CHAPTER 9

Game theory

Learning objectives

This chapter covers the following topics:

- pure strategies, dominance and Nash equilibrium
- simultaneous games, including the prisoner's dilemma
- mixed strategies
- sequential games: threats, credibility and commitment
- repeated games

Key terms

Battle of the sexes
Burning bridges
Constant-sum game
Coordination problem
Duopoly
Experimental economics
First-mover advantage
Interdependence
Mixed strategy
Mixed Strategy Nash Equilibrium
Non-constant sum game
Oligopoly
Prisoner's dilemma
Pure Strategy Nash Equilibrium
Second-mover advantage
Sequential game
Simultaneous game
Strict dominance
Subgame Perfect Equilibrium
Weak dominance
Zero-sum game

9.1 Introduction

Game theory is an approach to decision-making under conditions of uncertainty or situations of conflict, developed by the mathematicians John von Neumann and Oskar Morgenstern (1944). Game theory has many applications throughout the social, behavioural and physical sciences; and accordingly, its remit is much wider than just economics. Nevertheless, its focus on uncertainty, interdependence, conflict and strategy makes it ideally suited to the analysis of decision-making in oligopoly. Game theory shows how situations can arise in which firms take decisions that may appear rational from each firm's individual perspective, but lead to outcomes that are suboptimal when assessed according to criteria reflecting the collective interest of all the firms combined.

A game is a situation in which two or more decision-makers, or players, face choices between a number of possible actions at any stage of the game. A game that is played only once is a single-period game. A player's **strategy** is a set of rules telling him which action to choose under each possible set of circumstances that might exist at any stage in the game. Each player aims to select the strategy (or mix of strategies) that will maximise his own **payoff**. The players face a situation of interdependence. Each player is aware that the actions of other players can affect his payoff, but at the time the player chooses his own action he may not know which actions are being chosen by the other players. In a **constant-sum game**, the sum of the payoffs to all players is always the same, whatever strategies are chosen. In a **non-constant-sum game**, the sum of the payoffs depends on the strategies chosen. A **zero-sum game** is a constant-sum game in which the sum of the gains and losses of all players is always zero. A game of poker is a zero-sum game: one player's winnings are exactly matched by the losses of rival players. The outcome of a game is the set of strategies and actions that are actually chosen, and the resulting payoffs. An equilibrium is a combination of strategies, actions and payoffs that is optimal (in some sense) for all players.

A game in which all players choose their actions simultaneously, before knowing the actions chosen by other players, is called a **simultaneous game**, and explored in the first three sections of this chapter. Section 9.2 introduces the concepts of dominance and Nash equilibrium. Dominant strategies are those that are either unambiguously superior to, or at least as good as, all other strategies, no matter which strategy the other player selects. Pure Strategy Nash Equilibrium (sometimes known simply as Nash equilibrium) is the game theory equivalent of the Cournot–Nash equilibrium in the model of duopoly, in which the two firms compete on quantities produced and sold, and both assume that the other will stick to its current production plans. In the terminology of game theory, no player wishes to depart from his current strategy if the other players continue to pursue their current strategies. Section 9.3 examines an important type of game, known as prisoner's dilemma, in which the pursuit of dominant strategies produces an outcome that is inferior (from the players' perspective) to the outcome that could be achieved if the players were to depart from their dominant strategies. Section 9.4 shows that in some games there is no dominant strategy for any player, and no Pure Strategy Nash Equilibrium. The concept of a mixed strategy is introduced, in which players choose their actions randomly in accordance with specific probabilities assigned to each action.

A game in which the players choose their actions in turn, so that a player who moves later knows the actions that were chosen by players who moved earlier, is called a **sequential game**, discussed in Section 9.5. In a sequential game, sometimes the player who moves first is able to gain an advantage by influencing the future direction of the game in their favour. In other cases, the player who moves second gains an advantage through observing their opponent's actions before making a decision. In arriving at a solution to a sequential game, a key consideration is the credibility of any retaliatory threats by one player to react in a particular manner to a hostile action on the part of another player. A retaliatory threat is credible only if the player concerned will still wish to execute the threat, if and when the time comes to do so. Players wishing to boost the credibility of a retaliatory threat may sometimes take steps that will alter the payoff structure in a way that closes off the option of backing down. Other players, recognizing that the threat is credible, may decide to hold off from taking the action that would trigger execution of the threat. In other words, a retaliatory threat that is credible might never be executed.

Finally, a game that is played more than once is called a multiple-period or repeated game. A multiple-period game can be repeated either indefinitely, or a finite number of times. Section 9.6 shows that in a repeated game with a prisoner's dilemma structure, the players may be able to learn from their experience to cooperate by departing from their dominant strategies, with adherence to cooperative behaviour reinforced by the threat of punishment in the event that cooperation breaks down.

In many ways, the property of **interdependence** is the key defining characteristic of a game, and it is this property that makes game theory relevant to an understanding of decision-making for firms in oligopoly. However, game theory has many applications other than decision-making under oligopoly, including strategy and tactics in sports, military strategy and nuclear deterrence. Although Chapter 9 discusses some non-economics applications, in most of the game theory examples discussed in this chapter and elsewhere, the players are two or more oligopolistic firms. Payoffs are usually defined in terms of the implications for the firms' profitability of the chosen strategies; and strategies are the decisions taken by the firms about matters such as price, output, advertising, product differentiation, research and development, entry or location.

9.2 Dominance and Nash equilibrium

Production game with strictly dominant strategies

As an initial game theory example, Figure 9.1 shows the payoff matrix for two firms, A and B, that have to decide simultaneously whether to produce low or high levels of output. Firm A's strategies are denoted **Low** and **High** and, similarly, firm B's strategies are denoted *Low* and *High*. The elements in the matrix represent the payoffs (for example profit) to the two firms. Both firms' payoffs

		Firm B's strategies	
		Low	*High*
Firm A's strategies	**Low**	**4** *4*	**2** *3*
	High	**3** *2*	**1** *1*

Figure 9.1 Production game with strictly dominant strategies

depend on their own output level and on the output level of the other firm, since market price is a function of the combined output levels of both firms. Within each cell of Figure 9.1, the first figure is A's payoff and the second figure is B's payoff. For example, if A selects **High** and B selects *Low*, A's payoff (profit) is **3** and B's payoff is *2*.

First, consider the choice between strategies **Low** and **High** from firm A's perspective. One method A could use to make this choice would be to examine which of **Low** and **High** is best for A if B selects *Low*, and which of **Low** and **High** is best for A if B selects *High*:

- If B selects *Low*, **Low** yields a payoff of **4** for A, while **High** yields a payoff of **3**. Therefore if B selects *Low*, A's best response is **Low**.
- If B selects *High*, **Low** yields a payoff of **2** for A, while **High** yields a payoff of **1**. Therefore if B selects *High*, A's best response is **Low**.

In this game, no matter what strategy B selects, it is best for A to choose **Low** rather than **High**. **Low** is said to be a **strictly dominant strategy**, because it is a superior strategy for A no matter what strategy B selects.

Second, consider the choice between strategies *Low* and *High* from B's perspective, using a similar approach:

- If A selects **Low**, *Low* yields a payoff of *4* for B, while *High* yields a payoff of *3*. Therefore, if A selects **Low**, B's best response is *Low*.
- If A selects **High**, *Low* yields a payoff of *2* for B, while *High* yields a payoff of *1*. Therefore, if A selects **High**, B's best response is *Low*.

Accordingly, no matter what strategy A selects, it is best for B to select *Low* rather than *High*. Therefore *Low* is B's strictly dominant strategy. **Strict dominance** refers to the case where a player can identify a strategy that is the best response for all strategies the other player may choose. Following this approach, it appears A should select **Low** and B should select *Low*, so that both firms earn a payoff of 4. This outcome, denoted (**Low**, *Low*), is known as a **Dominant Strategy Equilibrium**. In fact, the game shown in Figure 9.1 is rather trivial, in the sense that 4 is the best payoff achievable by either player under any circumstances. It seems natural that the players should choose the combination of strategies that produces this payoff for both of them. Below, it is shown that not all games are structured in a way that always produces such a pleasing outcome for the players!

Entry game with a weakly dominant strategy

Another dominance concept is **weak dominance**, which refers to the case where a player can identify a strategy that is at least as good as any other strategy, for all strategies the other player may choose, and better than any other strategy for at least one strategy the other player may choose. Figure 9.2 illustrates a game in which firm A has a **weakly dominant strategy**, while firm B has neither a strictly dominant strategy nor a weakly dominant strategy. Firm A is an incumbent and firm B is a potential entrant. In the simultaneous version of the entry game, Firm B chooses whether or not to enter, and Firm A plans whether to accommodate B's arrival in the event that B does enter, or initiate a price war. To create the capability to fight a price war, A will have to take certain steps that will be irreversible, before A knows B's decision. However, these steps will not impose any additional costs upon A in the event that B decides not to enter and A keeps the entire market to itself.

Consider A's choice between the strategies **Accommodate** and **Fight**:

- If B selects *No entry*, A does not need to execute the threatened price war. A's payoff of **3** is the same, regardless whether A had decided to **Accommodate** or **Fight**.
- If B selects *Entry*, **Accommodate** yields a payoff of **2** for A, while **Fight** yields a payoff of **1**. If B selects *Entry*, A's best response is **Accommodate**.

Firm A has no strictly dominant strategy, but **Accommodate** is weakly dominant: A prefers **Accommodate** if B selects *Entry*, but if B selects *No entry* A is indifferent between **Accommodate** and **Fight**.

Then, from B's perspective:

- If A selects **Accommodate**, *Entry* yields a payoff of *4* for B, while *No entry* yields a payoff of *3*. If A selects **Accommodate**, B's best response is *Entry*.
- If A decides to **Fight**, *Entry* yields a payoff of *1* for B, while *No entry* yields a payoff of *3*. If A selects **Fight**, B's best response is *No entry*.

Firm B therefore has neither any strictly dominant strategy, nor any weakly dominant strategy. In comparison to the game with strictly dominant strategies shown in Figure 9.1, the outcome of the game shown in Figure 9.2 is more difficult to predict. Firm A might opt for the weakly dominant strategy of **Accommodate**, in which case B prefers *Entry*. Alternatively, A might decide to **Fight**, in which case B prefers *No entry*. It is hard to be certain which way this game will turn out. We will return to this example several times during this chapter.

		Firm B's strategies	
		Entry	*No entry*
Firm A's strategies	**Accommodate**	**2**, *4*	**3**, *3*
	Fight	**1**, *1*	**3**, *3*

Figure 9.2 Entry game with a weakly dominant strategy

Nash Equilibrium

There is one further desirable and key property of the game shown in Figure 9.1. At the dominant strategy equilibrium (**Low**, *Low*), neither firm can improve its payoff given the current strategy of the other firm. Given that B selects *Low*, if A switches from **Low** to **High**, A's payoff falls from **4** to **3**. And given that A selects **Low**, if B switches from *Low* to *High*, B's payoff also falls from *4* to *3*. Therefore, if A selects **Low** and B selects *Low*, both firms perceive that they are maximizing their own profit, based on the assumption that the other firm's output is fixed at its current level. Both firms maximise profit subject to a zero conjectural variation assumption.

An equilibrium of this kind has been identified previously, in the discussion of the Cournot duopoly model. In Section 7.3, it is known as a Cournot–Nash equilibrium. In game theory, it is known as a **Pure Strategy Nash Equilibrium** (sometimes known simply as a Nash Equilibrium). In a Nash Equilibrium, neither firm can improve its payoff by switching to a different strategy, assuming the strategy chosen by the other firm does not change.

While there is no Dominant Strategy Equilibrium for the game shown in Figure 9.2, this game has two Pure Strategy Nash Equilibria. First (**Accommodate**, *Entry*) is a Nash Equilibrium:

- If B selects *Entry*, A's payoff would drop from **2** to **1** if A switches to **Fight;**
- If A selects **Accommodate**, B's payoff would drop from *4* to *3* if B switches to *No entry*.

Likewise (**Fight**, *No entry*) is a Nash Equilibrium:

- If B selects *No entry*, A's payoff would remain at **3** if A switches to **Accommodate**;
- If A selects **Fight**, B's payoff would drop from *3* to *1* if B switches to *Entry*.

Both of the solutions (**Accommodate**, *Entry*) and (**Fight**, *No entry*) satisfy the requirement for a Nash Equilbrium; that is, neither firm can improve its payoff, assuming the current strategy of the other firm is fixed.

Relationship between Dominant Strategy Equilibrium and Nash Equilibrium

In the game shown in Figure 9.1, there is an exact correspondence between the Dominant Strategy Equilibrium and the (Pure Strategy) Nash Equilibrium. It can be shown that a Dominant Strategy Equilibrium in any game is always a Nash Equilibrium. If both players select their strictly dominant strategies, it is impossible for either to improve its own payoff by changing its strategy, given the current strategy of the other player. In games such as the one shown in Figure 9.2, however, there may exist one or more Nash Equilibria, although there are no strictly dominant strategies and no Dominant Strategy Equilibrium.

In the game shown in Figure 9.3, firms A and B must decide simultaneously their advertising expenditures. They have a choice between three levels

		Firm B's budget		
		Low	*Medium*	*High*
Firm A's budget	**Low**	**40** *40*	**35** *45*	**10** *25*
	Medium	**35** *35*	**45** *30*	**15** *20*
	High	**30** *25*	**25** *15*	**20** *30*

Figure 9.3 Payoff matrix for the advertising budgets of firms A and B

of expenditure: low, medium or high. Both firms' payoffs from the advertising campaign depend on their own expenditure and on the expenditure of the other firm.

As before, consider firm A's choices:

- If B chooses *Low*, A's best response is **Low**.
- If B chooses *Medium*, A's best response is **Medium**.
- If B chooses *High*, A's best response is **High**.

Similarly, consider firm B's choices:

- If A chooses **Low**, B's best response is *Medium*.
- If A chooses **Medium**, B's best response is *Low*.
- If A chooses **High**, B's best response is *High*.

There are no strictly dominant strategies, and no weakly dominant strategies, for either firm A or firm B. By inspection, however, it can be confirmed that (**High**, *High*) is a Nash Equilibrium. If B chooses *High*, then **High** is also A's best response; and if A chooses **High**, then *High* is also B's best response. Unfortunately, in the absence of strictly dominant strategies, there is no simple decision-making procedure that will enable the two firms to reach the Nash Equilibrium easily. If this solution is achieved by some means, however, it is stable in the sense that there is no incentive for either firm to depart from it, given the zero conjectural variation assumption.

It is important to notice that firms A and B could both be better off by cooperating or agreeing to choose (**Low**, *Low*) in Figure 9.3, rather than remaining at the Nash Equilibrium of (**High**, *High*). In contrast to the Nash Equilibrium, however, this cooperative solution is unstable. If A chooses **Low**, B has an incentive to 'cheat' and choose *Medium* instead of *Low*. But if B chooses *Medium*, A would also prefer **Medium**; and then if A chooses **Medium**, B would prefer *Low;* and so on. The cooperative solution is vulnerable to defection by one or both of the firms, and is likely to break down.

9.3 The prisoner's dilemma game

Production game with prisoner's dilemma structure

Figure 9.4 presents another simultaneous game, with a structure similar to Figure 9.1, but a different set of payoffs. Applying the same reasoning as before, from A's perspective:

- If B selects *Low*, **Low** yields a payoff of **3** for A, while **High** yields a payoff of **4**. If B selects *Low*, A's best response is **High**.
- If B selects *High*, **Low** yields a payoff of **1** for A, while **High** yields a payoff of **2**. If B selects *High*, A's best response is **High**.

And from B's perspective:

- If A selects **Low**, *Low* yields a payoff of *3* for B, while *High* yields a payoff of *4*. If A selects **Low**, B's best response is *High*.
- If A selects **High**, *Low* yields a payoff of *1* for B, while *High* yields a payoff of *2*. If A selects **High**, B's best response is *High*.

High is a strictly dominant strategy for A and *High* is a strictly dominant strategy for B. Accordingly, it seems that A should select **High** and B should select *High*, in which case both firms earn a payoff of 2. As before, the Dominant Strategy Equilibrium (**High**, *High*) is also a Nash Equilibrium. Given that B selects *High*, if A switches from **High** to **Low**, A's payoff falls from **2** to **1;** and given that A selects **High**, if B switches from *High* to *Low*, B's payoff also falls from *2* to *1*. However, this time something appears to be wrong. If both firms had selected the *other* strategy (**Low**, *Low*), either by cooperating or perhaps by acting independently, both firms would have earned a superior payoff of 3 each, rather than their actual payoff of 2 each.

Figure 9.5 is an example of a special class of single period non-constant-sum game, known as the **prisoner's dilemma**. In a prisoner's dilemma game, there are strictly dominant strategies for both players that produce a combined payoff that is worse than the combined payoff the players could achieve if they cooperate, with each player agreeing to choose a strategy other than his strictly dominant

		Firm B's strategies	
		Low	*High*
Firm A's strategies	**Low**	**3** *3*	**1** *4*
	High	**4** *1*	**2** *2*

Figure 9.4 Payoff matrix for firms A and B: prisoner's dilemma example

		Brian's strategies	
		Not confess	*Confess*
Alan's strategies	**Not confess**	**–2** *–2*	**–10** *–1*
	Confess	**–1** *–10*	**–5** *–5*

Figure 9.5 Payoff matrix for Alan and Brian: original prisoner's dilemma

strategy. In other words, in a prisoner's dilemma, gains can be made by both players if they cooperate or collude.

The original prisoner's dilemma

To see why this type of game is called prisoner's dilemma, consider a situation where the police hold two prisoners, Alan and Brian, who are suspected of having committed a serious crime together. The police have insufficient evidence to secure a conviction unless one or both prisoners confesses. The prisoners are separated physically and there is no communication between them. Each is told the following:

- If you both confess to the serious crime, you both receive a reduced punishment of five years in prison.
- If neither of you confesses to the serious crime, you are both convicted of a minor crime and you both receive the full sentence for the minor crime of two years in prison.
- If you confess to the serious crime and your fellow prisoner does not confess, you receive a reduced sentence of one year in prison for the minor crime (and your punishment for the serious crime is cancelled).
- If you do not confess to the serious crime and your fellow prisoner confesses, you receive the full sentence for the serious crime of ten years in prison.

The payoff matrix is shown in Figure 9.5, with all payoffs shown as negative numbers, because in this case a large payoff (prison sentence) is bad, not good. Alan's reasoning might be as follows: if Brian confesses, I should confess because five years is better than ten years; and if Brian does not confess, I should confess because one year is better than two years. Therefore I will confess. Brian's reasoning is the same, because the payoffs are symmetric between the two prisoners. Therefore both confess, and both receive sentences of five years. But if they had been able to cooperate, they could have agreed not to confess and both would have received sentences of two years. Even acting independently, they might be able to reach the cooperative solution. Alan knows that if he does not confess, he receives a two-year sentence as long as Brian does the same. However, Alan is worried because he knows there is a big incentive for Brian to 'cheat' on Alan by confessing. By doing so, Brian can earn the one-year sentence and leave Alan with a ten-year sentence!

Brian is in a similar position: if he does not confess, he receives the two-year sentence as long as Alan also does not confess. However, Brian also knows there is a big incentive for Alan to cheat. The cooperative solution might be achievable, especially if Alan and Brian can trust one another not to cheat, but it is also unstable and liable to break down.

The prisoner's dilemma and the Cournot duopoly model

Section 7.3 analysed the choices of output levels by two duopolists. Comparing the Cournot–Nash and the Chamberlin solutions to the duopoly model shown in Figure 7.9, it is apparent that if the two firms operate independently according to the zero conjectural variation assumption, and each firm produces a relatively high output level of 1/3, the Cournot–Nash equilibrium is attained. In the terminology of the present section, this is a non-cooperative outcome. If, on the other hand, the two firms recognise their interdependence and aim for joint profit maximization, and each firm produces the lower output level of 1/4, the Chamberlin equilibrium is attained. In present terminology, this is the cooperative outcome.

Figures 9.6 and 9.7 show that if the two duopolists have to make their output decisions simultaneously, without knowing the other firm's decision, effectively

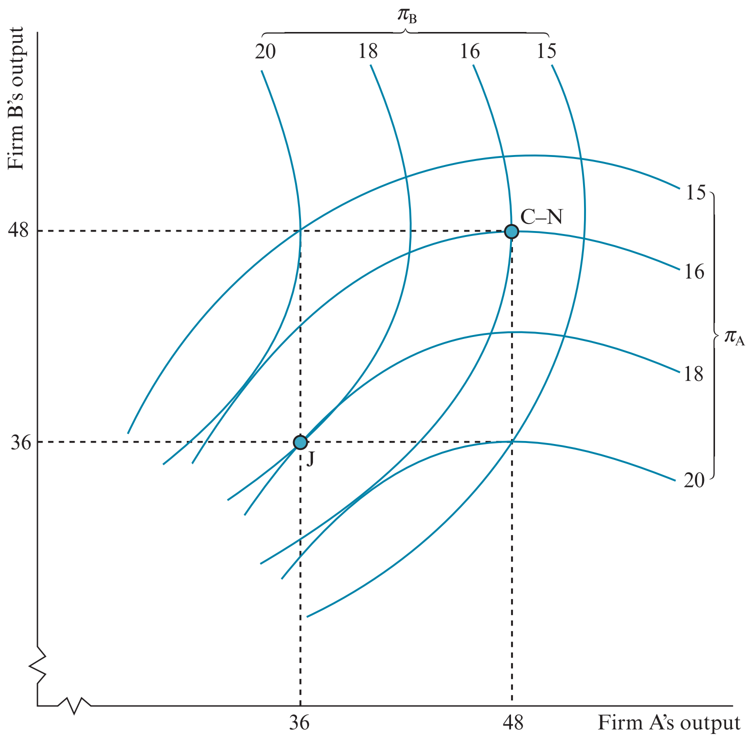

Figure 9.6 Isoprofit curves for firms A and B: Cournot–Nash versus Chamberlin's prisoner's dilemma

		Firm B's strategies	
		Low	*High*
Firm A's strategies	**Low**	**18** *18*	**15** *20*
	High	**20** *15*	**16** *16*

Figure 9.7 Payoff matrix for firms A and B: Cournot–Nash versus Chamberlin's prisoner's dilemma

they play a prisoner's dilemma game. The assumptions underlying Figures 9.6 and 9.7 are the same as in the original Cournot model developed in Section 7.3, with one exception. The two duopolists are assumed to produce an identical product and incur zero marginal costs. The one change involves a rescaling of the quantity axis for the market demand function, so that the maximum quantity that could be sold if the price falls to zero is 144 units (rather than one unit). As before, the price axis for market demand function is on a scale of $P = 0$ to $P = 1$, so when $P = 0$, $Q = 144$ and when $P = 1$, $Q = 0$. (Rescaling the quantity axis avoids the occurrence of fractional prices, quantities and profits.) You can verify that the prices, quantities and profits or payoffs shown in Figures 9.6 and 9.7 are equivalent to their counterparts in Figure 7.9 multiplied by a factor of 144.

In Figure 9.6, it is assumed that each firm has to choose between producing a high output of 48 units, or a low output of 36 units. If both firms select high, the Cournot–Nash equilibrium is attained, and both firms' profits are 16. If both firms select low, the Chamberlin joint profit-maximization equilibrium is attained, and both firms' profits are 18. If one firm selects low while the other selects high, the low-producing firm suffers and earns 15, while the high-producing firm prospers and earns 20. Figure 9.7 represents these outcomes in the form of a payoff matrix. Applying the same reasoning as before, from A's perspective:

- If B selects *Low*, **Low** yields a payoff of **18** for A, while **High** yields a payoff of **20**. If B selects *Low*, A's best response is **High**.
- If B selects *High*, **Low** yields a payoff of **15** for A, while **High** yields a payoff of **16**. If B selects *High*, A's best response is **High**.

Accordingly, it is best for A to select **High**, no matter what strategy B selects. The same is also true for B, because the two firms are identical. (**High**, *High*) is the Dominant Strategy Equilibrium, and is also a suboptimal non-cooperative Cournot–Nash outcome. As before, the cooperative or collusive outcome (**Low**, *Low*) might be achievable if the firms can trust each other to stick to the low output strategy and not defect and produce high output. This outcome is unstable, however, and is liable to break down. For the cooperative solution to hold in an oligopoly, any agreement between the firms might have to be accompanied by an enforceable contract (legal or otherwise).

Conflict *versus* cooperation

Not all prisoner's dilemma games generate suboptimal outcomes, especially when the assumptions are relaxed. First, the optimal (cooperative) outcome might be achieved if there is good communication between the players. If firms meet frequently, they can exchange information and monitor each other's actions. If the two prisoners, Alan and Brian, were not segregated, they could determine their best strategies by a continual examination of their options. The nuclear deterrence 'game' played by the United States and the Soviet Union in the 1960s and 1970s was likened to a prisoner's dilemma game. The choices were whether to attack the rival with a pre-emptive strike, or abide by the 'non-first use' agreement. Perhaps one reason why the optimal outcome (sticking to the agreement) was achieved was that the installation of a telephone hotline between Washington and Moscow permitted rapid communication and exchange of information at the highest levels of government. Alternatively, it might be possible to achieve a cooperative outcome if the players are able to recognise trustworthiness in other players through visual signals (Janssen, 2008).

Second, in practice an important characteristic of any game is the length of the reaction lag: the time it takes for a player who has been deceived to retaliate. The longer the reaction lags, the greater the temptation for either player to act as an aggressor. If Brian cheats on Alan, Alan may have to wait ten years to take revenge, unless he has friends outside the prison who are prepared to act more quickly. In cartels, the main deterrent to cheating is immediate discovery and punishment. In the nuclear deterrence game, short reaction lags were crucial to ensuring both sides kept to the agreement. Each side boasted that it could retaliate within minutes if attacked by the other, ensuring there was no first-mover advantage. This policy became known as mutually assured destruction (MAD).

Third, the dynamics of rivalry may also be relevant. Is the rivalry continuous, or 'one-off'? If rivalry is continuous in a repeated game, players learn over time that cooperation is preferable to aggression. Professional criminals have no problem with the prisoner's dilemma: experience has taught them that silence is the best option. In an oligopoly, firms change prices, alter product lines and determine advertising strategies, continuously. The firms may learn over time that aggressive behaviour leads to hostile (tit-for-tat) reactions from rivals, which tend to cancel out any short-term gains (see Case Study 9.1). Repeated or multiple-period games are examined in more detail below.

9.4 Mixed strategies

In some games, there is neither any Dominant Strategy Equilibrium nor any Pure Strategy Nash Equilibrium. In others there may be no Dominant Strategy Equilibrium, but more than one Pure Strategy Nash Equilibrium. In such cases, it may be beneficial for the firms (or other players) to adopt what are known as mixed strategies. A mixed strategy involves randomizing the choice between two or more options, with probabilities defined for each option.

Case study 9.1

Prison breakthrough

Economist, 20th August 2016

JOHN NASH arrived at Princeton University in 1948 to start his PhD with a one-sentence recommendation: "He is a mathematical genius". He did not disappoint. Aged 19 and with just one undergraduate economics course to his name, in his first 14 months as a graduate he produced the work that would end up, in 1994, winning him a Nobel prize in economics for his contribution to game theory.

On November 16th 1949, Nash sent a note barely longer than a page to the Proceedings of the National Academy of Sciences, in which he laid out the concept that has since become known as the "Nash equilibrium". This concept describes a stable outcome that results from people or institutions making rational choices based on what they think others will do. In a Nash equilibrium, no one is able to improve their own situation by changing strategy: each person is doing as well as they possibly can, even if that does not mean the optimal outcome for society. With a flourish of elegant mathematics, Nash showed that every "game" with a finite number of players, each with a finite number of options to choose from, would have at least one such equilibrium.

His insights expanded the scope of economics. In perfectly competitive markets, where there are no barriers to entry and everyone's products are identical, no individual buyer or seller can influence the market: none need pay close attention to what the others are up to. But most markets are not like this: the decisions of rivals and customers matter. From auctions to labour markets, the Nash equilibrium gave the dismal science a way to make real-world predictions based on information about each person's incentives.

One example in particular has come to symbolise the equilibrium: the prisoner's dilemma. Nash used algebra and numbers to set out this situation in an expanded paper published in 1951, but the version familiar to economics students is altogether more gripping. (Nash's thesis adviser, Albert Tucker, came up with it for a talk he gave to a group of psychologists.)

It involves two mobsters sweating in separate prison cells, each contemplating the same deal offered by the district attorney. If they both confess to a bloody murder, they each face ten years in jail. If one stays quiet while the other snitches, then the snitch will get a reward, while the other will face a lifetime in jail. And if both hold their tongue, then they each face a minor charge, and only a year in the clink. There is only one Nash-equilibrium solution to the prisoner's dilemma: both confess. Each is a best response to the other's strategy; since the other might have spilled the beans, snitching avoids a lifetime in jail. The tragedy is that if only they could work out some way of co-ordinating, they could both make themselves better off.

The example illustrates that crowds can be foolish as well as wise; what is best for the individual can be disastrous for the group. This tragic outcome is all too common in the real world. Left freely to plunder the sea, individuals will fish more than is best for the group, depleting fish stocks. Employees competing to impress their boss by staying longest in the office will encourage workforce exhaustion. Banks have an incentive to lend more rather than sit things out when house prices shoot up.

The Nash equilibrium would not have attained its current status without some refinements on the original idea. First, in plenty of situations, there is more than one possible Nash equilibrium. Drivers choose which side of the road to drive on as a best response to the behaviour of other drivers—with very different outcomes, depending on where they live; they stick to the left-hand side of the road in Britain, but to the right in America. Much to the disappointment of algebra-toting economists, understanding strategy requires knowledge of social norms and habits. Nash's theorem alone was not enough.

A second refinement involved accounting properly for non-credible threats. If a teenager threatens to run away from home if his mother separates him from his mobile phone, then there is a Nash equilibrium where she gives him the phone to retain peace of mind. But Reinhard Selten, a German economist who shared the 1994 Nobel prize with Nash and John Harsanyi, argued that this is not a plausible outcome. The mother should know that her child's threat is empty—no matter how tragic the loss of a phone would be, a night out on the streets would be worse. She should just confiscate the phone, forcing her son to focus on his homework.

Mr Selten's work let economists whittle down the number of possible Nash equilibria. Harsanyi addressed the fact that in many real-life games, people are unsure of what their opponent wants. Economists would struggle to analyse the best strategies for two lovebirds trying to pick a mutually acceptable location for a date with no idea of what the other prefers. By embedding each person's beliefs into the game (for example that they correctly think the other likes pizza just as much as sushi), Harsanyi made the problem solvable. A different problem continued to lurk. The predictive power of the Nash equilibrium relies on rational behaviour. Yet humans often fall short of this ideal. In experiments replicating the set-up of the prisoner's dilemma, only around half of people chose to confess. For the economists who had been busy embedding rationality (and Nash) into their models, this was problematic. What is the use of setting up good incentives, if people do not follow their own best interests?

All was not lost. The experiments also showed that experience made players wiser; by the tenth round only around 10% of players were refusing to confess. That taught economists to be more cautious about applying Nash's equilibrium. With complicated games, or ones where they do not have a chance to learn from mistakes, his insights may not work as well.

The Nash equilibrium nonetheless boasts a central role in modern microeconomics. Nash died in a car crash in 2015; by then his mental health had recovered, he had resumed teaching at Princeton and he had received that joint Nobel—in recognition that the interactions of the group contributed more than any individual.

Abridged
1051 words

http://www.economist.com/news/economics-brief/21705308-fifth-our-series-seminal-economic-ideas-looks-nash-equilibrium-prison

Advertising game with no Pure Strategy Nash Equilibrium

Consider first the case of two firms that need to decide simultaneously their advertising budgets (low or high). As before, both firms' payoffs from the advertising campaign depend both on their own expenditure and on the other firm's expenditure. The payoff matrix is shown in Figure 9.8. This is a constant-sum game. Whatever combination of strategies is chosen, the sum of the payoffs to both firms is 5. There is no strictly dominant strategy for either firm. From firm A's perspective:

- If B chooses *Low*, A's best response is **High**.
- If B chooses *High*, A's best response is **Low**.

		Firm B's strategies	
		Low	*High*
Firm A's strategies	**Low**	**1** *4*	**3** *2*
	High	**4** *1*	**0** *5*

Figure 9.8 Advertising game with no Pure Strategy Nash Equilibrium

And from B's perspective:

- If A chooses **Low**, B's best response is *Low*.
- If A chooses **High**, B's best response is *High*.

There is also no Pure Strategy Nash Equilibrium, because there is no pair of strategies from which neither firm would wish to defect if the other firm continues to follow the same strategy. A is in a difficult position. If A selects **Low**, B might select *Low* and A only earns a profit of 1. But, on the other hand, if A selects **High** and B selects *High*, A earns a profit of 0. Of course, B also faces a similar dilemma.

A possible solution lies in the concept of a **mixed strategy**, developed by von Neumann and Morgenstern (1944). A player follows a mixed strategy by choosing his action randomly, assigning fixed probabilities to the selection of each action. In contrast, previous examples have resulted in the choice of a **pure strategy** by both players. According to the non-cooperative solution to the prisoner's dilemma game shown in Figure 9.4, for example, A should only ever select **High** and B should only ever select *High*, because **High** and *High* are strictly dominant strategies.

At the Mixed Strategy Nash Equilibrium, each player has the same expected payoff from either action and from a mixed strategy that assigns specific probabilities to both actions; and this expected payoff is unaffected by the mixed strategy selected by the other player. Therefore neither player has any incentive to depart from his current mixed strategy, assuming the other player continues with *his* current mixed strategy.

Returning to Figure 9.8, suppose firm A assigns a probability of x to the choice of **Low**, and a probability of $(1 - x)$ to the choice of **High**. B's expected payoffs (in terms of x) are as follows:

- If B chooses *Low*, B's possible payoffs are *4* (if A chooses **Low**, with a probability of x) and *1* (if A chooses **High**, with a probability of $1 - x$). B's expected payoff is $4x + 1(1 - x) = 1 + 3x$.
- If B chooses *High*, B's possible payoffs are *2* (if A chooses **Low**, with a probability of x) and *5* (if A chooses **High**, with a probability of $1 - x$). B's expected payoff is $2x + 5(1 - x) = 5 - 3x$.

The right-hand diagram in Figure 9.9 plots firm B's expected payoffs against all possible values of x, for each of the two possible choices open to B. Setting $x = 0$ is equivalent to 'A always chooses **High**'. In this case, the best B can achieve is a payoff of 5 (if B chooses *High*). Similarly, setting $x = 1$ is equivalent to 'A always

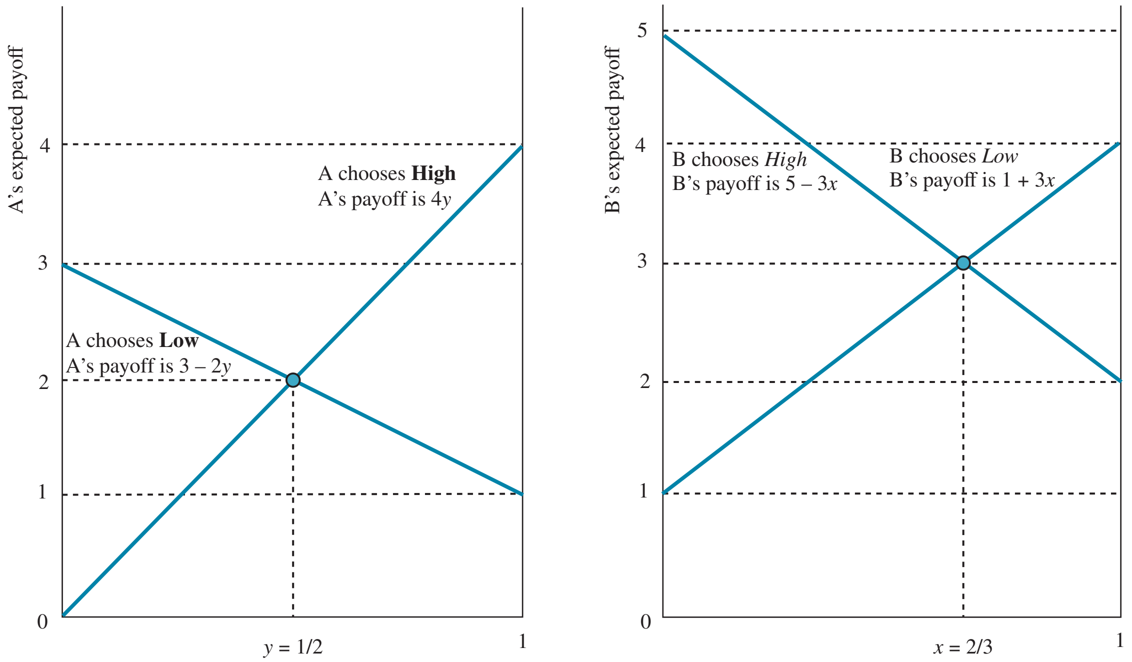

Figure 9.9 Expected payoffs for firms A and B in Mixed Strategy Nash Equilibrium

chooses **Low**'. In this case, the best B can achieve is a payoff of 4 (if B chooses *Low*). A should choose x to ensure B is indifferent between choosing *Low* or *High.* A does so by selecting a mixed strategy of $x = 2/3$, with the consequence that B's expected payoff of 3 is the same, regardless whether B selects *Low* or *High*, or any mixed strategy that combines *Low* and *High*.

In this case, A earns an expected profit of 2, whichever of *Low* and *High* is chosen by B. In fact, it can be shown that A still earns an expected profit of 2 if B selects any mixed strategy which involves choosing randomly between *Low* and *High*, no matter what probabilities B assigns to these two choices.

B's optimal mixed strategy can also be evaluated with reference to Figure 9.9. Let B assign a probability of y to the choice of *Low*, and a probability of $(1 - y)$ to the choice of *High*. A's expected payoffs (in terms of y) are as follows:

- If A chooses **Low**, A's possible payoffs are **1** (if B chooses *Low*, with a probability of y) and **3** (if B chooses *High*, with a probability of $1 - y$). A's expected payoff is $1y + 3(1 - y) = 3 - 2y$.
- If A chooses **High**, A's possible payoffs are **4** (if B chooses *Low*, with a probability of y) and **0** (if B chooses *High*, with a probability of $1 - y$). A's expected payoff is $4y$.

The left-hand diagram in Figure 9.9 plots A's expected payoffs against all possible values of y, for each of the two possible choices available to A. Setting $y = 1$ is equivalent to 'B always chooses *Low*'. In this case, the best A can achieve is a payoff of **4** (if A chooses **High**). Similarly, setting $y = 0$ is equivalent to 'B always chooses *High*'. In this case, the best A can achieve is a payoff of 3 (if A chooses **Low**). B should choose y to ensure A is indifferent between choosing **Low** or **High**. B does so by selecting a mixed strategy of $y = 1/2$, with the consequence that A's expected payoff of 2 is the same, regardless whether A selects **Low** or **High**, or any mixed strategy that combines **Low** and **High**.

If A sets $x = 2/3$ and B sets $y = 1/2$, the game shown in Figure 9.8 achieves a **Mixed Strategy Nash Equilibrium**. Each firm selects the probabilities that maximise its own expected payoff, given the mixed strategy that is being employed by the other firm. In fact, by selecting the probabilities in this way, each firm guarantees its own expected payoff, whatever the probabilities selected by the other firm. Selecting $x = 2/3$ guarantees A an expected payoff of 2 for any value of y selected by B; selecting $x = 2/3$ makes A indifferent to B's selection of probabilities. Likewise, selecting $y = 1/2$ guarantees B an expected payoff of 3 for any value of x selected by A; selecting $y = 1/2$ makes B indifferent to A's selection of probabilities. Although the mathematics is beyond the scope of this text, it has been shown that for any game with a fixed number of players, each of whom chooses between a fixed number of possible actions, a Nash equilibrium involving either pure strategies or mixed strategies always exists.

Entry game with a weakly dominant strategy

The concept of Mixed Strategy Nash Equilibrium is also relevant to the entry game with a weakly dominant strategy for firm A, discussed in Section 9.2 and summarised in Figure 9.2. We can search for a Mixed Strategy Nash Equilibrium using the same reasoning as above. Suppose firm A assigns a probability of x to the choice of **Accommodate**, and a probability of $(1 - x)$ to the choice of **Fight**. B's expected payoffs (in terms of x) are as follows:

- If B selects *Entry*, B's possible payoffs are *4* (if A chooses **Accommodate**, with a probability of x) and *1* (if A chooses **Fight**, with a probability of $1 - x$). B's expected payoff is $4x + 1(1 - x) = 1 + 3x$.
- If B chooses *No entry*, B's possible payoffs are *3* (if A chooses **Accommodate**) and *3* (if A chooses **Fight**). B's expected payoff is $3x + 3(1 - x) = 3$.

Since **Accommodate** and **Fight** are both best responses for A if B selects *No entry*, A is indifferent between any mixed strategy for any (feasible) value of x between 0 and 1 in the event that B selects *No entry*. To identify the Mixed Strategy Nash Equilibrium, we need to find the range of mixed strategies for A, or the range of values for x, that make *No entry* the best response for B. In other words, we need to find x such that $3 \geq 1 + 3x$. Accordingly, *No entry* is the best response for B for any $x \leq 2/3$. Any outcome in which A follows a mixed strategy with $x \leq 2/3$ and B follows a pure strategy of *No entry* is a Mixed Strategy Nash Equilibrium. The requirement $x \leq 2/3$ incorporates the Pure Strategy Nash Equilibrium (**Fight**, *No entry*), at which $x = 0$. The entry game with a weakly dominant strategy for firm A turns out to have not only two Pure Strategy Nash Equilibria, (**Accommodate**, *Entry*) and (**Fight**, *No entry*), but also an infinite number of Mixed Strategy Nash Equilibria.

Battle of the sexes game

Consider next the game represented by the payoff matrix shown in Figure 9.10, known as the **battle of the sexes** game. Arthur and Barbara are a couple with widely divergent preferences for an evening's live entertainment: Arthur prefers to watch

		Barbara's strategies	
		Football	*Ballet*
Arthur's strategies	**Football**	**4**, *2*	**1**, *1*
	Ballet	**1**, *1*	**2**, *4*

Figure 9.10 Battle of the sexes game

football, and Barbara prefers ballet. However, despite their differences in taste they both prefer each other's company to attending either event alone, and both will go straight home if the other does not show up. Owing to a temporary mobile phone outage, they are unable to communicate with each other prior to the start of either event. Both must decide which event to attend, without knowing the other's decision.

By inspection of Figure 9.10, it is clear that there are no strictly dominant strategies, but both (**Football**, *Football*) and (**Ballet**, *Ballet*) are Pure Strategy Nash Equilibria. In respect of (**Football**, *Football*), if Barbara selects *Football*, Arthur would not wish to switch from **Football** (because his payoff would drop from **4** to **1**); and if Arthur selects **Football**, Barbara would not wish to switch from *Football* (because her payoff would drop from *2* to *1*). The same reasoning applies to (**Ballet**, *Ballet*). But without communication, how can the couple be sure of achieving either of these two solutions?

The selection of mixed strategies by both partners might resolve the dilemma, but as we shall see, the Mixed Strategy Nash Equilibrium in this game is somewhat flawed. Let Arthur assign a probability of x to the choice of **Football**, and a probability of $(1 - x)$ to the choice of **Ballet**. Barbara's expected payoffs (in terms of x) are as follows:

- If Barbara chooses *Football*, Barbara's possible payoffs are *2* (with a probability of x) and *1* (with a probability of $1 - x$). Barbara's expected payoff is $2x + 1(1 - x) = 1 + x$.
- If Barbara chooses *Ballet*, Barbara's possible payoffs are *1* (with a probability of x) and *4* (if with a probability of $1 - x$). Barbara's expected payoff is $1x + 4(1 - x) = 4 - 3x$.

Let Barbara assign a probability of y to the choice of *Football*, and a probability of $(1 - y)$ to the choice of *Ballet*. Arthur's expected payoffs (in terms of y) are as follows:

- If Arthur chooses **Football**, Arthur's possible payoffs are **4** (with a probability of y) and **1** (with a probability of $1 - y$). Arthur's expected payoff is $4y + 1(1 - y) = 1 + 3y$.
- If Arthur chooses **Ballet**, Arthur's possible payoffs are **1** (with a probability of y) and **2** (with a probability of $1 - y$). Arthur's expected payoff is $1y + 2(1 - y) = 2 - y$.

Proceeding in the same way as before, Arthur should choose x so as to minimise Barbara's expected payoff. Arthur does so by selecting x such that Barbara's expected payoff is the same regardless whether Barbara selects *Football* or *Ballet*, or any mixed strategy that combines *Football* and *Ballet*. Solving $1 + x = 4 - 3x$ yields $x = 3/4$: Arthur should select **Football** with a probability of $x = 3/4$, and **Ballet** with a probability of $(1 - x) = 1/4$. Since the payoff matrix is symmetric, it is straightforward to verify that Barbara should select *Football* with a probability of $y = 1/4$ and *Ballet* with a probability of $(1 - y) = 3/4$.

A troublesome feature of the Mixed Strategy Nash Equilibrium in the Battle of the sexes game is that the expected payoffs to both players are lower than the payoffs at either of the two Pure Strategy Nash Equilibria. Arthur's expected payoff is calculated as follows:

P(Arthur chooses **Football**) × P(Barbara chooses *Football*) × 4
+ P(Arthur chooses **Football**) × P(Barbara chooses *Ballet*) × 1
+ P(Arthur chooses **Ballet**) × P(Barbara chooses *Football*) × 1
+ P(Arthur chooses **Ballet**) × P(Barbara chooses *Ballet*) × 2

$$= 4xy + x(1 - y) + (1 - x)y + 2(1 - x)(1 - y)$$
$$= 1.75 \text{ when } x = 3/4 \text{ and } y = 1/4.$$

Likewise Barbara's expected payoff is 1.75; so the expected payoffs at the Mixed Strategy Nash Equilibrium are (**1.75**, *1.75*). Both players would be better off at either of the two Pure Strategy Nash Equilibria, with payoffs of (**4**, *2*) at (**Football**, *Football*) or (**2**, *4*) at (**Ballet**, *Ballet*). As we have seen, however, the players encounter what is known as a **coordination problem** in reaching either of these solutions.

It has been suggested that the battle of the sexes game is relevant in describing the situation faced by two firms in deciding which of two alternative technological standards to adopt. Each firm has a competitive advantage with a different standard; but if the two firms fail to adopt the same standard, customers refuse to buy from either firm, and both firms are worse off than they would be if they made either standard the common standard. Substituting firms A and B for Arthur and Barbara, and standards 1 and 2 for Football and Ballet, the two firms face payoffs with the same relative magnitudes as those shown in Figure 9.9. The firms also face the same coordination problem: how do they achieve a consensus as to which standard to adopt? Some form of regulatory intervention, to enforce a common standard, might be justified, and even welcomed by the firms affected, under these circumstances.

9.5 Sequential games

In the games examined so far in Sections 9.2 to 9.4, the players act simultaneously and decide their strategies and actions before they know which strategies and actions have been chosen by their rivals. However, there are other games in which the players' decisions follow a sequence. One player makes a decision, and the other player observes this decision before making a response. For example, firm A decides to launch a new brand and firm B then decides how best to respond. Should B imitate A and launch a brand with identical characteristics, or should B aim for a segment in the market that is not serviced by A and launch a brand with different characteristics? For a **sequential game**, it is convenient to map the choices facing the players in the form of a game tree.

Firm B's strategies

Firm A's strategies	*Crunchy*	*Fruity*
Crunchy	**3** *3*	**5** *4*
Fruity	**4** *5*	**2** *2*

Figure 9.11 Breakfast cereals game: strategic form representation

Breakfast cereals game with first-mover advantage

Suppose two breakfast cereal producers are both considering a new product launch. They each have a choice of launching one of two products: one product's appeal is 'crunchiness' and the other's appeal is 'fruitiness'. Assume the crunchy cereal is more popular with consumers than the fruity cereal. Figure 9.11 shows the payoff matrix in the same form as before, assuming both firms move simultaneously, ignorant of what their rival is planning. The payoff structure is similar to the battle of the sexes game, with the exception that making different choices, rather than the same choices, is preferred by both firms. There are no strictly dominant strategies: if B produces *Crunchy* it is better for A to produce **Fruity**, but if B produces *Fruity* it is better for A to produce **Crunchy**. However, (**Fruity**, *Crunchy*) and (**Crunchy**, *Fruity*) are both Pure Strategy Nash Equilibria. Using the methods discussed in Section 9.4, you can verify that the Mixed Strategy Nash Equilibrium requires both firms to choose their actions randomly, with probabilities of 3/4 assigned to crunchy and 1/4 assigned to fruity.

In a sequential game, however, if A is the first to launch its new product and B then responds after having observed A's action, the outcome is different. Figure 9.12 shows the game tree representation of the payoffs of the breakfast cereal game, also known as the *extensive form representation*. (The equivalent

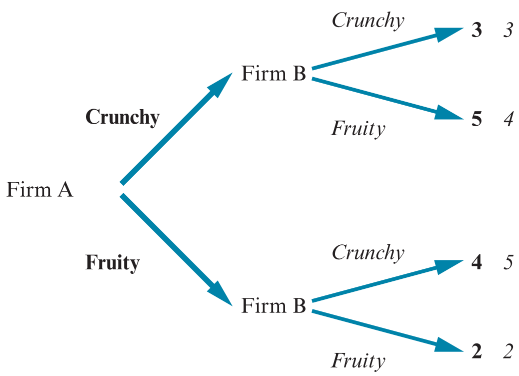

Figure 9.12 Breakfast cereals game: extensive form representation

terminology for the payoff matrix used previously is the *strategic form representation*.) Consider A's decision:

- If A produces **Crunchy**, B's *Fruity* payoff of *4* exceeds B's *Crunchy* payoff of *3*, so B will produce *Fruity* and A earns a payoff of **5**.
- If A produces **Fruity**, B's *Crunchy* payoff of *5* exceeds B's *Fruity* payoff of *2*, so B will produce *Crunchy* and A earns a payoff of **4**.

A realises that whatever product A launches, the rational response of B is to launch the alternative product. A's best action is to produce **Crunchy**, and A earns the higher payoff of **5**. B produces *Fruity* and earns the lower payoff of *4*. At (**Crunchy**, *Fruity*), A ends up with the higher payoff, because A benefits from a first-mover advantage. In many sequential games, the player who moves first gains an advantage, by influencing the shape of the game and forcing the other player to react to the first-mover's decision, rather than act in a way that is independent of the first-mover's presence. In the breakfast cereals game, from the symmetry of the payoff matrix it is obvious that if B were the first mover, the solution would be (**Fruity**, *Crunchy*), A would end up with the lower payoff of **4**, and B would end up with the higher payoff of *5*.

Assuming the first mover is A, can B take any steps that might deliver the other outcome, in which the firms produce different products (**Fruity**, *Crunchy*) yielding payoffs of **4** to A and *5* to B? Perhaps, prior to A's action, B could threaten to produce *Crunchy* regardless of A's decision. If A views B's threat as credible, A will calculate that by producing **Crunchy**, A will end up with a payoff of **3**, but by producing **Fruity** A will achieve a payoff of **4**. According to this calculation, B's threat should steer A towards the (**Fruity**, *Crunchy*) outcome that B prefers. However, is B's threat to produce *Crunchy* truly credible? A might calculate that if A produces **Crunchy** regardless, then B has no incentive to execute the threat. Once A has taken the decision to produce **Crunchy**, the only payoffs relevant to B are those at the top of the game tree: *3* if B executes the threat and produces *Crunchy*, and *4* if B reneges on the threat and produces *Fruity*. Faced with these alternatives, B reneges on the threat, and A's favoured outcome, (**Crunchy**, *Fruity*), is achieved.

In the breakfast cereals game with A as first mover, (**Crunchy**, *Fruity*) is a **Subgame Perfect Equilibrium** (SPE). Any SPE is a Nash Equilibrium in the strategic form representation, but not all Nash Equilibria are SPEs in the extensive form representation. SPEs exclude any Nash Equilibrium, such as (**Fruity**, *Crunchy*), whose attainment would require either player to make non-credible threats that they would not execute if/when the time comes to do so. In a sequential game the classification of an SPE depends upon the order of play: it is obvious that if B held the first-mover advantage and the payoffs were the same, (**Fruity**, *Crunchy*) would be an SPE, but (**Crunchy**, *Fruity*) would not be an SPE.

Technological standards game with second-mover advantage

Suppose two suppliers of consumer IT products are considering whether to stick with an old technological standard, or introduce a new standard, for a particular item of IT gadgetry. Firm A has produced using the old standard for a number

		Firm B's strategies	
		Old standard	*New standard*
Firm A's strategies	**Old standard**	**5**, *0*	**3**, *2*
	New standard	**2**, *3*	**4**, *1*

Figure 9.13 Technological standards game: strategic form representation

of years and has a reputation for reliability, while Firm B is a recent entrant with no established reputation. If both suppliers produce using the old standard, A's existing customers, aware of A's reputation, will see no reason to purchase from B rather than A. A earns a payoff of 5 and B earns 0. If A switches to the new standard while B remains with the old, a minority of A's existing customers will, out of loyalty, switch as well, but a majority, content to remain with the old standard if they have the option of doing so, will purchase from B. A earns a payoff of 2, and B earns 3. If A produces using the old standard while B produces using the new, the majority of A's existing customers will purchase from A, while a minority, willing to switch to the new standard out of curiosity, will purchase from B. Finally, if A and B both produce using the new standard, forcing all customers to switch, a large majority will purchase from A out of loyalty, but a small minority, viewing B as a more credible producer for the new standard, will purchase from B. Figure 9.13 shows the payoff matrix in the same format as before, assuming both firms move simultaneously. Figure 9.14 shows the extensive form representation of the sequential game, drawn under alternative assumptions that A is first mover (left-hand panel) and B is first mover (right-hand panel).

Suppose first that A is first mover, and consider A's decision:

- If A selects the **Old** standard, B's payoff *of 2* for *New* exceeds B's payoff of *0* for *Old*, so B selects *New* and A earns a payoff of **3**.
- If A selects the **New** standard, B's payoff of *3* for *Old* exceeds B's payoff of *1* for *New*, so B selects *Old* and A earns a payoff of **2**.

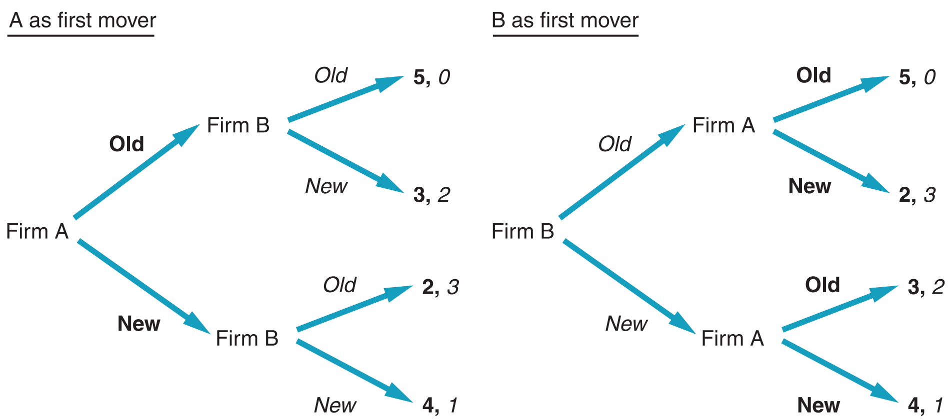

Figure 9.14 Technological standards game: extensive form representation

Therefore A selects **Old**, the outcome is (**Old**, *New*), and the payoffs are (**3**, *2*).

Suppose instead that B is first mover, and consider B's decision:

- If B selects the *Old* standard, A's payoff of **5** for **Old** exceeds A's payoff of **2** for **New**, so A selects **Old** and B earns a payoff of *0*.
- If B selects the *New* standard, A's payoff of **4** for **New** exceeds A's payoff of **3** for **Old**, so A selects **New** and B earns a payoff of *1*.

Therefore B selects *New*, the outcome is (**New**, *New*), and the payoffs are (**4**, *1*).

Note that A does better when B is first mover; and B does better when A is first mover. In each case there is a second-mover advantage: the second mover gains from being able to observe the first mover's decision before making his own decision.

Entry game with a weakly dominant strategy

Next we return to the entry game with a weakly dominant strategy for firm A, described in Sections 9.2 and 9.4, and presented in strategic form representation in Figure 9.2. On our previous encounter with this game, we established that there were multiple feasible outcomes, based on the criterion of Nash Equilibrium. Suppose, however, this game is played as a sequential game, in which either firm A or firm B moves first, and the other firm observes the first move before taking its decision. It turns out to be far easier to predict an outcome for the sequential game than for the equivalent simultaneous game.

Suppose firm A is the first mover: A must decide whether to prepare for a price war in the event that firm B enters, or accommodate by allowing B a share of the market. The left-hand panel of Figure 9.15 shows the extensive form representation.

- If A selects **Accommodate**, B's payoff of *4* for *Entry* exceeds B's payoff of *3* for *No entry*, so B selects *Entry* and A earns a payoff of **2**.
- If A selects **Fight**, B's payoff of *3* for *No entry* exceeds B's payoff of *1* for *Entry*, so B selects *No entry* and A earns a payoff of **3**.

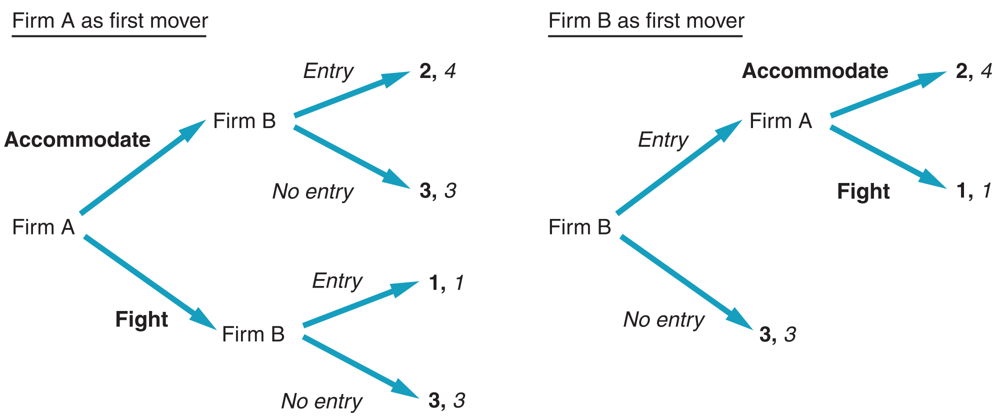

Figure 9.15 Entry game: extensive form representation

Therefore A selects **Fight**, the outcome is (**Fight**, *No entry*), and the payoffs are (**3**, *3*).

Suppose instead B is first mover: A can hold off from taking the decision as to whether to invest in new capacity until A has observed B's decision whether to enter or not. The right-hand panel of Figure 9.15 shows the extensive form representation.

- If B selects *Entry*, A's payoff of **2** for **Accommodate** exceeds A's payoff of **1** for **Fight**, so A selects **Accommodate** and B earns a payoff of *4*.
- If B selects *No entry*, A's payoff of **3** is the same for **Accommodate** and **Fight**, and B earns a payoff of *3* regardless which strategy A selects.

Therefore B selects *Entry*, the outcome is (**Accommodate**, *Entry*), and the payoffs are (**2**, *4*). This outcome is an SPE for the sequential game with B as first mover. By comparing the payoffs when either A or B is the first mover, it is clear that the entry game confers a first-mover advantage.

In the sequential game with B as first mover, could A have steered B towards a different outcome by threatening to fight, even in the event that B decides to enter? Clearly if B believes A will fight regardless of B's decision, B will prefer not to enter, rather than enter and fight a price war. From A's perspective, however, the problem with the threatened price war is that if B does choose to enter, A's threat to fight is no longer credible. If A executes the threat after B has already entered, a price war breaks out and A's payoff is **1;** but if A does not execute the threat, A's payoff is **2**. Recognizing that A's threat to fight is not credible once B has entered, B goes ahead and does so, and B's preferred outcome of (**Accommodate**, *Entry*) is achieved. Since this outcome is a Nash Equilibrium reliant only on credible threats, (**Accommodate**, *Entry*) is an SPE. Using similar reasoning, you can easily verify that in the sequential game with A as first mover, (**Fight**, *No entry*) is an SPE.

The structure of the entry game shown in Figures 9.2 and 9.15 is equivalent to what has become known as Selten's game, named after the Nobel Prize-winning game theorist Reinhard Selten (1975). Selten's contribution to game theory was the development of the concept of Subgame Perfect Equilibrium, and its application in demonstrating that some Nash Equilibria are more likely to occur than others, namely, those that are reinforced by credible retaliatory threats.

Burning bridges and building credibility

Thomas Schelling (1960), another Nobel Prize winner, focused on social interactions that contained elements of both conflict and common interest. One of Schelling's best known contributions is his description of the game known as **burning bridges**. Two warring countries are separated by a small island, and each country has one bridge providing the only means of access. The first country crosses its bridge and occupies the island, and must then decide whether to burn its bridge behind it. Subsequently the second country decides whether or not to invade. If the second country invades and the bridge has not been burnt, the first country must choose whether to fight or retreat. If the bridge has been

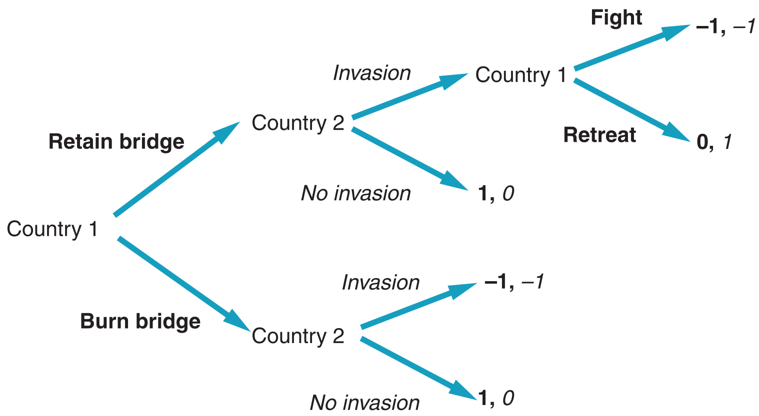

Figure 9.16 Burning bridges

burnt, the first country has no choice other than to fight if the second country invades. Figure 9.16 shows the game tree.

Suppose the first country does not burn its bridge, and retains the option to retreat. If the second country does not invade, the first country keeps the island and the payoffs are (**1**, *0*). If the second country invades, the first country has the option to either fight, which is damaging to both countries, reflected in payoffs of (**−1**, −*1*); or retreat and allow the second country to retain the island, with payoffs of (**0**, *1*). Faced with this choice, the first country prefers to retreat. Aware that the first country's threat to fight in the event of invasion is not credible, the second country invades and gains the island with payoffs of (**0**, *1*).

It is interesting to discover that the first country achieves a better outcome by limiting its options in the event that the second country invades. This may seem counter-intuitive, or contrary to the received wisdom that 'keeping your options open' is better than limiting options. If the first country burns its bridge, however, forfeiting the option of retreat, the second country knows that if it invades the first country will fight, with payoffs of (**−1**, −*1*); In this case the second country decides not to invade, and the outcome is that the first country retains the island with payoffs of (**1**, *0*). By burning its bridge, the first country demonstrates its commitment to remaining on the island, and establishes the credibility of the retaliatory threat if the second country invades. The second country, recognizing that the retaliatory threat is credible, abstains from invasion.

A similar approach for building credibility may assist firm A in achieving a better outcome in the entry game described previously, in the case where firm B enjoys first-mover advantage. The top-right section of the game tree shown in Figure 9.17 replicates the right-hand panel of Figure 9.15, in which B chooses between *Entry* and *No entry*, and then A decides whether to **Accommodate** or **Fight**.

Now suppose A has the opportunity, before B takes any decision, to undertake an investment that will reduce A's payoff in the event that B enters and A accommodates, but leave all other payoffs unchanged. The investment might entail switching to an alternative larger-scale technology, which is equally cost-effective

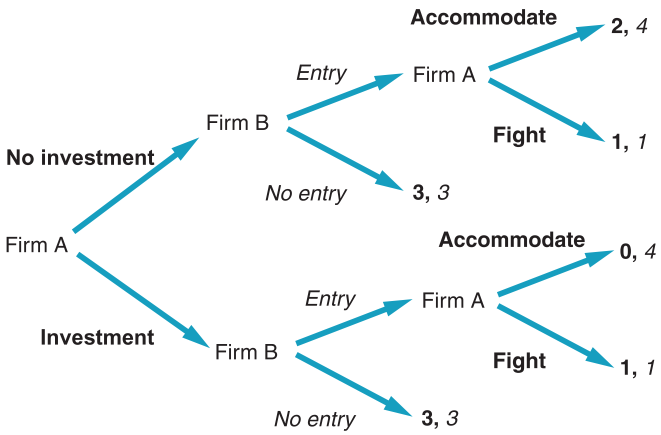

Figure 9.17 Building credibility in the entry game

if A produces at high volume (if A has the market to itself, or if A engages in a price war), but more expensive if A produces at low volume (if A shares the market with B).

The availability of this investment opportunity extends the game tree, and changes the final outcome. If A does not invest, the rest of the game plays out in the same way as before, and the final outcome is (**Accommodate**, *Entry*) with payoffs of (**2**, *4*). If A undertakes the investment, however, B recognises that if B enters, A's preference will now be to **Fight** rather than **Accommodate**, leading to a price war and a payoff for B of *1*. If B does not enter, B achieves a superior payoff of *3*. B prefers to select *No entry*, and the final payoffs are (**3**, *3*). As in the burning bridges game, by (effectively) eliminating the option of accommodating B in the event that B enters, A demonstrates commitment and establishes the credibility of the threat to fight, and deters B from entering. The final outcome leaves A with a superior payoff of **3** having undertaken the investment, by comparison with the original payoff of **2** if A does not invest. The notion that game theory players can themselves take steps to enhance the credibility of threatened retaliatory action was further developed by Spence (1977) and Dixit (1980).

9.6 Repeated games

In the previous discussion of single-period prisoner's dilemma and other games, it is assumed that the game is played only once. However, some games may be played repeatedly by the same players. Suppose firms A and B are hotdog sellers located outside a sports stadium. If the occasion is a one-off event such as the Olympics, and the two hotdog sellers are unlikely to ever see each other again, the game between them is a single-period game. In this case, the two hotdog sellers are less likely to cooperate. Suppose, however, the event is one that is repeated at regular intervals. Suppose the stadium is Old Trafford, the event is Manchester United home matches and the hotdog sellers see one another at regular, fortnightly intervals. In this case, it is more likely that cooperative behaviour will evolve as the two sellers observe and learn from each other's behaviour. In a **repeated** or **multiple-period game**, each firm may attempt to influence its rival's behaviour by sending signals that promise to reward cooperative behaviour and threaten to punish non-cooperative behaviour.

With reference to the production game with prisoner's dilemma structure shown in Figure 9.4, (**High**, *High*) is the (suboptimal) non-cooperative solution, which produces payoffs of 2 for firms A and B; and (**Low**, *Low*) is the (optimal) cooperative choice, which produces payoffs of 3 for both firms. In a single-period game, in which the firms act independently, **High** and *High* are the dominant strategies, and the non-cooperative outcome is likely to occur. However, suppose the game is to be repeated over an indefinite number of periods. Firm A could adopt the following strategy, known as **tit-for-tat**, in an attempt to encourage firm B to always select the cooperative choice:

- In period 1 A chooses **Low**.
- If B chose *Low* in period $t - 1$, in period t (for $t > 1$) A chooses **Low**.
- If B chose *High* in period $t - 1$, in period t (for $t > 1$) A chooses **High**.

In each period after the first, provided B chose the cooperative strategy last time, A rewards B by choosing the cooperative strategy this time. But if B chose the non-cooperative strategy last time, A punishes B by choosing the non-cooperative strategy this time. For as long as B cooperates, A also cooperates and the (optimal) cooperative solution is achieved. But if B attempts to exploit A's cooperation for short-term gain by defecting from *Low* to *High*, A punishes B in the following period by also switching from **Low** to **High**. However, B's punishment does not necessarily have to be long-lasting. Provided B learns from his error and switches back from *High* to *Low*, A also switches back from **High** to **Low** and cooperation is restored.

Since it is difficult to observe situations which replicate the structure of many theoretical games in practice, a subfield of economics known as **experimental economics** has been developed to test the predictions of game theory. Laboratory experimentation allows economists to determine the structure of games and test relevant hypotheses. Some economists are particularly optimistic about the future of this development:

> [A] hundred years from now, game theory will have become the backbone of a kind of microeconomics engineering that will have roughly the relation to the economic theory and laboratory experimentation of the time that chemical engineering has to chemical theory and bench chemistry.
>
> *(Roth, 1991, p. 107)*

In the present context, experiments have shown the adoption of a tit-for-tat strategy by one or both players is a highly effective method for ensuring adherence to cooperative behaviour in repeated games with a prisoner's dilemma structure. Usually, both players rapidly learn it is best for them to adhere to the cooperative strategy on each occasion the game is repeated.

However, there is one important caveat. Tit-for-tat is effective in infinitely repeated games in which there is no period when the game is played for the last time. Tit-for-tat may also be effective in games that are repeated only a finite number of times, but on each occasion neither player knows whether or not this *is* the last time the game will be played. However, tit-for-tat is likely to be ineffective in games that are repeated only a finite number of times, and on the final occasion the players know they will not play the game again.

Suppose the game is played for the last time in period T. In period T, B knows 'defecting' from *Low* to *High* will go unpunished, because the game will not be played again in $T + 1$. Therefore there is no deterrence and B defects. Realizing that B will behave in this way, A may as well abandon tit-for-tat in T, and also defect from **Low** to **High**. Therefore the non-cooperative outcome occurs in T.

From this reasoning, it might be supposed that the usefulness of the tit-for-tat strategy now finishes in period $T - 1$. In fact, the situation is actually worse than this, because in period $T - 1$ the same difficulty occurs. In $T - 1$, B knows 'defecting' from *Low* to *High* will go unpunished, because non-cooperation is going to happen anyway in T. Therefore there is no deterrence in $T - 1$ either, and B defects. Realizing that B will behave in this way, A may as well abandon tit-for-tat in $T - 1$, and also defect. Therefore the non-cooperative outcome also occurs in $T - 1$.

Similar reasoning will also apply in periods $T - 2$, $T - 3$ and so on, all the way back to the start of the game. In other words, the usefulness of tit-for-tat as a means for ensuring adherence to cooperative behaviour unravels completely due to the finite lifetime of the repeated game. A has no means of punishing B for non-cooperative behaviour in period T, so the tit-for-tat strategy fails in period T. But if tit-for-tat fails in period T, it also fails in $T - 1$; and if it fails in $T - 1$, it also fails in $T - 2$; and so on.

9.7 Summary

Game theory is an approach to decision-making in which two or more decision-makers or players face choices between a number of possible courses of action or actions at any stage of the game. The property of interdependence is the key defining characteristic of a game. Although game theory has many applications throughout the social and physical sciences, it is the treatment of interdependence that makes game theory relevant to an understanding of decision-making in oligopoly.

A game in which all players choose their actions simultaneously, before knowing the actions chosen by other players, is called a simultaneous game. In a two-person simultaneous game, when both players adopt a strictly dominant strategy, which is their best response no matter what strategy the other player chooses, a Dominant Strategy Equilibrium is achieved. Nash Equilibrium describes a broader range of game theory solutions, which satisfy the criterion that neither player wishes to depart from his current strategy if the other player continues to pursue his current strategy. A Dominant Strategy Equilibrium is always a Nash Equilibrium; but games in which one or both players have weakly dominant strategies, at least as good as any other strategy for all strategies the other player may choose, and a best response for at least one of the other player's strategies, may also exhibit one or more Nash Equilibria. In such cases, it may be difficult or impossible to identify any unambiguous solution to the game.

Analysis of prisoner's dilemma games shows that situations can arise in which players take decisions that appear rational from an individual perspective, but lead to outcomes that are suboptimal when assessed according to criteria reflecting the players' collective interest. Good communication between the players might assist in fostering cooperation, or rapid and effective punishment may be an effective deterrent against defection from cooperative behaviour.

For games with no Dominant Strategy Equilibrium, and games with either no Nash Equilibrium or multiple Nash Equilibria, players might decide to adopt mixed strategies, involving randomization of the choice between two or more actions, with specific probabilities defined for each action. In a Mixed Strategy Nash Equilibrium, each player selects the probabilities that maximise his own expected payoff, given the mixed strategy that is being employed by the other player.

In a sequential game, the players' decisions follow a sequence. One player makes his decision, and the other player observes this decision before making his response. Threats of a retaliatory nature, issued by the second player in an attempt to deter the first player from acting in a manner detrimental to the second player's interests, are only credible if the second player would still be willing to execute the threat in the event that the first player took the action the threat was intended to deter. Nash Equilibria that rely only on credible threats are known as Subgame Perfect Equilibria (SPE). In games with multiple Nash Equilibria, some of which rely on non-credible threats, the concept of SPE is helpful in identifying which of the Nash Equilibria is most likely to occur. In some cases it may be possible for a player to strengthen the credibility of threatened retaliatory action by taking steps to alter the payoffs in a manner that effectively shuts down the option of backing away from execution of the threat.

A game played more than once, known as a multiple-period or repeated game, can be repeated either indefinitely or a finite number of times. In a repeated game with a prisoner's dilemma structure, the players may be able to learn from their experience to cooperate by departing from their dominant strategies. Adherence to cooperative behaviour may be reinforced by the threat of punishment the next time the game is repeated, in the event that cooperation breaks down.

Discussion questions

1. Explain the relationship between Cournot's solution to the problem of output determination in duopoly and the game theory concept of the Nash equilibrium.
2. In repeated games, it is often assumed that rivals are more likely to cooperate with one another than to compete. Under what conditions might competition be likely to break out in a repeated game?

3. According to game theory, why might a player sometimes decide to randomise his strategies? Explain with reference to the concept of Mixed Strategy Nash Equilibrium.
4. In a technological standards game, two firms each have a competitive advantage with a different standard, but both will earn higher profits if they adopt a common standard, rather than each firm adopt the standard that each one prefers. Explain why the firms face a coordination problem in determining which standard to adopt, and suggest how this coordination problem might be resolved.
5. With reference to the relative magnitudes of the payoffs in the strategic form representation, examine the circumstances in which a sequential game would confer either a first-mover advantage, or a second-mover advantage.
6. In what ways might a firm that is threatened by an aggressive expansion strategy on the part of a competitor establish the credibility of its own threatened retaliatory action?
7. Explain the distinction between the concepts of Pure Strategy Nash Equilibrium and Subgame Perfect Equilibrium.
8. With reference to Case Study 9.1, outline the contribution of the model of the prisoner's dilemma to our understanding of strategic behaviour.

Computational questions

1. The following payoff matrix shows the profits to two firms, A and B, that need to decide whether to set a high or a low price for identical products produced by both firms:

		B's campaign focused on:	
		Low	*High*
A's price:	**Low**	**20**, *30*	**50**, *20*
	High	**10**, *60*	**40**, *50*

   (a) Identify the Dominant Strategy Equilibrium for this game, and justify your answer.
   (b) Is there an alternative cooperative solution that could offer both firms a higher return than they achieve at the Dominant Strategy Equilibrium?
   (c) What factors might be helpful in allowing the firms to adhere to the cooperative solution?
2. Firm A is an established monopoly supplier of a particular product. Firm B is a supplier to a different market, and is considering a diversification strategy that involves entry into A's market. Faced with the threat of entry from B, A must decide whether to accommodate B's presence by cutting back its own output in the event that B does enter, or plan to fight a price war with B. If B does not enter, A earns a profit of 500 and B earns a profit of 300. If B enters and A decides to accommodate, A earns 200 and B earns 600. If A decides to fight B in a price war, both firms earn a profit of 100.

(a) If A and B decide their strategies simultaneously, identify any Pure Strategy Nash Equilibria in this game.

(b) Suppose A can observe B's decision as to whether to enter or not, before A takes the decision to accommodate or fight. Write down the extensive form representation of the sequential game.

(c) Which of the Pure Strategy Nash Equilibria identified in part (a) is a Subgame Perfect Equilibrium?

(d) Is there a first-mover advantage or a second-mover advantage in the sequential game?

3. With reference to the sequential entry game described in Q2, before B decides whether or not to enter, A is presented with an opportunity to invest in an expansion of capacity. If A subsequently keeps the market to itself, or if A fights a price war with B, the investment will break even and A's payoffs will be unaffected. If B enters and A accommodates by sharing the market with B, however, the investment will be loss-making, and A's payoff will be reduced from 200 to 0.

(a) Write down the extensive form representation of the sequential game.

(b) Which outcome would you expect to occur?

(c) Explain why the outcome is different to the outcome identified in Q2 part (c).

4. You are given the following information, based on past observation of penalty kicks in football. 60% of penalties where a right-footed penalty-taker shot towards the left-hand side of the goal (his stronger side), and the goalkeeper dived to the same side, were successfully converted. 90% of penalties where a right-footed penalty-taker shot towards the left-hand side and the goalkeeper dived in the opposite direction were successfully converted. Where a right-footed penalty taker shot towards the right-hand side of the goal (his weaker side), 30% of penalties were successfully converted when the goalkeeper dived in the correct direction, and 80% were converted when the goalkeeper dived in the wrong direction.

(a) By applying the concept of Mixed Strategy Nash Equilibrium, show that the optimal mixed strategy for the penalty-taker is to randomise his choice of direction, by shooting to the left-hand side with probability 0.625 and to the right-hand side with probability 0.375.

(b) Show that the goalkeeper's optimal mixed strategy is also to randomise, diving towards the left-hand side of the goal (from the kicker's perspective) with probability 0.75, and towards the right-hand side with probability 0.25.

(c) Why should the goalkeeper dive to the left-hand side of the goal (from the kicker's perspective) more frequently than the penalty-taker kicks towards the left-hand side?

5. Firms A and B are simultaneously planning an advertising campaign for the launch of a new product that will have similar characteristics when produced by either firm. Firm A is an established company, and is well known to consumers. Firm B is a relatively unknown recent entrant. Both firms have to decide whether to focus their advertising campaigns primarily on traditional media (TV and print), or on social media. If both firms decide to advertise using the

same media, A's reputational advantage will tend to dominate the campaign, and B will struggle to establish a foothold in the market. If both firms decide to advertise using different media, B's campaign will succeed in enabling B to establish a foothold. The payoff matrix, expressed in terms of expected market share after the campaign, is as follows:

		B's campaign focused on:	
		Traditional media	*Social media*
A's campaign focused on:	**Traditional media**	**90**, *10*	**60**, *40*
	Social media	**50**, *50*	**80**, *20*

(a) Does this game have any Pure Strategy Nash Equilibria?

(b) Identify the Mixed Strategy Nash Equilibrium in the simultaneous game.

(c) Suppose B can observe A's decision before B decides how to focus its own campaign. Write down the extensive form representation of the sequential game.

(d) What is the likely outcome of the sequential game? Is there a first-mover advantage or a second-mover advantage?

Further reading

Aumann, R.J. (2008) Game theory, in Durlauf, S.N and Blume, L. (eds) *New Palgrave Dictionary of Economics*, 2nd edn. Basingstoke: Palgrave.

Binmore, K. (2007) *Game Theory: A very short introduction*. Oxford: Oxford University Press.

Chiappori, P., Levitt, S. and Groseclose, T. (2002) Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer, *American Economic Review,* 92, 1138–1151.

The Economist (2015) Applying game theory: Three's a charm, by R.L.G., 3 February.

The Economist (2016) Game theory: prison breakthrough – The Nash Equilibrium, 20 August.

Gibbons, R. (1997) An introduction to applicable game theory, *Journal of Economic Perspectives*, 11, 127–49.

Gul, F. (1997) A Nobel Prize for game theorists: the contributions of Harsanyi, Nash and Selten, *Journal of Economic Perspectives*, 11, 159–74.

Roth, A.E. (1991) Game theory as a part of empirical economics, *Economic Journal*, 101, 107–14.

Spaniel, W. (2011) *Game Theory 101: The complete textbook*. CreateSpace Independent Publishing Platform.

Turocy, T.L. and von Stengel, B. (2002) Game theory, in *Encyclopedia of Information Systems*, Vol 2. San Diego: Elsevier, pp. 403–20.