs the market that the ship builder will provide. And if the ship builder

[1] C.f. M. McCartney, 'Game Theory: A Refinement or an Alternative to Neo-Classical Economics?', in E. Fullbrook (ed.), *Real World Economics* (Anthem Press, London, 2007), pp. 156-157.

is to open a new workshop it will need the materials the steel firm will provide. A profit payoff matrix in this situation might be as follows:

		Ship Builder	
		Not Invest	Invest
Steel Firm	Not Invest	√0, 0√	0, -7
	Invest	-7,0	√3, 3√

Figure 35. Payoffs to Investment for two Interdependent Firms

The figures here are profits. If neither firm invests, their mutual profits are zero. If either invests in plant, and the other does not, the investing firm makes large losses of -7. But if they both Invest they make their best joint returns of 3 since each provides services to the other. The problem is that there are two Nash Equilibria, Invest and Not Invest. The Game could eventuate in either outcome, but Not Invest is more likely. This is because, although both players will know that Invest yields their joint highest returns, there is a risk for any one in Investing if the other for some reason (say a failure to reach terms with a bank) fails to invest. Then they will be left with large unused capacity and large losses (-7). Any risk averse firm wishing to minimise potential losses (MaxiMin) will simply not invest – as then the worst that can happen is their profits are unchanged.

This example highlights a dilemma in economic development. Development is an interdependent process with many types of industry and infrastructure needing to move forward together if they are to move forward at all. It's no use building a new factory if there are no roads or electricity supplies to serve it – just there is no use building a road or electricity plant if there are no consumers to use its services. This is why development often has to be of the 'Great Spurt' variety analysed by Alexander Gerschenkron and why the state can have a role to play: in coordinating investment, building infrastructure and providing guaranteed orders. This development model Game suggests that a free market might not be able to generate large-scale economic growth.

Conclusion

What Game Theory provides is a technique for analysing the rational behaviour of decision takers who are formulating strategies in situations where the outcomes of their decisions depend on the decisions taken by other parties. It generally turns out that, while the outcomes of such interactions would appear to be unpredictable, in reality we can predict how a Game will play out and what will be the likely equilibrium result of the Game. This result will not always be competitive. It will not always be zero-sum. It can often be the case – especially in repeated Games – that it is in the long-run interests of both players to cooperate: to stay silent in the Prisoner's Dilemma, or to mutually charge high prices in an oligopolistic pricing Game. But even in these cases, the logic of the situation will usually offer a temptation to one or other player to defect and try to steal a march on the other. Hence cooperation can always turn into conflict. Yet, whenever the players value their long-term payoffs, they will generally be led to cooperate again for this is in their mutual interest.

The predictions of Game Theory often appear rather disappointing. For one thing, its assumptions are unrealistic. It assumes players are rational, which studies show is often not the case. For example, even when cooperation is clearly in the interests in both players, there is a continual temptation to defect and cheat. In a well-known set of Game Theory experiments at the University of Ohio in the early 1960s players were given the choice of pushing a Black or Red button. The payoffs in terms of cents per Game were as follows.

		Player 2	
		Black	**Red**
Player 1	**Black**	**√4c, 4c√**	**√1c, 3c**
	Red	**3c, 1c√**	**0, 0**

Figure 36. The Payoff Matrix in the University of Ohio Cooperate or Defect Game Experiment

This Game has a clear Nash Equilibrium: both players should click Black and make 4 cents each per play. Indeed, for both players Red is a Dominated Strategy. Both players should click Black whatever the other player does. And yet when this Game was played, the players clicked Red 47 per cent of the time! Why? It seems the reason was both players couldn't resist trying to actually *win* a play. So, if Player 1 goes Black but Player 2 goes Red, then although Player 2 makes only 3c and not 4c, they will beat Player 1 by two cents (since Player 1 now only makes 1c). It seems that humans for some reason, biological or cultural, have an instinct to try and beat another player even when it is in their mutual best interests to cooperate.[1] Winning yields its own particular satisfaction, even if we damage our long-term interests to do so.

[1] Poundstone, *Prisoner's Dilemma*, pp. 173-177.

Game Theory is also unrealistic in assuming that the players can give specific values, not only to their own payoffs, but to their rivals. In some cases this is possible. Where the payoffs are money values, as in the above example, or in some zero-sum Games. But generally it is not. Attaching values to payoffs is usually a subjective exercise, so the scientific precision of the predictions arrived at is rather delusional.

Third, Game Theory rarely generates striking insights that challenge what we might regard as expected behaviour in a given situation. Its predicted outcomes are usually intuitive. The fact is we are all natural Game Theory players. And this is not surprising: we are all schooled from childhood to develop strategies for dealing with situations where we are interacting with small number of other persons to realise our goals subject to the other players seeking to realise their goals. From one perspective, family life involves little else (What time does the child go to bed? How much vegetables and ice-cream will the child eat? What gift does the child ask for as a present, and so on). All these are Game Theory situations, and school and work life remain replete with them too. We all develop an instinct of what to do when confronted by diverse payoffs and a range of strategies. It doesn't require Game Theory for a Prisoner to know it's in their interest to confess, or for a player of Chicken to know they had better swerve if their rival is a 'mad-man', and footballers really do kick left and right randomly in penalty shoot outs.

What Game Theory provides is an exceptionally rigorous approach to understanding strategic behaviour in such situations. It doesn't so much shape our behaviour, as allow us to better understand it. It simplifies and clarifies and provides a series of concepts the better to analyse choices – MaxiMin, Dominated Strategies, Nash Equilibria, Simultaneous versus Sequential Games, Tit-for-Tat, Strategic Moves, Credible Commitment and so forth. Like much of economics, it seeks to develop theories to explain what humans *actually* do, not shape what they *ought* to do. It is positive rather than normative. But, of course, this is what science seeks to do in general: Newton sought to understand gravity and trace its operation in more exact ways, not subvert it. The apple in his orchard fell to the ground well enough before Newton accounted for it. Indeed, economists used Game theoretic concepts even before Game Theory was invented. The Nash Equilibrium had already been developed by the likes of Cournot before John Nash produced his formal proof of the proposition. Yet, once having studied Game Theory, we become alive to just how ubiquitous Game situations are in our lives generally and in economic behaviour especially. And once confronted with such situations, it helps us to think logically and precisely about them, bringing new rigour to our study of oligopoly behaviour or the free-rider problem and public goods. For this reason it is an essential part of modern economic analysis, even if it did not revolutionise economic theory (or military strategy) as Von Neumann and Morgenstern initially hoped it would.