

3 Mixed and Continuous Strategies

3.1 Mixed Strategies: The Welfare Game

The games we have looked at have so far been simple in at least one respect: the number of moves in the action set has been finite. In this chapter we allow a continuum of moves, such as when a player chooses a price between 10 and 20 or a purchase probability between 0 and 1. Chapter 3 begins by showing how to find mixed-strategy equilibria for a game with no pure-strategy equilibria. In Section 3.2 the mixed-strategy equilibria are found by the payoff-equating method, and mixed strategies are applied to two dynamic games, the War of Attrition and Patent Race for a New Market. Section 3.3 takes a more general look at mixed strategy equilibria and extends the analysis to three or more players. Section 3.4 distinguishes between mixed strategies and random actions in the importance class of “auditing games”. Section 3.5 switches from the continuous strategy spaces of mixed strategies to strategy spaces that are continuous even with pure strategies, using the Cournot duopoly model, in which two firms choose output on the continuum between zero and infinity. Section 3.6 looks at the Bertrand model and strategic substitutes. Section 3.7 switches gears a bit and talks about four reasons why a Nash equilibrium might not exist. These last sections introduce a number of ideas besides simply how to find equilibria, ideas that will be built upon in later chapters—dynamic games in Chapter 4, auditing and agency in Chapters 7 and 8, and Cournot oligopoly in Chapter 14.

We invoked the concept of Nash equilibrium to provide predictions of outcomes without dominant strategies, but some games lack even a Nash equilibrium. It is often useful and realistic to expand the strategy space to include random strategies, in which case a Nash equilibrium almost always exists. These random strategies are called “mixed strategies.”

A **pure strategy** maps each of a player’s possible information sets to one action. $s_i : \omega_i \rightarrow a_i$.

A **mixed strategy** maps each of a player’s possible information sets to a probability distribution over actions.

$$s_i : \omega_i \rightarrow m(a_i), \text{ where } m \geq 0 \text{ and } \int_{A_i} m(a_i) da_i = 1.$$

A **completely mixed strategy** puts positive probability on every action, so $m > 0$.

The version of a game expanded to allow mixed strategies is called the **mixed extension** of the game.

A pure strategy constitutes a rule that tells the player what action to choose, while a mixed strategy constitutes a rule that tells him what dice to throw in order to choose an

action. If a player pursues a mixed strategy, he might choose any of several different actions in a given situation, an unpredictability which can be helpful to him. Mixed strategies occur frequently in the real world. In American football games, for example, the offensive team has to decide whether to pass or to run. Passing generally gains more yards, but what is most important is to choose an action not expected by the other team. Teams decide to run part of the time and pass part of the time in a way that seems random to observers but rational to game theorists.

The Welfare Game

The Welfare Game models a government that wishes to aid a pauper if he searches for work but not otherwise, and a pauper who searches for work only if he cannot depend on government aid.

Table 1 shows payoffs which represent the situation. “Work” represents trying to find work, and “Loaf” represents not trying. The government wishes to help a pauper who is trying to find work, but not one who does not try. Neither player has a dominant strategy, and with a little thought we can see that no Nash equilibrium exists in pure strategies either.

Table 1: The Welfare Game

		Pauper	
		<i>Work</i> (γ_w)	<i>Loaf</i> ($1 - \gamma_w$)
Government	<i>Aid</i> (θ_a)	3, 2	→ -1, 3
	<i>No Aid</i> ($1 - \theta_a$)	↑ -1, 1	← 0, 0

Payoffs to: (Government, Pauper). Arrows show how a player can increase his payoff.

Each strategy profile must be examined in turn to check for Nash equilibria.

1 The strategy profile (*Aid*, *Work*) is not a Nash equilibrium, because the pauper would respond with *Loaf* if the government picked *Aid*.

2 (*Aid*, *Loaf*) is not Nash, because the government would switch to *No Aid*.

3 (*No Aid*, *Loaf*) is not Nash, because the pauper would switch to *Work*.

4 (*No Aid*, *Work*) is not Nash, because the government would switch to *Aid*, which brings us back to (1).

The Welfare Game does have a mixed-strategy Nash equilibrium, which we can calculate. The players’ payoffs are the expected values of the payments from Table 1. If the government plays *Aid* with probability θ_a and the pauper plays *Work* with probability γ_w , the government’s expected payoff is

$$\begin{aligned}
 \pi_{\text{Government}} &= \theta_a[3\gamma_w + (-1)(1 - \gamma_w)] + [1 - \theta_a][-1\gamma_w + 0(1 - \gamma_w)] \\
 &= \theta_a[3\gamma_w - 1 + \gamma_w] - \gamma_w + \theta_a\gamma_w \\
 &= \theta_a[5\gamma_w - 1] - \gamma_w.
 \end{aligned} \tag{1}$$

If only pure strategies are allowed, θ_a equals zero or one, but in the mixed extension of the game, the government's action of θ_a lies on the continuum from zero to one, the pure strategies being the extreme values. If we followed the usual procedure for solving a maximization problem, we would differentiate the payoff function with respect to the choice variable to obtain the first-order condition. That procedure is actually not the best way to find mixed-strategy equilibria, which is the "payoff-equating method" I will describe in the next section. Let us use the maximization approach here, though, because it is good for helping you understand how mixed strategies work. The first-order condition for the government would be

$$0 = \frac{d\pi_{Government}}{d\theta_a} = 5\gamma_w - 1 \quad (2)$$

$$\Rightarrow \gamma_w = 0.2.$$

In the mixed-strategy equilibrium, the pauper selects *Work* 20 percent of the time. This is a bit strange, though: we obtained the pauper's strategy by differentiating the government's payoff! That is because we have not used maximization in the standard way. The problem has a corner solution, because depending on the pauper's strategy, one of three strategies maximizes the government's payoff: (i) Do not aid ($\theta_a = 0$) if the pauper is unlikely enough to try to work; (ii) Definitely aid ($\theta_a = 1$) if the pauper is likely enough to try to work; (iii) any probability of aid, if the government is indifferent because the pauper's probability of work is right on the border line of $\gamma_w = 0.2$.

It is possibility (iii) which allows a mixed strategy equilibrium to exist. To see this, go through the following 4 steps:

1 I assert that an optimal mixed strategy exists for the government.

2 If the pauper selects *Work* more than 20 percent of the time, the government always selects *Aid*. If the pauper selects *Work* less than 20 percent of the time, the government never selects *Aid*.

3 If a mixed strategy is to be optimal for the government, the pauper must therefore select *Work* with probability exactly 20 percent.

To obtain the probability of the government choosing *Aid*, we must turn to the pauper's payoff function, which is

$$\begin{aligned} \pi_{Pauper} &= \gamma_w(2\theta_a + 1[1 - \theta_a]) + (1 - \gamma_w)(3\theta_a + [0][1 - \theta_a]) \\ &= 2\gamma_w\theta_a + \gamma_w - \gamma_w\theta_a + 3\theta_a - 3\gamma_w\theta_a \\ &= -\gamma_w(2\theta_a - 1) + 3\theta_a. \end{aligned} \quad (3)$$

The first-order condition is

$$\begin{aligned} \frac{d\pi_{Pauper}}{d\theta_a} &= -(2\theta_a - 1) = 0, \\ \Rightarrow \theta_a &= 1/2. \end{aligned} \quad (4)$$

If the pauper selects *Work* with probability 0.2, the government is indifferent among selecting *Aid* with probability 100 percent, 0 percent, or anything in between. If the strategies are to form a Nash equilibrium, however, the government must choose $\theta_a = 0.5$. In the mixed-strategy Nash equilibrium, the government selects *Aid* with probability 0.5 and the pauper selects *Work* with probability 0.2. The equilibrium outcome could be any of the four entries in the outcome matrix. The entries having the highest probability of occurrence are *(No Aid, Loaf)* and *(Aid, Loaf)*, each with probability 0.4 ($= 0.5[1 - 0.2]$).

Interpreting Mixed Strategies

Mixed strategies are not as intuitive as pure strategies, and many modellers prefer to restrict themselves to pure-strategy equilibria in games which have them. One objection to mixed strategies is that people in the real world do not take random actions. That is not a compelling objection, because all that a model with mixed strategies requires to be a good description of the world is that the actions appear random to observers, even if the player himself has always been sure what action he would take. Even explicitly random actions are not uncommon, however— the Internal Revenue Service randomly selects which tax returns to audit, and telephone companies randomly monitor their operators' conversations to discover whether they are being polite.

A more troubling objection is that a player who selects a mixed strategy is always indifferent between two pure strategies. In the Welfare Game, the pauper is indifferent between his two pure strategies and a whole continuum of mixed strategies, given the government's mixed strategy. If the pauper were to decide not to follow the particular mixed strategy $\gamma_w = 0.2$, the equilibrium would collapse because the government would change its strategy in response. Even a small deviation in the probability selected by the pauper, a deviation that does not change his payoff if the government does not respond, destroys the equilibrium completely because the government does respond. A mixed-strategy Nash equilibrium is weak in the same sense as the *(North, North)* equilibrium in the Battle of the Bismarck Sea: to maintain the equilibrium a player who is indifferent between strategies must pick a particular strategy from out of the set of strategies.

One way to reinterpret the Welfare Game is to imagine that instead of a single pauper there are many, with identical tastes and payoff functions, all of whom must be treated alike by the government. In the mixed-strategy equilibrium, each of the paupers chooses *Work* with probability 0.2, just as in the one- pauper game. But the many-pauper game has a pure-strategy equilibrium: 20 percent of the paupers choose the pure strategy *Work* and 80 percent choose the pure strategy *Loaf*. The problem persists of how an individual pauper, indifferent between the pure strategies, chooses one or the other, but it is easy to imagine that individual characteristics outside the model could determine which actions are chosen by which paupers.

The number of players needed so that mixed strategies can be interpreted as pure strategies in this way depends on the equilibrium probability γ_w , since we cannot speak of a fraction of a player. The number of paupers must be a multiple of five in the Welfare Game in order to use this interpretation, since the equilibrium mixing probability is a multiple of $1/5$. For the interpretation to apply no matter how we vary the parameters of a model we would need a *continuum* of players.

Another interpretation of mixed strategies, which works even in the single-pauper game, assumes that the pauper is drawn from a population of paupers, and the government does not know his characteristics. The government only knows that there are two types of paupers, in the proportions (0.2, 0.8): those who pick *Work* if the government picks $\theta_a = 0.5$, and those who pick *Loaf*. A pauper drawn randomly from the population might be of either type. Harsanyi (1973) gives a careful interpretation of this situation.

Mixed Strategies Can Dominate Otherwise Undominated Pure Strategies

Before we continue with methods of calculating mixed strategies, it is worth taking a moment to show how they can be used to simplify the set of rational strategies players might use in a game. Chapter 1 talked about using the ideas of dominated strategies and iterated dominance as an alternative to Nash equilibrium, but it ignored the possibility of mixed strategies. That is a meaningful omission, because some pure strategy in a game may be strictly dominated by a mixed strategy, even if it is not dominated by any of the other pure strategies. The example in Table 2 illustrates this.

Table 2: Pure Strategies Dominated by a Mixed Strategy

		Column	
		<i>North</i>	<i>South</i>
Row	<i>North</i>	0,0	4,-4
	<i>South</i>	4,-4	0,0
	<i>Defense</i>	1,-1	1,-1

Payoffs to: (Row, Column)

In the zero-sum game of Table 2, Row's army can attack in the North, attack in the South, or remain on the defense. An unexpected attack gives Row a payoff of 4, an expected attack a payoff of 0, and defense a payoff of 1. Column can respond by preparing to defend in the North or in the South.

Row could guarantee himself a payoff of 1 if he chose *Defense*. But suppose he plays *North* with probability 0.5 and *South* with probability 0.5. His expected payoff from this mixed strategy if Column plays *North* with probability N is

$$0.5(N)(0) + 0.5(1 - N)(4) + 0.5(N)(4) + 0.5(1 - N)(0) = 2, \quad (5)$$

so whatever response Column picks, Row's expected payoff is higher from the mixed strategy than his payoff of 1 from *Defense*. For Row, *Defense* is strictly dominated by (0.5 *North*, 0.5 *South*).

"What if Row is risk averse?", you may ask. "Might he not prefer the sure payoff of 1 from playing *Defense*?" No. Payoffs are specified in units of utility, not of money or some other input into a utility function. In Table 2, it might be that Row's payoff of 0 represents

gaining no territory, 1 represents 100 square miles, and 4 represents 800 square miles, so the marginal payoff of territory acquisition (or loss, for *South*) is declining. When using mixed strategies it is particularly important to keep track of the difference between utility and the inputs into utility.

Thus, regardless of risk aversion, in the unique Nash equilibrium of Pure Strategies Dominated by a Mixed Strategy, Row and Column would both choose *North* with probability $N = 0.5$ and *South* with probability 0.5. This is a player's unique equilibrium action because any other choice would cause the other player to deviate to whichever direction was not being guarded as often.

3.2 The Payoff-Equating Method and Games of Timing

The next game illustrates why we might decide that a mixed-strategy equilibrium is best even if pure-strategy equilibria also exist. In the game of *Chicken*, the players are two Malibu teenagers, Smith and Jones. Smith drives a hot rod south down the middle of Route 1, and Jones drives north. As collision threatens, each decides whether to *Continue* in the middle or *Swerve* to the side. If a player is the only one to *Swerve*, he loses face, but if neither player picks *Swerve* they are both killed, which has an even lower payoff. If a player is the only one to *Continue*, he is covered with glory, and if both *Swerve* they are both embarrassed. (We will assume that to *Swerve* means by convention to *Swerve* right; if one swerved to the left and the other to the right, the result would be both death and humiliation.) Table 3 assigns numbers to these four outcomes.

Table 3: *Chicken*

		Jones	
		<i>Continue</i> (θ)	<i>Swerve</i> ($1 - \theta$)
Smith:	<i>Continue</i> (θ)	-3, -3	→ 2, 0
		↓	↑
	<i>Swerve</i> ($1 - \theta$)	0, 2	← 1, 1

Payoffs to: (Smith, Jones). Arrows show how a player can increase his payoff.

Chicken has two pure-strategy Nash equilibria, (*Swerve*, *Continue*) and (*Continue*, *Swerve*), but they have the defect of asymmetry. How do the players know which equilibrium is the one that will be played out? Even if they talk before the game started, it is not clear how they could arrive at an asymmetric result. We encountered the same dilemma in choosing an equilibrium for *Battle of the Sexes*. As in that game, the best prediction in *Chicken* is perhaps the mixed-strategy equilibrium, because its symmetry makes it a focal point of sorts, and does not require any differences between the players.

The **payoff-equating** method used here to calculate the mixing probabilities for *Chicken* will be based on the logic followed in Section 3.1, but it does not use the calculus of maximization. The basis of the payoff-equating method is that **when a player uses a mixed strategy in equilibrium, he must be getting the same payoff from each of the pure strategies used in the mixed strategy**. If one of his mixing strategies has a higher payoff, he should deviate to use just that one instead of mixing. If one has a lower payoff, he should deviate by dropping it from his mixing.

In Chicken, therefore, Smith's payoffs from the pure strategies of *Swerve* and *Continue* must be equal. Moreover, Chicken, unlike the Welfare Game, is a symmetric game, so we can guess that in equilibrium each player will choose the same mixing probability. If that is the case, then, since the payoffs from each of Jones' pure strategies must be equal in a mixed-strategy equilibrium, it is true that

$$\begin{aligned}\pi_{Jones}(Swerve) &= (\theta_{Smith}) \cdot (0) + (1 - \theta_{Smith}) \cdot (1) \\ &= (\theta_{Smith}) \cdot (-3) + (1 - \theta_{Smith}) \cdot (2) = \pi_{Jones}(Continue).\end{aligned}\tag{6}$$

From equation (6) we can conclude that $1 - \theta_{Smith} = 2 - 5\theta_{Smith}$, so $\theta_{Smith} = 0.25$. In the symmetric equilibrium, both players choose the same probability, so we can replace θ_{Smith} with simply θ . As for the question which represents the greatest interest to their mothers, the two teenagers will survive with probability $1 - (\theta \cdot \theta) = 0.9375$.

The payoff-equating method is easier to use than the calculus method if the modeller is sure which strategies will be mixed, and it can also be used in asymmetric games. In the Welfare Game, it would start with $V_g(Aid) = V_g(No\ Aid)$ and $V_p(Loaf) = V_p(Work)$, yielding two equations for the two unknowns, θ_a and γ_w , which when solved give the same mixing probabilities as were found earlier for that game. The reason why the payoff-equating and calculus maximization methods reach the same result is that the expected payoff is linear in the possible payoffs, so differentiating the expected payoff equalizes the possible payoffs. The only difference from the symmetric-game case is that two equations are solved for two different mixing probabilities instead of a single equation for the one mixing probability that both players use.

It is interesting to see what happens if the payoff of -3 in the northwest corner of Table 3 is generalized to x . Solving the analog of equation (6) then yields

$$\theta = \frac{1}{1 - x}.\tag{7}$$

If $x = -3$, this yields $\theta = 0.25$, as was just calculated, and if $x = -9$, it yields $\theta = 0.10$. This makes sense; increasing the loss from crashes reduces the equilibrium probability of continuing down the middle of the road. But what if $x = 0.5$? Then the equilibrium probability of continuing appears to be $\theta = 2$, which is impossible; probabilities are bounded by zero and one.

When a mixing probability is calculated to be greater than one or less than zero, the implication is either that the modeller has made an arithmetic mistake or, as in this case, that he is wrong in thinking that the game has a mixed-strategy equilibrium. If $x = 0.5$, one can still try to solve for the mixing probabilities, but, in fact, the only equilibrium is in pure strategies— (*Continue, Continue*) (the game has become a prisoner's dilemma). The absurdity of probabilities greater than one or less than zero is a valuable aid to the fallible modeller because such results show that he is wrong about the qualitative nature of the equilibrium— it is pure, not mixed. Or, if the modeller is not sure whether the equilibrium is mixed or not, he can use this approach to prove that the equilibrium is not in mixed strategies.

The War of Attrition

After the start of the book with the Dry Cleaners Game, we have been looking at games that are either simultaneous or have the players move in sequence. Some situations, however, are naturally modelled as flows of time during which players repeatedly choose their moves. The War of Attrition is one of these. It is a game something like Chicken stretched out over time, where both players start with *Continue*, and the game ends when the first one picks *Swerve*. Until the game ends, both earn a negative amount per period, and when one exits, he earns zero and the other player earns a reward for outlasting him.

We will look at a war of attrition in discrete time. We will continue with Smith and Jones, who have both survived to maturity and now play games with more expensive toys: they control two firms in an industry which is a natural monopoly, with demand strong enough for one firm to operate profitably, but not two. The possible actions are to *Exit* or to *Continue*. In each period that both *Continue*, each earns -1 . If a firm exits, its losses cease and the remaining firm obtains the value of the market's monopoly profit, which we set equal to 3. We will set the discount rate equal to $r > 0$, although that is inessential to the model, even if the possible length of the game is infinite (discount rates will be discussed in detail in Section 4.3).

The War of Attrition has a continuum of Nash equilibria. One simple equilibrium is for Smith to choose (*Continue* regardless of what Jones does) and for Jones to choose (*Exit* immediately), which are best responses to each other. But we will solve for a symmetric equilibrium in which each player chooses the same mixed strategy: a constant probability θ that the player picks *Exit* given that the other player has not yet exited.

We can calculate θ as follows, adopting the perspective of Smith. Denote the expected discounted value of Smith's payoffs by V_{stay} if he stays and V_{exit} if he exits immediately. These two pure strategy payoffs must be equal in a mixed strategy equilibrium (which was the basis for the payoff-equating method). If Smith exits, he obtains $V_{exit} = 0$. If Smith stays in, his payoff depends on what Jones does. If Jones stays in too, which has probability $(1 - \theta)$, Smith gets -1 currently and his expected value for the following period, which is discounted using r , is unchanged. If Jones exits immediately, which has probability θ , then Smith receives a payment of 3. In symbols,

$$V_{stay} = \theta \cdot (3) + (1 - \theta) \left(-1 + \left[\frac{V_{stay}}{1 + r} \right] \right), \quad (8)$$

which, after a little manipulation, becomes

$$V_{stay} = \left(\frac{1 + r}{r + \theta} \right) (4\theta - 1). \quad (9)$$

Once we equate V_{stay} to V_{exit} , which equals zero, equation (9) tells us that $\theta = 0.25$ in equilibrium, and that this is independent of the discount rate r .

Returning from arithmetic to ideas, why does Smith *Exit* immediately with positive probability, given that Jones will exit first if Smith waits long enough? The reason is that Jones might choose to continue for a long time and both players would earn -1 each period until Jones exited. The equilibrium mixing probability is calculated so that both of them are likely to stay in long enough so that their losses soak up the gain from being the

survivor. Papers on the War of Attrition include Fudenberg & Tirole (1986b), Ghemawat & Nalebuff (1985), Maynard Smith (1974), Nalebuff & Riley (1985), and Riley (1980). All are examples of “rent-seeking” welfare losses. As Posner (1975) and Tullock (1967) have pointed out, the real costs of acquiring rents can be much bigger than the second-order triangle losses from allocative distortions, and the war of attrition shows that the big loss from a natural monopoly might be not the reduced trade that results from higher prices, but the cost of the battle to gain the monopoly.

We are likely to see wars of attrition in business when new markets open up, either new geographic markets for old goods or new goods, especially when it appears the market may be a natural monopoly, as in situations of network externalities. McAfee (2002, p. 76, 364) cites as examples the fight between Sky Television and British Satellite Broadcasting for the British satellite TV market; Amazon versus Barnes and Noble for the Internet book market; and Windows CE versus Palm in the market for handheld computers. Wars of attrition can also arise in declining industries, as a contest for which firm can exit the latest. In the United States, for example, the number of firms making rockets declined from six in 1990 to two by 2002 (McAfee [2002, p. 104]).

In the War of Attrition, the reward goes to the player who does not choose the move which ends the game, and a cost is paid each period that both players refuse to end it. Various other **timing games** also exist. The opposite of a war of attrition is a **pre-emption game**, in which the reward goes to the player who chooses the move which ends the game, and a cost is paid if both players choose that move, but no cost is incurred in a period when neither player chooses it. The game of **Grab the Dollar** is an example. A dollar is placed on the table between Smith and Jones, who each must decide whether to grab for it or not. If both grab, both are fined one dollar. This could be set up as a one-period game, a T period game, or an infinite- period game, but the game definitely ends when someone grabs the dollar. Table 4 shows the payoffs.

Table 4: Grab the Dollar

		Jones	
		<i>Grab</i>	<i>Don't Grab</i>
Smith:	<i>Grab</i>	-1, -1 →	1, 0
	<i>Don't Grab</i>	↓ 0, 1 ←	↑ 0, 0

Payoffs to: (Smith, Jones). Arrows show how a player can increase his payoff.

Like the War of Attrition, Grab the Dollar has asymmetric equilibria in pure strategies, and a symmetric equilibrium in mixed strategies. In the infinite- period version, the equilibrium probability of grabbing is 0.5 per period in the symmetric equilibrium.

Still another class of timing games are duels, in which the actions are discrete occurrences which the players locate at particular points in continuous time. Two players with guns approach each other and must decide when to shoot. In a **noisy duel**, if a player shoots and misses, the other player observes the miss and can kill the first player at his leisure. An equilibrium exists in pure strategies for the noisy duel. In a **silent duel**, a player does not know when the other player has fired, and the equilibrium is in mixed

strategies. Karlin (1959) has details on duelling games, and Chapter 4 of Fudenberg & Tirole (1991a) has an excellent discussion of games of timing in general. See also Shubik (1954) on the rather different problem of who to shoot first in a battle with three or more sides.

We will go through one more game of timing to see how to derive a continuous mixed strategies probability distribution, instead of just the single number derived earlier. In presenting this game, a new presentation scheme will be useful. If a game has a continuous strategy set, it is harder or impossible to depict the payoffs using tables or the extensive form using a tree. Tables of the sort we have been using so far would require a continuum of rows and columns, and trees a continuum of branches. A new format for game descriptions of the players, actions, and payoffs will be used for the rest of the book. The new format will be similar to the way the rules of the Dry Cleaners Game were presented in Section 1.1.

Patent Race for a New Market

Players

Three identical firms, Apex, Brydox, and Central.

The Order of Play

Each firm simultaneously chooses research spending $x_i \geq 0$, ($i = a, b, c$).

Payoffs

Firms are risk neutral and the discount rate is zero. Innovation occurs at time $T(x_i)$ where $T' < 0$. The value of the patent is V , and if several players innovate simultaneously they share its value. Let us look at the payoff of firm $i = a, b, c$, with j and k indexing the other two firms:

$$\pi_i = \begin{cases} V - x_i & \text{if } T(x_i) < \text{Min}\{T(x_j), T(x_k)\} & \text{(Firm } i \text{ gets the patent)} \\ \frac{V}{2} - x_i & \text{if } T(x_i) = \text{Min}\{T(x_j), T(x_k)\} & \text{(Firm } i \text{ shares the patent with} \\ & < \text{Max}\{T(x_j), T(x_k)\} & \text{1 other firm)} \\ \frac{V}{3} - x_i & \text{if } T(x_i) = T(x_j) = T(x_k) & \text{(Firm } i \text{ shares the patent with} \\ & & \text{2 other firms)} \\ -x_i & \text{if } T(x_i) > \text{Min}\{T(x_j), T(x_k)\} & \text{(Firm } i \text{ does not get the patent)} \end{cases}$$

The format first assigns the game a title, after which it lists the players, the order of play (together with who observes what), and the payoff functions. Listing the players is redundant, strictly speaking, since they can be deduced from the order of play, but it is useful for letting the reader know what kind of model to expect. The format includes very little explanation; that is postponed, lest it obscure the description. This exact format is not standard in the literature, but every good article begins its technical section by specifying the same information, if in a less structured way, and the novice is strongly advised to use all the structure he can.

The game Patent Race for a New Market does not have any pure strategy Nash equilibria, because the payoff functions are discontinuous. A slight difference in research by one player can make a big difference in the payoffs, as shown in Figure 1 for fixed values of x_b and x_c . The research levels shown in Figure 1 are not equilibrium values. If Apex chose any research level x_a less than V , Brydox would respond with $x_a + \varepsilon$ and win the patent. If Apex chose $x_a = V$, then Brydox and Central would respond with $x_b = 0$ and $x_c = 0$, which would make Apex want to switch to $x_a = \varepsilon$.

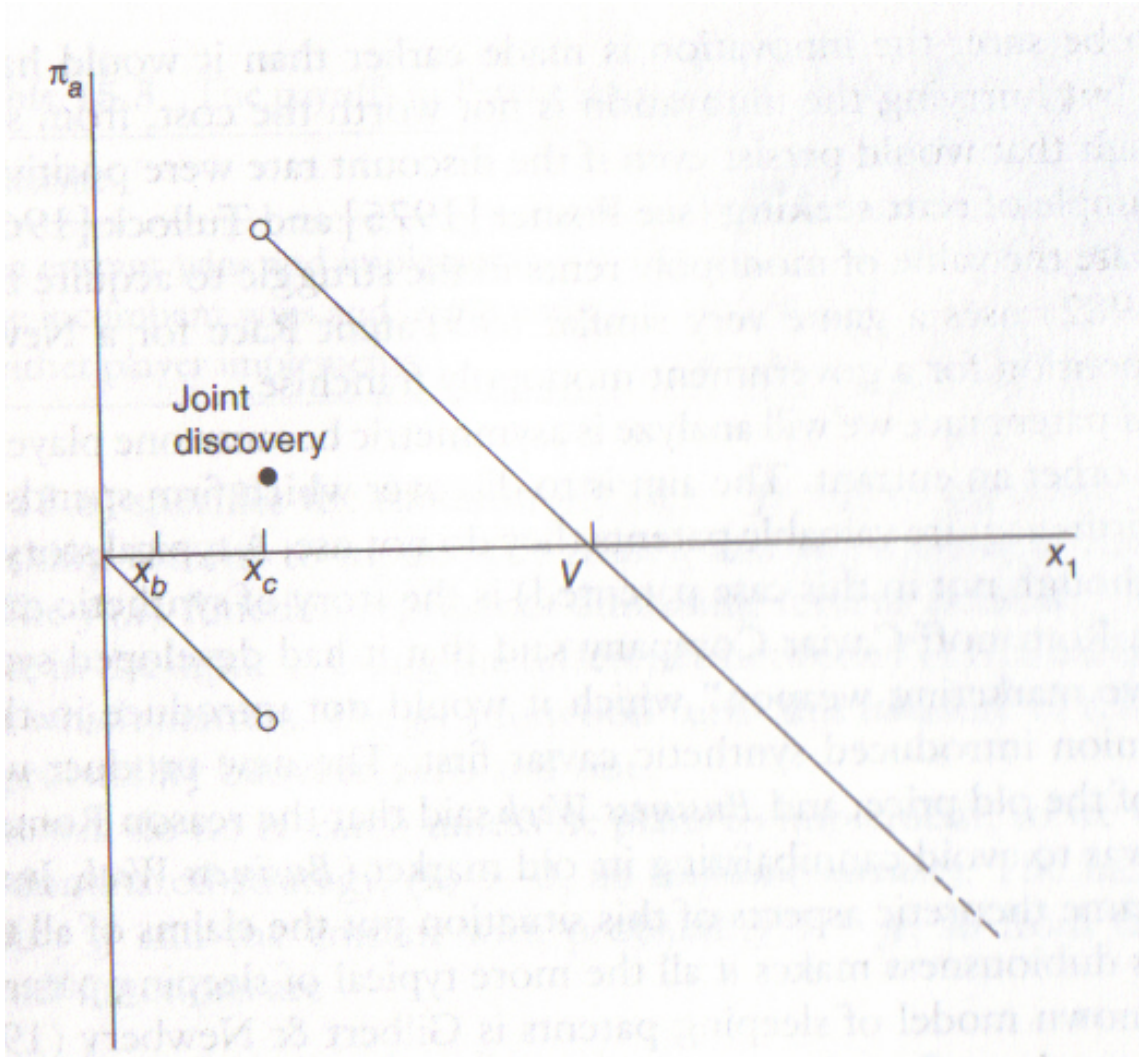


Figure 1: The Payoffs in Patent Race for a New Market

There does exist a symmetric mixed strategy equilibrium. Denote the probability that firm i chooses a research level less than or equal to x as $M_i(x)$. This function describes the firm's mixed strategy. In a mixed- strategy equilibrium a player is indifferent between any of the pure strategies among which he is mixing (the basis of Section 3.2's payoff-equating method). Since we know that the pure strategies $x_a = 0$ and $x_a = V$ yield zero payoffs, if Apex mixes over the support $[0, V]$ then the expected payoff for every strategy mixed between must also equal zero. The expected payoff from the pure strategy x_a is the

expected value of winning minus the cost of research. Letting x stand for nonrandom and X for random variables, this is

$$\pi_a(x_a) = V \cdot Pr(x_a \geq X_b, x_a \geq X_c) - x_a = 0 = \pi_a(x_a = 0), \quad (10)$$

which can be rewritten as

$$V \cdot Pr(X_b \leq x_a)Pr(X_c \leq x_a) - x_a = 0, \quad (11)$$

or

$$V \cdot M_b(x_a)M_c(x_a) - x_a = 0. \quad (12)$$

We can rearrange equation (12) to obtain

$$M_b(x_a)M_c(x_a) = \frac{x_a}{V}. \quad (13)$$

If all three firms choose the same mixing distribution M , then

$$M(x) = \left(\frac{x}{V}\right)^{1/2} \text{ for } 0 \leq x \leq V. \quad (14)$$

What is noteworthy about a patent race is not the nonexistence of a pure strategy equilibrium but the overexpenditure on research. All three players have expected payoffs of zero, because the patent value V is completely dissipated in the race. As in Brecht's *Threepenny Opera* (Act III, Scene 7), "When all race after happiness/Happiness comes in last." To be sure, the innovation is made earlier than it would have been by a monopolist, but hurrying the innovation is not worth the cost, from society's point of view, a result that would persist even if the discount rate were positive. Rogerson (1982) uses a game very similar to Patent Race for a New Market to analyze competition for a government monopoly franchise. Also, we will see in Chapter 13 that this is an example of an "all-pay auction", and the techniques and findings of auction theory can be quite useful when modelling this kind of conflict.

Correlated Strategies

One example of a war of attrition is setting up a market for a new security, which may be a natural monopoly for reasons to be explained in Section 8.5. Certain stock exchanges have avoided the destructive symmetric equilibrium by using lotteries to determine which of them would trade newly listed stock options under a system similar to the football draft.¹ Rather than waste resources fighting, these exchanges use the lottery as a coordinating device, even though it might not be a binding agreement.

Aumann (1974, 1987) has pointed out that it is often important whether players can use the same randomizing device for their mixed strategies. If they can, we refer to the resulting strategies as **correlated strategies**. Consider the game of Chicken. The only mixed-strategy equilibrium is the symmetric one in which each player chooses *Continue*

¹"Big Board will Begin Trading of Options on 4 Stocks it Lists," *Wall Street Journal*, p. 15 (4 October 1985).

with probability 0.25 and the expected payoff is 0.75. A correlated equilibrium would be for the two players to flip a coin and for Smith to choose *Continue* if it comes up heads and for Jones to choose *Continue* otherwise. Each player's strategy is a best response to the other's, the probability of each choosing *Continue* is 0.5, and the expected payoff for each is 1.0, which is better than the 0.75 achieved without correlated strategies.

Usually the randomizing device is not modelled explicitly when a model refers to correlated equilibrium. If it is, uncertainty over variables that do not affect preferences, endowments, or production is called **extrinsic uncertainty**. Extrinsic uncertainty is the driving force behind **sunspot models**, so called because the random appearance of sunspots might cause macroeconomic changes via correlated equilibria (Maskin & Tirole [1987]) or bets made between players (Cass & Shell [1983]).

One way to model correlated strategies is to specify a move in which Nature gives each player the ability to commit first to an action such as *Continue* with equal probability. This is often realistic because it amounts to a zero probability of both players entering the industry at exactly the same time without anyone knowing in advance who will be the lucky starter. Neither firm has an a priori advantage, but the outcome is efficient.

The population interpretation of mixed strategies cannot be used for correlated strategies. In ordinary mixed strategies, the mixing probabilities are statistically independent, whereas in correlated strategies they are not. In *Chicken*, the usual mixed strategy can be interpreted as populations of Smiths and Joneses, each population consisting of a certain proportion of pure swervers and pure stayers. The correlated equilibrium has no such interpretation.

Another coordinating device, useful in games that, like *Battle of the Sexes*, have a coordination problem, is **cheap talk** (Crawford & Sobel [1982], Farrell [1987]). Cheap talk refers to costless communication before the game proper begins. In *Ranked Coordination*, cheap talk instantly allows the players to make the desirable outcome a focal point. In *Chicken*, cheap talk is useless, because it is dominant for each player to announce that he will choose *Continue*. But in *The Battle of the Sexes*, coordination and conflict are combined. Without communication, the only symmetric equilibrium is in mixed strategies. If both players know that making inconsistent announcements will lead to the wasteful mixed-strategy outcome, then they are willing to mix announcing whether they will go to the ballet or the prize fight. With many periods of announcements before the final decision, their chances of coming to an agreement are high. Thus communication can help reduce inefficiency even if the two players are in conflict.

*3.3 Mixed Strategies with General Parameters and N Players: The Civic Duty Game

Having looked at a number of specific games with mixed-strategy equilibria, let us now apply the method to the general game of Table 5.

Table 5: The General 2-by-2 Game

		Column	
		<i>Left</i> (θ)	<i>Right</i> ($1 - \theta$)
Row:	<i>Up</i> (γ)	a, w	b, x
	<i>Down</i> ($1 - \gamma$)	c, y	d, z

Payoffs to: (Row, Column)

To find the game's equilibrium, equate the payoffs from the pure strategies. For Row, this yields

$$\pi_{Row}(Up) = \theta a + (1 - \theta)b \quad (15)$$

and

$$\pi_{Row}(Down) = \theta c + (1 - \theta)d. \quad (16)$$

Equating (15) and (16) gives us

$$\theta(a + d - b - c) + b - d = 0, \quad (17)$$

which yields

$$\theta^* = \frac{d - b}{(d - b) + (a - c)}. \quad (18)$$

Similarly, equating the payoffs for Column gives

$$\pi_{Column}(Left) = \gamma w + (1 - \gamma)y = \pi_{Column}(Right) = \gamma x + (1 - \gamma)z, \quad (19)$$

which yields

$$\gamma^* = \frac{z - y}{(z - y) + (w - x)}. \quad (20)$$

The equilibrium represented by (18) and (20) illustrates a number of features of mixed strategies.

First, it is possible, but wrong, to follow the payoff-equating method for finding a mixed strategy even if no mixed strategy equilibrium actually exists. Suppose, for example, that *Down* is a strictly dominant strategy for Row, so $c > a$ and $d > b$. Row is unwilling to mix, so the equilibrium is not in mixed strategies. Equation (18) would be misleading, though some idiocy would be required to stay misled for very long since the equation implies that $\theta^* > 1$ or $\theta^* \leq 0$ in cases like that.

Second, the exact features of the equilibrium in mixed strategies depend heavily on the cardinal values of the payoffs, not just on their ordinal values like the pure strategy equilibria in other 2-by-2 games. Ordinal rankings are all that is needed to know that an equilibrium exists in mixed strategies, but cardinal values are needed to know the exact mixing probabilities. If the payoff to Column from (*Confess*, *Confess*) is changed slightly in the Prisoner's Dilemma it makes no difference at all to the equilibrium. If the payoff of z to Column from (*Down*, *Right*) is increased slightly in the General 2-by-2 Game, equation (20) says that the mixing probability γ^* will change also.

Third, the payoffs can be changed by affine transformations without changing the game substantively, even though cardinal payoffs do matter (which is to say that monotonic but

non-affine transformations do make a difference). Let each payoff π in Table 5 become $\alpha + \beta\pi$. Equation (20) then becomes

$$\begin{aligned}\gamma^* &= \frac{\alpha + \beta z - \alpha - \beta y}{(\alpha + \beta z - \alpha - \beta y) + (\alpha + \beta w - \alpha - \beta x)} \\ &= \frac{z - y}{(z - y) + (w - x)}.\end{aligned}\tag{21}$$

The affine transformation has left the equilibrium strategy unchanged.

Fourth, as was mentioned earlier in connection with the Welfare Game, each player’s mixing probability depends only on the payoff parameters of the other player. Row’s strategy γ^* in equation (20) depends on the parameters w, x, y and z , which are the payoff parameters for Column, and have no direct relevance for Row.

Categories of Games with Mixed Strategies

Table 6 uses the players and actions of Table 5 to depict three major categories of 2-by-2 games in which mixed-strategy equilibria are important. Some games fall in none of these categories— those with tied payoffs, such as The Swiss Cheese Game in which all eight payoffs equal zero— but the three games in Table 6 encompass a wide variety of economic phenomena.

Table 6: 2-by-2 Games with Mixed Strategy Equilibria

$\begin{array}{cc} a, w & \rightarrow & b, x \\ \uparrow & & \downarrow \\ c, y & \leftarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \leftarrow & b, x \\ \downarrow & & \uparrow \\ c, y & \rightarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \leftarrow & b, x \\ \uparrow & & \downarrow \\ c, y & \rightarrow & d, z \end{array}$	$\begin{array}{cc} a, w & \rightarrow & b, x \\ \downarrow & & \uparrow \\ c, y & \leftarrow & d, z \end{array}$
Discoordination Games	Coordination Games	Contribution Games	

Payoffs to: (Row, Column). Arrows show how a player can increase his payoff.

Discoordination games have a single equilibrium, in mixed strategies. The payoffs are such that either (a) $a > c$, $d > b$, $x > w$, and $y > z$, or (b) $c > a$, $b > d$, $w > x$, and $z > y$. The Welfare Game is a discoordination game, as are Auditing Game I in the next section and Matching Pennies in problem 3.3.

Coordination games have three equilibria: two symmetric equilibria in pure strategies and one symmetric equilibrium in mixed strategies. The payoffs are such that $a > c$, $d > b$, $w > x$, and $z > y$. Ranked Coordination and The Battle of the Sexes are two varieties of coordination games in which the players have the same and opposite rankings of the pure-strategy equilibria.

Contribution games have three equilibria: two asymmetric equilibria in pure strategies and one symmetric equilibrium in mixed strategies. The payoffs are such that $c > a$, $b > d$, $x > w$, and $y > z$. Also, it must be true that $c < b$ and $y > x$, or $c > b$ and $y < x$.

I have used the name “contribution game” because the type of game described by this term is often used to model a situation in which two players each have a choice of taking some action that contributes to the public good, but would each like the other player to bear

the cost. The difference from the Prisoner’s Dilemma is that each player in a contribution game is willing to bear the cost alone if necessary.

Contribution games appear to be quite different from the Battle of the Sexes, but they are essentially the same. Both of them have two pure-strategy equilibria, ranked oppositely by the two players. In mathematical terms, the fact that contribution games have the equilibria in the southwest and northeast corners of the outcome matrix whereas coordination games have them in the northwest and southeast, is unimportant; the location of the equilibria could be changed by just switching the order of Row’s strategies. We do view real situations differently, however, depending on whether players choose the same actions or different actions in equilibrium.

Let us take a look at a particular contribution game to show how to extend two-player games to games with several players. A notorious example in social psychology is the murder of Kitty Genovese, who was killed in New York City in 1964 despite the presence of numerous neighbors.

“For more than half an hour 38 respectable, law-abiding citizens in Queens watched a killer stalk and stab a woman in three separate attacks in Kew Gardens.... Twice the sound of their voices and the sudden glow of their bedroom lights interrupted him and frightened him off. Each time he returned, sought her out, and stabbed her again. Not one person telephoned the police during the assault; one witness called after the woman was dead.” (Martin Gansberg, “38 Who Saw Murder Didn’t Call Police,” *The New York Times*, March 27, 1964, p. 1.)

Even as hardened an economist as myself finds it somewhat distasteful to call this a “game,” but game theory does explain what happened.

I will use a less appalling story for our model. In the Civic Duty Game of Table 7, Smith and Jones observe a burglary taking place. Each would like someone to call the police and stop the burglary because having it stopped adds 10 to his payoff, but neither wishes to make the call himself because the effort subtracts 3. If Smith can be assured that Jones will call, Smith himself will ignore the burglary. Table 7 shows the payoffs.

Table 7: The Civic Duty Game

		Jones		
		<i>Ignore</i> (γ)	<i>Telephone</i> ($1 - \gamma$)	
Smith:	<i>Ignore</i> (γ)	0, 0	→	10, 7
	<i>Telephone</i> ($1 - \gamma$)	↓	←	↑ 7, 7
		7, 10		

Payoffs to: (Row, Column). Arrows show how a player can increase his payoff.

The Civic Duty Game has two asymmetric pure-strategy equilibria and a symmetric mixed-strategy equilibrium. In solving for the mixed-strategy equilibrium, let us move from two players to N players. In the N -player version of the game, the payoff to Smith is

0 if nobody calls, 7 if he himself calls, and 10 if one or more of the other $N - 1$ players calls. This game also has asymmetric pure-strategy and a symmetric mixed-strategy equilibrium. If all players use the same probability γ of *Ignore*, the probability that the other $N - 1$ players besides Smith all choose *Ignore* is γ^{N-1} , so the probability that one or more of them chooses *Telephone* is $1 - \gamma^{N-1}$. Thus, equating Smith's pure-strategy payoffs using the payoff-equating method of equilibrium calculation yields

$$\pi_{Smith}(Telephone) = 7 = \pi_{Smith}(Ignore) = \gamma^{N-1}(0) + (1 - \gamma^{N-1})(10). \quad (22)$$

Equation (22) tells us that

$$\gamma^{N-1} = 0.3 \quad (23)$$

and

$$\gamma^* = 0.3^{\frac{1}{N-1}}. \quad (24)$$

If $N = 2$, Smith chooses *Ignore* with a probability of 0.30. As N increases, Smith's expected payoff remains equal to 7 whether $N = 2$ or $N = 38$, since his expected payoff equals his payoff from the pure strategy of *Telephone*. The probability of *Ignore*, γ^* , however, increases with N . If $N = 38$, the value of γ^* is about 0.97. When there are more players, each player relies more on somebody else calling.

The probability that nobody calls is γ^{*N} . Equation (23) shows that $\gamma^{*N-1} = 0.3$, so $\gamma^{*N} = 0.3\gamma^*$, which is increasing in N because γ^* is increasing in N . If $N = 2$, the probability that neither player phones the police is $\gamma^{*2} = 0.09$. When there are 38 players, the probability rises to γ^{*38} , about 0.29. The more people that watch a crime, the less likely it is to be reported.

As in the Prisoner's Dilemma, the disappointing result in the Civic Duty Game suggests a role for active policy. The mixed-strategy outcome is clearly bad. The expected payoff per player remains equal to 7 whether there is 1 player or 38, whereas if the equilibrium played out was the equilibrium in which one and only one player called the police, the average payoff would rise from 7 with 1 player to about 9.9 with 38 ($= [1(7) + 37(10)]/38$). A situation like this requires something to make one of the pure-strategy equilibria a focal point. The problem is divided responsibility. One person must be made responsible for calling the police, whether by tradition (e.g., the oldest person on the block always calls the police) or direction (e.g., Smith shouts to Jones: "Call the police!").

*3.4 Randomizing Is Not Always Mixing: The Auditing Game

The next three games will illustrate the difference between mixed strategies and random actions, a subtle but important distinction. In all three games, the Internal Revenue Service must decide whether to audit a certain class of suspect tax returns to discover whether they are accurate or not. The goal of the IRS is to either prevent or catch cheating at minimum cost. The suspects want to cheat only if they will not be caught. Let us assume that the benefit of preventing or catching cheating is 4, the cost of auditing is C , where $C < 4$, the cost to the suspects of obeying the law is 1, and the cost of being caught is the fine $F > 1$.

Even with all of this information, there are several ways to model the situation. Table 8 shows one way: a 2-by-2 simultaneous-move game.

Table 8: Auditing Game I

		Suspects	
		<i>Cheat</i> (θ)	<i>Obey</i> ($1 - \theta$)
IRS:	<i>Audit</i> (γ)	$4 - C, -F$	\rightarrow $4 - C, -1$
	<i>Trust</i> ($1 - \gamma$)	$0, 0$	\leftarrow $4, -1$

Payoffs to: (IRS, Suspects). Arrows show how a player can increase his payoff.

Auditing Game I is a discoordination game, with only a mixed strategy equilibrium. Equations (18) and (20) or the payoff- equating method tell us that

$$\begin{aligned} \text{Probability}(\text{Cheat}) = \theta^* &= \frac{4-(4-C)}{(4-(4-C))+((4-C)-0)} \\ &= \frac{C}{4} \end{aligned} \tag{25}$$

and

$$\begin{aligned} \text{Probability}(\text{Audit}) = \gamma^* &= \frac{-1-0}{(-1-0)+(-F--1)} \\ &= \frac{1}{F}. \end{aligned} \tag{26}$$

Using (25) and (26), the payoffs are

$$\begin{aligned} \pi_{IRS}(\text{Audit}) = \pi_{IRS}(\text{Trust}) &= \theta^*(0) + (1 - \theta^*)(4) \\ &= 4 - C. \end{aligned} \tag{27}$$

and

$$\begin{aligned} \pi_{Suspect}(\text{Obey}) = \pi_{Suspect}(\text{Cheat}) &= \gamma^*(-F) + (1 - \gamma^*)(0) \\ &= -1. \end{aligned} \tag{28}$$

A second way to model the situation is as a sequential game. Let us call this Auditing Game II. The simultaneous game implicitly assumes that both players choose their actions without knowing what the other player has decided. In the sequential game, the IRS chooses government policy first, and the suspects react to it. The equilibrium in Auditing Game II is in pure strategies, a general feature of sequential games of perfect information. In equilibrium, the IRS chooses *Audit*, anticipating that the suspect will then choose *Obey*. The payoffs are $(4 - C)$ for the IRS and -1 for the suspects, the same for both players as in Auditing Game I, although now there is more auditing and less cheating and fine-paying.

We can go a step further. Suppose the IRS does not have to adopt a policy of auditing or trusting every suspect, but instead can audit a random sample. This is not necessarily a mixed strategy. In Auditing Game I, the equilibrium strategy was to audit all suspects with probability $1/F$ and none of them otherwise. That is different from announcing in advance that the IRS will audit a random sample of $1/F$ of the suspects. For Auditing Game III, suppose the IRS move first, but let its move consist of the choice of the proportion α of tax returns to be audited.

We know that the IRS is willing to deter the suspects from cheating, since it would be willing to choose $\alpha = 1$ and replicate the result in Auditing Game II if it had to. It chooses α so that

$$\pi_{suspect}(Obey) \geq \pi_{suspect}(Cheat), \quad (29)$$

i.e.,

$$-1 \geq \alpha(-F) + (1 - \alpha)(0). \quad (30)$$

In equilibrium, therefore, the IRS chooses $\alpha = 1/F$ and the suspects respond with *Obey*. The IRS payoff is $(4 - \alpha C)$, which is better than the $(4 - C)$ in the other two games, and the suspect's payoff is -1 , exactly the same as before.

The equilibrium of Auditing Game III is in pure strategies, even though the IRS's action is random. It is different from Auditing Game I because the IRS must go ahead with the costly audit even if the suspect chooses *Obey*. Auditing Game III is different in another way also: its action set is continuous. In Auditing Games I and Auditing Game II the action set is $\{Audit, Trust\}$, although the strategy set becomes $\gamma \in [0, 1]$ once mixed strategies are allowed. In Auditing Game III, the action set is $\alpha \in [0, 1]$, and the strategy set would allow mixing of any of the elements in the action set, although mixed strategies are pointless for the IRS because the game is sequential.

Games with mixed strategies are like games with continuous strategies since a probability is drawn from the continuum between zero and one. Auditing Game III also has a strategy drawn from the interval between zero and one, but it is not a mixed strategy to pick an audit probability of, say, 70 percent. An example of a mixed strategy would be the choice of a probability 0.5 of an audit probability of 60 percent and 0.5 of 80 percent. The big difference between the pure strategy choice of an audit probability of 0.70 and the mixed strategy choice of (0.5— 60%audit, 0.5— 80% audit), both of which yield an audit probability of 70%, is that the pure strategy is an irreversible choice that might be used even when the player is not indifferent between pure strategies, but the mixed strategy is the result of a player who in equilibrium is indifferent as to what he does. The next section will show another difference between mixed strategies and continuous strategies: the payoffs are linear in the mixed-strategy probability, as is evident from payoff equations (15) and (16), but they can be nonlinear in continuous strategies generally.

I have used auditing here mainly to illustrate what mixed strategies are and are not, but auditing is interesting in itself and optimal auditing schemes have many twists to them. An example is the idea of **cross-checking**. Suppose an auditor is supposed to check the value of some variable $x \in [0, 1]$, but his employer is worried that he will not report the true value. This might be because the auditor will be lazy and guess rather than go to the effort of finding x , or because some third party will bribe him, or that certain values of x will trigger punishments or policies the auditor dislikes (this model applies even if x is the auditor's own performance on some other task). The idea of cross-checking is to hire a second auditor and ask him to simultaneously report x . Then, if both auditors report the same x , they are both rewarded, but if they report different values they are both punished. There will still be multiple equilibria, because anything in which they report the same value is an equilibrium. But at least truthful reporting becomes a possible equilibrium. See Kandori & Matsushima (1998) for details (and compare to the Maskin Matching Scheme of Chapter 10).

3.5 Continuous Strategies: The Cournot Game

Most of the games so far in the book have had discrete strategy spaces: *Aid* or *No Aid*, *Confess* or *Deny*. Quite often when strategies are discrete and moves are simultaneous, no pure-strategy equilibrium exists. The only sort of compromise possible in the Welfare Game, for instance, is to choose *Aid* sometimes and *No Aid* sometimes, a mixed strategy. If “*A Little Aid*” were a possible action, maybe there would be a pure-strategy equilibrium. The simultaneous-move game we discuss next, the Cournot Game, has a continuous strategy space even without mixing. It models a duopoly in which two firms choose output levels in competition with each other.

The Cournot Game

Players

Firms Apex and Brydox

The Order of Play

Apex and Brydox simultaneously choose quantities q_a and q_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $Q = q_a + q_b$, and we will assume it to be linear (for generalization see Chapter 14), and, in fact, will use the following specific function:

$$p(Q) = 120 - q_a - q_b. \quad (31)$$

Payoffs are profits, which are given by a firm's price times its quantity minus its costs, i.e.,

$$\begin{aligned} \pi_{Apex} &= (120 - q_a - q_b)q_a - cq_a = (120 - c)q_a - q_a^2 - q_aq_b; \\ \pi_{Brydox} &= (120 - q_a - q_b)q_b - cq_b = (120 - c)q_b - q_aq_b - q_b^2. \end{aligned} \quad (32)$$




Figure 2: Reaction Curves in the Cournot Game

If this game were cooperative (see Section 1.2), firms would end up producing somewhere on the 45° line in Figure 2, where total output is the monopoly output and maximizes the sum of the payoffs. The monopoly output maximizes $pQ - cQ = (120 - Q - c)Q$ with respect to the total output of Q , resulting in the first-order condition

$$120 - c - 2Q = 0, \quad (33)$$

which implies a total output of $Q = 54$ and a price of 66. Deciding how much of that output of 54 should be produced by each firm—where the firm’s output should be located on the 45° line—would be a zero-sum cooperative game, an example of bargaining. But since the Cournot Game is noncooperative, the strategy profiles such that $q_a + q_b = 54$ are not necessarily equilibria despite their Pareto optimality (where Pareto optimality is defined from the point of view of the two players, not of consumers, and under the implicit assumption that price discrimination cannot be used).

Cournot noted in Chapter 7 of his 1838 book that this game has a unique equilibrium when demand curves are linear. To find that “Cournot-Nash” equilibrium, we need to refer to the **best-response functions** for the two players. If Brydox produced 0, Apex would produce the monopoly output of 54. If Brydox produced $q_b = 108$ or greater, the market

price would fall to 12 and Apex would choose to produce zero. The best response function is found by maximizing Apex's payoff, given in equation (32), with respect to his strategy, q_a . This generates the first-order condition $120 - c - 2q_a - q_b = 0$, or

$$q_a = 60 - \left(\frac{q_b + c}{2} \right) = 54 - \left(\frac{1}{2} \right) q_b. \quad (34)$$

Another name for the best response function, the name usually used in the context of the Cournot Game, is the **reaction function**. Both names are somewhat misleading since the players move simultaneously with no chance to reply or react, but they are useful in imagining what a player would do if the rules of the game did allow him to move second. The reaction functions of the two firms are labelled R_a and R_b in Figure 2. Where they cross, point E, is the **Cournot-Nash equilibrium**, which is simply the Nash equilibrium when the strategies consist of quantities. Algebraically, it is found by solving the two reaction functions for q_a and q_b , which generates the unique equilibrium, $q_a = q_b = 40 - c/3 = 36$. The equilibrium price is then 48 ($= 120 - 36 - 36$).

In the Cournot Game, the Nash equilibrium has the particularly nice property of **stability**: we can imagine how starting from some other strategy profile the players might reach the equilibrium. If the initial strategy profile is point X in Figure 2, for example, Apex's best response is to decrease q_a and Brydox's is to increase q_b , which moves the profile closer to the equilibrium. But this is special to The Cournot Game, and Nash equilibria are not always stable in this way.

Stackelberg Equilibrium

There are many ways to model duopoly. The three most prominent are Cournot, Stackelberg, and Bertrand. Stackelberg equilibrium differs from Cournot in that one firm gets to choose its quantity first. If Apex moved first, what output would it choose? Apex knows how Brydox will react to its choice, so it picks the point on Brydox's reaction curve that maximizes Apex's profit (see Figure 3).

The Stackelberg Game

Players

Firms Apex and Brydox

The Order of Play

- 1 Apex chooses quantity q_a from the set $[0, \infty)$.
- 2 . Brydox chooses quantity q_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $Q = q_a + q_b$:

$$p(Q) = 120 - q_a - q_b. \quad (35)$$

Payoffs are profits, which are given by a firm's price times its quantity minus its costs, i.e.,

$$\begin{aligned}\pi_{Apex} &= (120 - q_a - q_b)q_a - cq_a = (120 - c)q_a - q_a^2 - q_aq_b; \\ \pi_{Brydox} &= (120 - q_a - q_b)q_b - cq_b = (120 - c)q_b - q_aq_b - q_b^2.\end{aligned}\tag{36}$$




Figure 3: Stackelberg Equilibrium

Apex, moving first, is called the **Stackelberg leader** and Brydox is the **Stackelberg follower**. The distinguishing characteristic of a Stackelberg equilibrium is that one player gets to commit himself first. In Figure 3, Apex moves first intertemporally. If moves were simultaneous but Apex could commit himself to a certain strategy, the same equilibrium would be reached as long as Brydox was not able to commit himself. Algebraically, since Apex forecasts Brydox's output to be $q_b = 60 - \frac{q_a + c}{2}$ from the analog of equation (34), Apex can substitute this into his payoff function in (32) to obtain

$$\pi_a = (120 - c)q_a - q_a^2 - q_a\left(60 - \frac{q_a + c}{2}\right).\tag{37}$$

Maximizing his payoff with respect to q_a yields the first-order condition

$$(120 - c) - 2q_a - 60 + q_a + \frac{c}{2} = 0,\tag{38}$$

which generates Apex's reaction function, $q_a = 60 - c/2 = 54$ (which only equals the monopoly output by coincidence, due to the particular numbers in this example). Once Apex chooses this output, Brydox chooses his output to be $q_b = 27$. (That Brydox chooses exactly half the monopoly output is also accidental.) The market price is $120 - 54 - 27 = 39$ for both firms, so Apex has benefited from his status as Stackelberg leader, but industry profits have fallen compared to the Cournot equilibrium.

3.6 Continuous Strategies: The Bertrand Game, Strategic Complements, and Strategic Substitutes

A natural alternative to a duopoly model in which the two firms pick outputs simultaneously is a model in which they pick prices simultaneously. This is known as **Bertrand equilibrium**, because the difficulty of choosing between the two models was stressed in Bertrand (1883), a review discussion of Cournot's book. We will use the same two-player linear-demand world as before, but now the strategy spaces will be the prices, not the quantities. We will also use the same demand function, equation (31), which implies that if p is the lowest price, $q = 120 - p$. In the Cournot model, firms chose quantities but allowed the market price to vary freely. In the Bertrand model, they choose prices and sell as much as they can.

The Bertrand Game

Players

Firms Apex and Brydox

The Order of Play

Apex and Brydox simultaneously choose prices p_a and p_b from the set $[0, \infty)$.

Payoffs

Marginal cost is constant at $c = 12$. Demand is a function of the total quantity sold, $Q(p) = 120 - p$. The payoff function for Apex (Brydox's would be analogous) is

$$\pi_a = \begin{cases} (120 - p_a)(p_a - c) & \text{if } p_a \leq p_b \\ \frac{(120 - p_a)(p_a - c)}{2} & \text{if } p_a = p_b \\ 0 & \text{if } p_a > p_b \end{cases}$$

The Bertrand Game has a unique Nash equilibrium: $p_a = p_b = c = 12$, with $q_a = q_b = 54$. That this is a weak Nash equilibrium is clear: if either firm deviates to a higher price, it loses all its customers and so fails to increase its profits to above zero. In fact, this is an example of a Nash equilibrium in weakly dominated strategies. That the equilibrium is unique is less clear. To see why it is, divide the possible strategy profiles into four groups:

$p_a < c$ or $p_b < c$. In either of these cases, the firm with the lowest price will earn negative profits, and could profitably deviate to a price high enough to reduce its demand to zero.

$p_a > p_b > c$ or $p_b > p_a > c$. In either of these cases the firm with the higher price could deviate to a price below its rival and increase its profits from zero to some positive value.

$p_a = p_b > c$. In this case, Apex could deviate to a price ϵ less than Brydox and its profit would rise, because it would go from selling half the market quantity to selling all of it with an infinitesimal decline in profit per unit sale.

$p_a > p_b = c$ or $p_b > p_a = c$. In this case, the firm with the price of c could move from zero profits to positive profits by increasing its price slightly while keeping it below the other firm's price.

This proof is a good example of one common method of proving uniqueness of equilibrium in game theory: partition the strategy profile space and show area by area that deviations would occur. It is such a good example that I recommend it to anyone teaching from this book as a good test question.²

Like the surprising outcome of Prisoner's Dilemma, the Bertrand equilibrium is less surprising once one thinks about the model's limitations. What it shows is that duopoly profits do not arise just because there are two firms. Profits arise from something else, such as multiple periods, incomplete information, or differentiated products.

Both the Bertrand and Cournot models are in common use. The Bertrand model can be awkward mathematically because of the discontinuous jump from a market share of 0 to 100 percent after a slight price cut. The Cournot model is useful as a simple model that avoids this problem and which predicts that the price will fall gradually as more firms enter the market. There are also ways to modify the Bertrand model to obtain intermediate prices and gradual effects of entry. Let us proceed to look at one such modification.

The Differentiated Bertrand Game

The Bertrand model generates zero profits because only slight price discounts are needed to bid away customers. The assumption behind this is that the two firms sell identical goods, so if Apex's price is slightly higher than Brydox's all the customers go to Brydox. If customers have brand loyalty or poor price information, the equilibrium is different. Let us now move to a different duopoly market, where the demand curves facing Apex and Brydox are

$$q_a = 24 - 2p_a + p_b \quad (39)$$

and

$$q_b = 24 - 2p_b + p_a, \quad (40)$$

and they have constant marginal costs of $c = 3$.

²Is it still a good question given that I have just provided a warning to the students? Yes. First, it will prove a filter for discovering which students have even skimmed the assigned reading. Second, questions like this are not always easy even if one knows they are on the test. Third, and most important, even if in equilibrium every student answers the question correctly, that very fact shows that the incentive to learn this particular item has worked – and that is our main goal, is it not?

The greater the difference in the coefficients on prices in a demand curve like (39) or (40), the less substitutable are the products. As with standard demand functions such as equation (31), we have made implicit assumptions about the extreme points of equations (39) and (40). These equations only apply if the quantities demanded turn out to be nonnegative, and we might also want to restrict them to prices below some ceiling, since otherwise the demand facing one firm becomes infinite as the other's price rises to infinity. A sensible ceiling here is 12, since if $p_a > 12$ and $p_b = 0$, equation (39) would yield a negative quantity demanded for Apex. Keeping in mind these limitations, the payoffs are

$$\pi_a = (24 - 2p_a + p_b)(p_a - c) \quad (41)$$

and

$$\pi_b = (24 - 2p_b + p_a)(p_b - c). \quad (42)$$

The order of play is the same as in the Bertrand Game (or Undifferentiated Bertrand Game, as we will call it when that is necessary to avoid confusion): Apex and Brydax simultaneously choose prices p_a and p_b from the set $[0, \infty)$.

Maximizing Apex's payoff by choice of p_a , we obtain the first-order condition,

$$\frac{d\pi_a}{dp_a} = 24 - 4p_a + p_b + 2c = 0, \quad (43)$$

and the reaction function,

$$p_a = 6 + \left(\frac{1}{2}\right)c + \left(\frac{1}{4}\right)p_b = 7.5 + \left(\frac{1}{4}\right)p_b. \quad (44)$$

Since Brydax has a parallel first-order condition, the equilibrium occurs where $p_a = p_b = 10$. The quantity each firm produces is 14, which is below the 21 each would produce at prices of $p_a = p_b = c = 3$. Figure 4 shows that the reaction functions intersect. Apex's demand curve has the elasticity

$$\left(\frac{\partial q_a}{\partial p_a}\right) \cdot \left(\frac{p_a}{q_a}\right) = -2 \left(\frac{p_a}{q_a}\right), \quad (45)$$

which is finite even when $p_a = p_b$, unlike in the undifferentiated-goods Bertrand model.




Figure 4: Bertrand Reaction Functions with Differentiated Products

The differentiated-good Bertrand model is important because it is often the most descriptively realistic model of a market. A basic idea in marketing is that selling depends on “The Four P’s”: Product, Place, Promotion, and Price. Economists have concentrated heavily on price differences between products, but we realize that differences in product quality and characteristics, where something is sold, and how the sellers get information about it to the buyers also matter. Sellers use their prices as control variables more often than their quantities, but the seller with the lowest price does not get all the customers.

Why, then, did I bother to even describe the Cournot and undifferentiated Bertrand models? Aren’t they obsolete? No, because descriptive realism is not the *summum bonum* of modelling. Simplicity matters a lot too. The Cournot and undifferentiated Bertrand models are simpler, especially when we go to three or more firms, so they are better models in many applications.

Strategic Substitutes and Strategic Complements

You may have noticed an interesting difference between the Cournot and Differentiated Bertrand reaction curves in Figures 2 and 4: the reaction curves have opposite slopes. Figure 5 puts the two together for easier comparison.

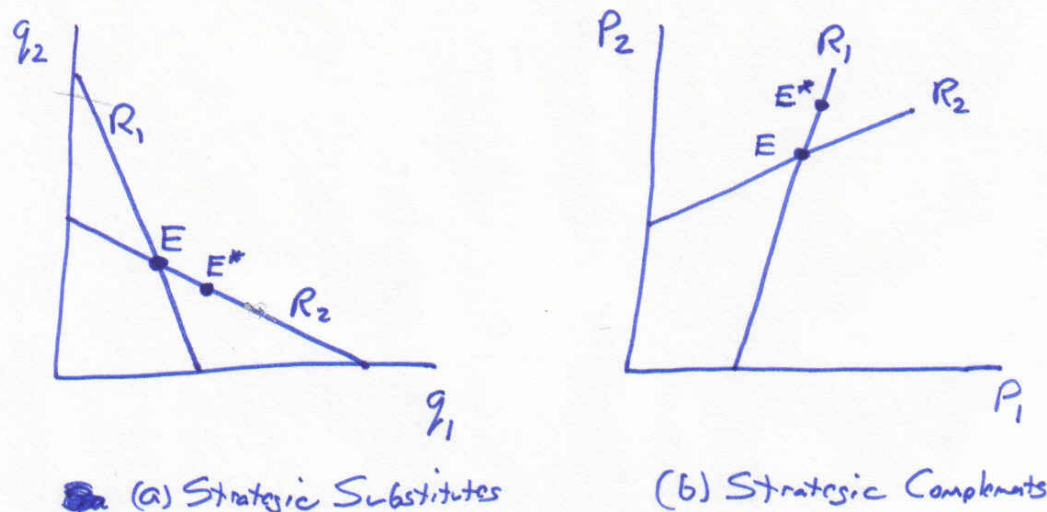


Figure 5: Cournot vs. Differentiated Bertrand Reaction Functions (Strategic Substitutes vs. Strategic Complements)

In both models, the reaction curves cross once, so there is a unique Nash equilibrium. Off the equilibrium path, though, there is an interesting difference. If a Cournot firm increases its output, its rival will do the opposite and reduce its output. If a Bertrand firm increases its price, its rival will do the same thing, and increase its price too.

We can ask of any game: “If the other players do more of their strategy, will I do more of my own strategy, or less?” In some games, the answer is “do more” and in others it is “do less”. Jeremy Bulow, John Geanakoplos & Paul Klemperer (1985) apply the term “strategic complements” to the strategies in the “do more” kind of game, because when Player 1 does more of his strategy that increases Player 2’s marginal payoff from 2’s strategy, just as when I buy more bread it increases my marginal utility from buying more butter. If strategies are strategic complements, their reaction curves are upward sloping, as in the Differentiated Bertrand Game.

On the other hand, in the “do less” kind of game, when Player 1 does more of his strategy that *reduces* Player 2’s marginal payoff from 2’s strategy, just as my buying potato chips reduces my marginal utility from buying more corn chips. The strategies are therefore “strategic substitutes” and their reaction curves are downward sloping, as in the Cournot Game.

Which way the reaction curves slope also affects whether a player wants to move first or second. Esther Gal-Or (1985) notes that if reaction curves slope down (as with strategic substitutes and Cournot) there is a first-mover advantage, whereas if they slope upwards (as with strategic complements and Differentiated Bertrand) there is a second-mover advantage.

We can see that in Figure 5, the Cournot Game in which Player 1 moves first is simply the Stackelberg Game, which we have already analyzed using Figure 3. The equilibrium moves from E to E^* in Figure 5a, Player 1's payoff increases, and Player 2's payoff falls. Note, too, that the total industry payoff is lower in Stackelberg than in Cournot— not only does one player lose, but he loses more than the other player gains.

We have not analyzed the Differentiated Bertrand Game when Player 1 moves first, but since price is a strategic complement, the effect of sequentiality is very different from in the Cournot Game (and, actually, from the sequential undifferentiated Bertrand Game— see the end-of-chapter notes). We cannot tell what Player 1's optimal strategy is from the diagram alone, but Figure 5 illustrates one possibility. Player 1 chooses a price p^* higher than he would in the simultaneous-move game, predicting that Player 2's response will be a price somewhat lower than p^* , but still greater than the simultaneous Bertrand price at E . The result is that Player 2's payoff is higher than Player 1's—a second-mover advantage. Note, however, that both players are better off at E^* than at E , so both players would favor converting the game to be sequential.

Both sequential games could be elaborated further by adding moves beforehand which would determine which player would choose his price or quantity first, but I will leave that to you. The important point for now is that whether a game has strategic complements or strategic substitutes is hugely important to the incentives of the players.

The point is simple enough and important enough that I devote an entire session of my MBA game theory course to strategic complements and strategic substitutes. In the practical game theory that someone with a Master of Business Administration degree ought to know, the most important thing is to learn how to describe a situation in terms of players, actions, information, and payoffs. Often there is not enough data to use a specific functional form, but it is possible to figure out with a mixture of qualitative and quantitative information whether the relation between actions and payoffs is one of strategic substitutes or strategic complements. The businessman then knows whether, for example, he should try to be a first mover or a second mover, and whether he should keep his action secret or proclaim his action to the entire world.

To understand the usefulness of the idea of strategic complements and substitutes, think about how you would model situations like the following (note that there is no universally right answer for any of them):

1. Two firms are choosing their research and development budgets. Are the budgets strategic complements or strategic substitutes?
2. Smith and Jones are both trying to be elected President of the United States. Each must decide how much he will spend on advertising in California. Are the advertising budgets strategic complements or strategic substitutes?
3. Seven firms are each deciding whether to make their products more special, or more suited to the average consumer. Is the degree of specialness a strategic complement or a strategic substitute?
4. India and Pakistan are each deciding whether to make their armies larger or smaller.

Is army size a strategic complement or a strategic substitute?

Economists have a growing appreciation of how powerful the ideas of substitution and complementarity can be in thinking about the deep structure of economic behavior. The mathematical idea of supermodularity, to be discussed in Chapter 14, is all about complementarity. For an inspiring survey, see Vives (2005).

***3.7 Existence of Equilibrium**

One of the strong points of Nash equilibria is that they exist in practically every game one is likely to encounter. There are four common reasons why an equilibrium might not exist or might only exist in mixed strategies.

(1) An unbounded strategy space

Suppose in a stock market game that Smith can borrow money and buy as many shares x of stock as he likes, so his strategy set, the amount of stock he can buy, is $[0, \infty)$, a set which is unbounded above. (Note, by the way, that we thus assume that he can buy fractional shares, e.g. $x = 13.4$, but cannot sell short, e.g. $x = -100$.)

If Smith knows that the price is lower today than it will be tomorrow, his payoff function will be $\pi(x) = x$ and he will want to buy an infinite number of shares, which is not an equilibrium purchase. If the amount he buys is restricted to be less than or equal to 1,000, however, then the strategy set is bounded (by 1,000), and an equilibrium exists— $x = 1,000$.

Sometimes, as in the Cournot Game discussed earlier in this chapter, the unboundedness of the strategy sets does not matter because the optimum is an interior solution. In other games, though, it is important, not just to get a determinate solution but because the real world is a rather bounded place. The solar system is finite in size, as is the amount of human time past and future.

(2) An open strategy space

Again consider Smith. Let his strategy be $x \in [0, 1,000)$, which is the same as saying that $0 \leq x < 1,000$, and his payoff function be $\pi(x) = x$. Smith's strategy set is bounded (by 0 and 1,000), but it is open rather than closed, because he can choose any number less than 1,000, but not 1,000 itself. This means no equilibrium will exist, because he wants to buy 999.999... shares. This is just a technical problem; we ought to have specified Smith strategy space to be $[0, 1,000]$, and then an equilibrium would exist, at $x = 1,000$.

(3) A discrete strategy space (or, more generally, a nonconvex strategy space)

Suppose we start with an arbitrary pair of strategies s_1 and s_2 for two players. If the players' strategies are strategic complements, then if player 1 increases his strategy in response to s_2 , then player 2 will increase his strategy in response to that. An equilibrium

will occur where the players run into diminishing returns or increasing costs, or where they hit the upper bounds of their strategy sets. If, on the other hand, the strategies are strategic substitutes, then if player 1 increases his strategy in response to s_2 , player 2 will in turn want to reduce his strategy. If the strategy spaces are continuous, this can lead to an equilibrium, but if they are discrete, player 2 cannot reduce his strategy just a little bit— he has to jump down a discrete level. That could then induce Player 1 to increase his strategy by a discrete amount. This jumping of responses can be never- ending—there is no equilibrium.

That is what is happening in the Welfare Game of Table 1 in this chapter. No compromise is possible between a little aid and no aid, or between working and not working— until we introduce mixed strategies. That allows for each player to choose a continuous amount of his strategy.

This problem is not limited to games such as 2-by-2 games that have discrete strategy spaces. Rather, it is a problem of “gaps” in the strategy space. Suppose we had a game in which the government was not limited to amount 0 or 100 of aid, but could choose any amount in the space $\{[0, 10], [90, 100]\}$. That is a continuous, closed, and bounded strategy space, but it is non-convex— there is gap in it. (For a space $\{x\}$ to be convex, it must be true that if x_1 and x_2 are in the space, so is $\theta x_1 + (1 - \theta)x_2$ for any $\theta \in [0, 1]$.) Without mixed strategies, an equilibrium to the game might well not exist.

(4) A discontinuous reaction function arising from nonconcave or discontinuous payoff functions

Even if the strategy spaces are closed, bounded, and convex, a problem remains. For a Nash equilibrium to exist, we need for the reaction functions of the players to intersect. If the reaction functions are discontinuous, they might not intersect.

Figure 6 shows this for a two-player game in which each player chooses a strategy from the interval between 0 and 1. Player 1’s reaction function, $s_1(s_2)$, must pick one or more value of s_1 for each possible value of s_2 , so it must cross from the bottom to the top of the diagram. Player 2’s reaction function, $s_2(s_1)$, must pick one or more value of s_2 for each possible value of s_1 , so it must cross from the left to the right of the diagram. If the strategy sets were unbounded or open, the reaction functions might not exist, but that is not a problem here: they do exist. And in Panel (a) a Nash equilibrium exists, at the point, E , where the two reaction functions intersect.

In Panel (b), however, no Nash equilibrium exists. The problem is that Firm 2’s reaction function $s_2(s_1)$ is discontinuous at the point $s_1 = 0.5$. It jumps down from $s_2(0.5) = 0.6$ to $s_2(0.50001) = 0.4$. As a result, the reaction curves never intersect, and no equilibrium exists.

If the two players can use mixed strategies, then an equilibrium will exist even for the game in Panel (b), though I will not prove that here. I would, however, like to say why it is that the reaction function might be discontinuous. A player’s reaction functions, remember, is derived by maximizing his payoff as a function of his own strategy given the strategies of the other players.

Thus, a first reason why Player 1's reaction function might be discontinuous in the other players' strategies is that his payoff function is discontinuous in either his own or the other players' strategies. This is what happens in Chapter 14's Hotelling Pricing Game, where if Player 1's price drops enough (or Player 2's price rises high enough), all of Player 2's customers suddenly rush to Player 1.

A second reason why Player 1's reaction function might be discontinuous in the other players' strategies is that his payoff function is not concave. The intuition is that if an objective function is not concave, then there might be a number of maxima that are local but not global, and as the parameters change, which maximum is the global one can suddenly change. This means that the reaction function will suddenly jump from one maximizing choice to another one that is far-distant, rather than smoothly changing as it would in a more nicely behaved problem.




Figure 6: Continuous and Discontinuous Reaction Functions

Problems (1) and (2) are really problems in decision theory, not game theory, because unboundedness and openness lead to nonexistence of the solution to even a one-player maximization problem. Problems (3) and (4) are special to game theory. They arise because although each player has a best response to the other players, no profile of best choices is such that everybody has chosen his best response to everybody else. They are similar to the decision theory problem of nonexistence of an interior solution, but if only one player were involved, we would at least have a corner solution.

In this chapter, I have introduced a number of seemingly disparate ideas— mixed strategies, auditing, continuous strategy spaces, reaction curves, complementary substitutes and complements, existence of equilibrium... What ties them together? The unifying theme is the possibility of reaching equilibrium by small changes in behavior, whether that be by changing the probability in a mixed strategy or an auditing game or by changing the level of a continuous price or quantity. Continuous strategies free us from the need to use n -by- n

tables to predict behavior in games, and with a few technical assumptions they guarantee we will find equilibria.

NOTES

N3.1 Mixed Strategies: The Welfare Game

- Waldegrave (1713) is a very early reference to mixed strategies.
- Mixed strategies come up constantly in recreational games. The Scissors- Paper-Stone choosing game, for example, has a unique mixed-strategy equilibrium, as shown in Fisher & Ryan (1992). There are usually analogs in other spheres; it turns out that three types of Californian side-blotched lizard males play the same game, as reported in Sinevero & Lively (1996). It is interesting to try to see how closely players come to the theoretically optimal mixing strategies. Chiappori, Levitt & Groseclose (2002) conclude that players choosing whether to kick right or left in soccer penalty kicks are following optimal mixed strategies, but that kickers are heterogeneous in their abilities to kick in each direction.
- The January 1992 issue of *Rationality and Society* is devoted to attacks on and defenses of the use of game theory in the social sciences, with considerable discussion of mixed strategies and multiple equilibria. Contributors include Harsanyi, Myerson, Rapaport, Tullock, and Wildavsky. The Spring 1989 issue of the *RAND Journal of Economics* also has an exchange on the use of game theory, between Franklin Fisher and Carl Shapiro. I also recommend the Peltzman (1991) attack on the game theory approach to industrial organization in the spirit of the “Chicago School”.
- In this book it will always be assumed that players remember their previous moves. Without this assumption of **perfect recall**, the definition in the text is not that for a mixed strategy, but for a **behavior strategy**. As historically defined, a player pursues a mixed strategy when he randomly chooses between pure strategies at the starting node, but he plays a pure strategy thereafter. Under that definition, the modeller cannot talk of random choices at any but the starting node. Kuhn (1953) showed that the definition of mixed strategy given in the text is equivalent to the original definition if the game has perfect recall. Since all important games have perfect recall and the new definition of mixed strategy is better in keeping with the modern spirit of sequential rationality, I have abandoned the old definition.

The classic example of a game without perfect recall is **bridge**, where the four players of the actual game can be cutely modelled as two players who forget what half their cards look like at any one time in the bidding. A more useful example is a game that has been simplified by restricting players to Markov strategies (see Section 5.4), but usually the modeller sets up such a game with perfect recall and then rules out non-Markov equilibria after showing that the Markov strategies form an equilibrium for the general game.

- It is *not* true that when two pure-strategy equilibria exist a player would be just as willing to use a strategy mixing the two even when the other player is using a pure strategy. In Battle of the Sexes, for instance, if the man knows the woman is going to the ballet he is not indifferent between the ballet and the prize fight.
- A continuum of players is useful not only because the modeller need not worry about fractions of players, but because he can use more modelling tools from calculus— taking the integral of the quantities demanded by different consumers, for example, rather than the sum. But using a continuum is also mathematically more difficult: see Aumann (1964a, 1964b).
- There is an entire literature on the econometrics of estimating game theory models. Suppose we would like to estimate the payoff numbers in a 2-by-2 game, where we observe the actions

taken by each of the two players and various background variables. The two actions might be, for example, to enter or not enter, and the background variables might be such things as the size of the market or the cost conditions facing one of the players. We will of course need multiple repetitions of the situation to generate enough data to use econometrics. There is an identification problem, because there are eight payoffs in a 2-by-2 payoff matrix, but only four possible action profiles— and if mixed strategies are being used, the four mixing probabilities have to add up to one, so there are really only three independent observed outcomes. How can we estimate 8 parameters with only 3 possible outcomes? For identification, it must be that some environmental variables affect only one of the players, as Bajari, Hong & Ryan (2004) note. In addition, there is the problem that there may be multiple equilibria being played out, so that additional identifying assumptions are needed to help us know which equilibria are being played out in which observations. The foundational articles in this literature are Bresnahan & Reiss (1990, 1991a), and it is an active area of research.

N3.2 The Payoff-Equating Method and Games of Timing

- The game of Chicken discussed in the text is simpler than the game acted out in the movie *Rebel Without a Cause*, in which the players race towards a cliff and the winner is the player who jumps out of his car last. The pure-strategy space in the movie game is continuous and the payoffs are discontinuous at the cliff’s edge, which makes the game more difficult to analyze technically. (Recall, too, the importance in the movie of a disastrous mistake—the kind of “tremble” that Section 4.1 will discuss.)
- Technical difficulties arise in some models with a continuum of actions and mixed strategies. In the Welfare Game, the government chose a single number, a probability, on the continuum from zero to one. If we allowed the government to mix over a continuum of aid levels, it would choose a function, a probability density, over the continuum. The original game has a finite number of elements in its strategy set, so its mixed extension still has a strategy space in \mathbf{R}^n . But with a continuous strategy set extended by a continuum of mixed strategies for each pure strategy, the mathematics become difficult. A finite number of mixed strategies can be allowed without much problem, but usually that is not satisfactory.

Games in continuous time frequently run into this problem. Sometimes it can be avoided by clever modelling, as in Fudenberg & Tirole’s (1986b) continuous-time war of attrition with asymmetric information. They specify as strategies the length of time firms would proceed to *Continue* given their beliefs about the type of the other player, in which case there is a pure strategy equilibrium.

- **Differential games** are played in continuous time. The action is a function describing the value of a state variable at each instant, so the strategy maps the game’s past history to such a function. Differential games are solved using dynamic optimization. A book-length treatment is Bagchi (1984).
- Fudenberg & Levine (1986) show circumstances under which the equilibria of games with infinite strategy spaces can be found as the limits of equilibria of games with finite strategy spaces.

N3.4 Randomizing versus Mixing: The Auditing Game

- Auditing Game I is similar to a game called The Police Game. Care must be taken in such games that one does not use a simultaneous- move game when a sequential game is appropriate. Also, discrete strategy spaces can be misleading. In general, economic analysis assumes that costs rise convexly in the amount of an activity and benefits rise concavely. Modelling a situation with a 2-by-2 game uses just two discrete levels of the activity, so the concavity or convexity is lost in the simplification. If the true functions are linear, as in auditing costs which rise linearly with the probability of auditing, this is no great loss. If the true costs rise convexly, as in the case where the hours a policeman must stay on the street each day are increased, then a 2-by-2 model can be misleading. Be especially careful not to press the idea of a mixed-strategy equilibrium too hard if a pure-strategy equilibrium would exist when intermediate strategies are allowed. See Tsebelis (1989) and the criticism of it in Jack Hirshleifer & Rasmusen (1992).
- Douglas Diamond (1984) shows the implications of monitoring costs for the structure of financial markets. A fixed cost to monitoring investments motivates the creation of a financial intermediary to avoid repetitive monitoring by many investors.
- Baron & Besanko (1984) study auditing in the context of a government agency which can at some cost collect information on the true production costs of a regulated firm.
- Mookherjee & Png (1989) and Border & Sobel (1987) have examined random auditing in the context of taxation. They find that if a taxpayer is audited he ought to be more than compensated for his trouble if it turns out he was telling the truth. Under the optimal contract, the truth-telling taxpayer should be delighted to hear that he is being audited. The reason is that a reward for truthfulness widens the differential between the agent's payoff when he tells the truth and when he lies.

Why is such a scheme not used? It is certainly practical, and one would think it would be popular with the voters. One reason might be the possibility of corruption; if being audited leads to a lucrative reward, the government might purposely choose to audit its friends. The current danger seems even worse, though, since the government can audit its enemies and burden them with the trouble of an audit even if they have paid their taxes properly.

- Government action strongly affects what information is available as well as what is contractible. In 1988, for example, the United States passed a law sharply restricting the use of lie detectors for testing or monitoring. Previous to the restriction, about two million workers had been tested each year. (“Law Limiting Use of Lie Detectors is Seen Having Widespread Effect” *Wall Street Journal*, p. 13, 1 July 1988), “American Polygraph Association,” <http://www.polygraph.org/betasite/menu8.html>, Eric Rasmusen, “Bans on Lie Detector Tests,” <http://mypage.iu.edu/~erasmuse/archives1.htm#august10a>.)
- Section 3.4 shows how random actions come up in auditing and in mixed strategies. Another use for randomness is to reduce transactions costs. In 1983, for example, Chrysler was bargaining over how much to pay Volkswagen for a Detroit factory. The two negotiators locked themselves into a hotel room and agreed not to leave till they had an agreement. When they narrowed the price gap from \$100 million to \$5 million, they agreed to flip a coin. (Chrysler won.) How would you model that? “Chrysler Hits Brakes, Starts Saving Money after Shopping Spree,” *Wall Street Journal*, p. 1, 12 January 1988. See also David Friedman's ingenious idea in Chapter 15 of *Law's Order* of using a 10% probability of death to replace a 6-year prison term (http://www.daviddfriedman.com/Academic/Course/Pages/L_and_E_LS_9_8/Why_Is_Law/Why_Is_Law_Chapter_15/Why_Is_Law_Chapter_15.html)

N3.5 Continuous Strategies: The Cournot Game

- An interesting class of simple continuous payoff games are the **Colonel Blotto games** (Tukey [1949], McDonald & Tukey [1949]). In these games, two military commanders allocate their forces to m different battlefields, and a battlefield contributes more to the payoff of the commander with the greater forces there. A distinguishing characteristic is that player i 's payoff increases with the value of player i 's particular action relative to player j 's, and i 's actions are subject to a budget constraint. Except for the budget constraint, this is similar to the tournaments of Section 8.2.
- “Stability” is a word used in many different ways in game theory and economics. The natural meaning of a stable equilibrium is that it has dynamics which cause the system to return to that point after being perturbed slightly, and the discussion of the stability of Cournot equilibrium was in that spirit. The uses of the term by von Neumann & Morgenstern (1944) and Kohlberg & Mertens (1986) are entirely different.
- The term “Stackelberg equilibrium” is not clearly defined in the literature. It is sometimes used to denote equilibria in which players take actions in a given order, but since that is just the perfect equilibrium (see Section 4.1) of a well-specified extensive form, I prefer to reserve the term for the Nash equilibrium of the duopoly quantity game in which one player moves first, which is the context of Chapter 3 of Stackelberg (1934).

An alternative definition is that a Stackelberg equilibrium is a strategy profile in which players select strategies in a given order and in which each player's strategy is a best response to the fixed strategies of the players preceding him and the yet-to-be-chosen strategies of players succeeding him, i.e., a situation in which players precommit to strategies in a given order. Such an equilibrium would not generally be either Nash or perfect.

- Stackelberg (1934) suggested that sometimes the players are confused about which of them is the leader and which the follower, resulting in the disequilibrium outcome called **Stackelberg warfare**.
- With linear costs and demand, total output is greater in Stackelberg equilibrium than in Cournot. The slope of the reaction curve is less than one, so Apex's output expands more than Brydoux's contracts. Total output being greater, the price is less than in the Cournot equilibrium.
- A useful application of Stackelberg equilibrium is to an industry with a dominant firm and a **competitive fringe** of smaller firms that sell at capacity if the price exceeds their marginal cost. These smaller firms act as Stackelberg leaders (not followers), since each is small enough to ignore its effect on the behavior of the dominant firm. The oil market could be modelled this way with OPEC as the dominant firm and producers such as Britain on the fringe.

N3.6 Continuous Strategies: The Bertrand Game, Strategic Complements, and Strategic Substitutes

- The text analyzed the simultaneous undifferentiated Bertrand game but not the sequential one. $p_a = p_c = c$ remains an equilibrium outcome, but it is no longer unique. Suppose Apex moves first, then Brydoux, and suppose, for a technical reason to be apparent shortly, that if $p_a = p_b$ Brydoux captures the entire market. Apex cannot achieve more than a payoff of zero, because either $p_a = c$ or Brydoux will choose $p_b = p_a$ and capture the entire market. Thus, Apex is indifferent between any $p_a \geq c$.

The game needs to be set up with this tiebreaking rule because if split the market between Apex and Brydox when $p_a = p_b$, Brydox's best response to $p_a > c$ would be to choose p_b to be the biggest number less than p_a — but with a continuous space, no such number exists, so Brydox's best response is ill- defined. Giving all the demand to Brydox in case of price ties gets around this problem.

- The demand curves (39) and (40) can be generated by a quadratic utility function. Dixit (1979) tells us that with respect to three goods 0, 1, and 2, the utility function

$$U = q_0 + \alpha_1 q_1 + \alpha_2 q_2 - \frac{1}{2} (\beta_1 q_1^2 + 2\gamma q_1 q_2 + \beta_2 q_2^2) \quad (46)$$

(where the constants $\alpha_1, \alpha_2, \beta_1$, and β_2 are positive and $\gamma^2 \leq \beta_1 \beta_2$) generates the inverse demand functions

$$p_1 = \alpha_1 - \beta_1 q_1 - \gamma q_2 \quad (47)$$

and

$$p_2 = \alpha_2 - \beta_2 q_2 - \gamma q_1. \quad (48)$$

- We can also work out the Cournot equilibrium for demand functions (39) and (40), but product differentiation does not affect it much. Start by expressing the price in the demand curve in terms of quantities alone, obtaining

$$p_a = 12 - \left(\frac{1}{2}\right) q_a + \left(\frac{1}{2}\right) p_b \quad (49)$$

and

$$p_b = 12 - \left(\frac{1}{2}\right) q_b + \left(\frac{1}{2}\right) p_a. \quad (50)$$

After substituting from (50) into (49) and solving for p_a , we obtain

$$p_a = 24 - \left(\frac{2}{3}\right) q_a - \left(\frac{1}{3}\right) q_b. \quad (51)$$

The first-order condition for Apex's maximization problem is

$$\frac{d\pi_a}{dq_a} = 24 - 3 - \left(\frac{4}{3}\right) q_a - \left(\frac{1}{3}\right) q_b = 0, \quad (52)$$

which gives rise to the reaction function

$$q_a = 15.75 - \left(\frac{1}{4}\right) q_b. \quad (53)$$

We can guess that $q_a = q_b$. It follows from (53) that $q_a = 12.6$ and the market price is 11.4. On checking, you would find this to indeed be a Nash equilibrium. But reaction function (53) has much the same shape as if there were no product differentiation, unlike when we moved from undifferentiated to differentiated Bertrand competition.

- For more on the technicalities of strategic complements and strategic substitutes, see Bulow, Geanakoplos & Klemperer (1985) and Milgrom & Roberts (1990). If the strategies are strategic complements, Milgrom & Roberts (1990) and Vives (1990) show that pure-strategy equilibria exist. These models often explain peculiar economic phenomenon nicely, as in Peter Diamond (1982) on search and business cycles and Douglas Diamond & Dybvig (1983) on bank runs. If the strategies are strategic substitutes, existence of pure-strategy equilibria is more troublesome; see Dubey, Haimanko & Zapechelnyuk (2005).

Problems

3.1. Presidential Primaries (medium)

Smith and Jones are fighting it out for the Democratic nomination for President of the United States. The more months they keep fighting, the more money they spend, because a candidate must spend one million dollars a month in order to stay in the race. If one of them drops out, the other one wins the nomination, which is worth 11 million dollars. The discount rate is r per month. To simplify the problem, you may assume that this battle could go on forever if neither of them drops out. Let θ denote the probability that an individual player will drop out each month in the mixed-strategy equilibrium.

- (a) In the mixed-strategy equilibrium, what is the probability θ each month that Smith will drop out? What happens if r changes from 0.1 to 0.15?
- (b) What are the two pure-strategy equilibria?
- (c) If the game only lasts one period, and the Republican wins the general election if both Democrats refuse to give up (resulting in Democrat payoffs of zero), what is the probability γ with which each Democrat drops out in a symmetric equilibrium?

3.2. Running from the Police (medium)

Two risk-neutral men, Schmidt and Braun, are walking south along a street in Nazi Germany when they see a single policeman coming to check their papers. Only Braun has his papers (unknown to the policeman, of course). The policeman will catch both men if both or neither of them run north, but if just one runs, he must choose which one to stop—the walker or the runner. The penalty for being without papers is 24 months in prison. The penalty for running away from a policeman is 24 months in prison, on top of the sentences for any other charges, but the conviction rate for this offense is only 25 percent. The two friends want to maximize their joint welfare, which the policeman wants to minimize. Braun moves first, then Schmidt, then the policeman.

- (a) What is the outcome matrix for outcomes that might be observed in equilibrium? (Use θ for the probability that the policeman chases the runner and γ for the probability that Braun runs.)
- (b) What is the probability that the policeman chases the runner, (call it θ^*)?
- (c) What is the probability that Braun runs, (call it γ^*)?
- (d) Since Schmidt and Braun share the same objectives, is this a cooperative game?

3.3: Uniqueness in Matching Pennies (easy)

In the game Matching Pennies, Smith and Jones each show a penny with either heads or tails up. If they choose the same side of the penny, Smith gets both pennies; otherwise, Jones gets them.

- (a) Draw the outcome matrix for Matching Pennies.
- (b) Show that there is no Nash equilibrium in pure strategies.

- (c) Find the mixed-strategy equilibrium, denoting Smith's probability of *Heads* by γ and Jones's by θ .
- (d) Prove that there is only one mixed-strategy equilibrium.

3.4. Mixed Strategies in the Battle of the Sexes (medium)

Refer back to The Battle of the Sexes and Ranked Coordination. Denote the probabilities that the man and woman pick *Prize Fight* by γ and θ .

- (a) Find an expression for the man's expected payoff.
- (b) What are the equilibrium values of γ and θ , and the expected payoffs?
- (c) Find the most likely outcome and its probability.
- (d) What is the equilibrium payoff in the mixed-strategy equilibrium for Ranked Coordination?
- (e) Why is the mixed-strategy equilibrium a better focal point in The Battle of the Sexes than in Ranked Coordination?

3.5. A Voting Paradox (medium)

Adam, Karl, and Vladimir are the only three voters in Podunk. Only Adam owns property. There is a proposition on the ballot to tax property-holders 120 dollars and distribute the proceeds equally among all citizens who do not own property. Each citizen dislikes having to go to the polling place and vote (despite the short lines), and would pay 20 dollars to avoid voting. They all must decide whether to vote before going to work. The proposition fails if the vote is tied. Assume that in equilibrium Adam votes with probability θ and Karl and Vladimir each vote with the same probability γ , but they decide to vote independently of each other.

- (a) What is the probability that the proposition will pass, as a function of θ and γ ?
- (b) What are the two possible equilibrium probabilities γ_1 and γ_2 with which Karl might vote? Why, intuitively, are there two symmetric equilibria?
- (c) What is the probability θ that Adam will vote in each of the two symmetric equilibria?
- (d) What is the probability that the proposition will pass?

3.6. Rent Seeking (hard)

I mentioned that Rogerson (1982) uses a game very similar to "Patent Race for a New Market" to analyze competition for a government monopoly franchise. See if you can do this too. What can you predict about the welfare results of such competition?

3.7. Nash Equilibrium (easy)

Find the unique Nash equilibrium of the game in Table 9.

Table 9: A Meaningless Game

		Column		
		<i>Left</i>	<i>Middle</i>	<i>Right</i>
	<i>Up</i>	1,0	10, -1	0, 1
Row:	<i>Sideways</i>	-1,0	-2,-2	-12, 4
	<i>Down</i>	0, 2	823,-1	2, 0

Payoffs to: (Row, Column).

3.8. Triopoly (easy)

Three companies provide tires to the Australian market. The total cost curve for a firm making Q tires is $TC = 5 + 20Q$, and the demand equation is $P = 100 - N$, where N is the total number of tires on the market.

According to the Cournot model, in which the firms's simultaneously choose quantities, what will the total industry output be?

3.9. Cournot with Heterogeneous Costs (hard)

On a seminar visit, Professor Schaffer of Michigan told me that in a Cournot model with a linear demand curve $P = \alpha - \beta Q$ and constant marginal cost C_i for firm i , the equilibrium industry output Q depends on $\sum_i C_i$, but not on the individual levels of C_i . I may have misremembered. Prove or disprove this assertion. Would your conclusion be altered if we made some other assumption on demand? Discuss.

3.10. Alba and Rome: Asymmetric Information and Mixed Strategies (medium)

A Roman, Horatius, unwounded, is fighting the three Curiatius brothers from Alba, each of whom is wounded. If Horatius continues fighting, he wins with probability 0.1, and the payoffs are (10,-10) for (Horatius, Curiatii) if he wins, and (-10,10) if he loses. With probability $\alpha = 0.5$, Horatius is panic-stricken and runs away. If he runs and the Curiatii do not chase him, the payoffs are (-20, 10). If he runs and the Curiatius brothers chase and kill him, the payoffs are (-21, 20). If, however, he is not panic-stricken, but he runs anyway and the Curiatii give chase, he is able to kill the fastest brother first and then dispose of the other two, for payoffs of (10,-10). Horatius is, in fact, not panic-stricken.

- (a) With what probability θ would the Curiatii give chase if Horatius were to run?
- (b) With what probability γ does Horatius run?
- (c) How would θ and γ be affected if the Curiatii falsely believed that the probability of Horatius being panic-stricken was 1? What if they believed it was 0.9?

3.11. Finding Nash Equilibria (easy)

Find all of the Nash equilibria for the game of Table 10.

Table 10: A Takeover Game

		Target		
		<i>Hard</i>	<i>Medium</i>	<i>Soft</i>
Raider:	<i>Hard</i>	-3, -3	-1, 0	4, 0
	<i>Medium</i>	0, 0	2, 2	3, 1
	<i>Soft</i>	0, 0	2, 4	3, 3

Payoffs to: (Raider, Target).

3.12. Risky Skating (hard)

Elena and Mary are the two leading figure skaters in the world. Each must choose during her training what her routine is going to look like. They cannot change their minds later and try to alter any details of their routines. Elena goes first in the Olympics, and Mary goes next. Each has five minutes for her performance. The judges will rate the routines on three dimensions, beauty, how high they jump, and whether they stumble after they jump. A skater who stumbles is sure to lose, and if both Elena and Mary stumble, one of the ten lesser skaters will win, though those ten skaters have no chance otherwise.

Elena and Mary are exactly equal in the beauty of their routines, and both of them know this, but they are not equal in their jumping ability. Whoever jumps higher without stumbling will definitely win. Elena’s probability of stumbling is $P(h)$, where h is the height of the jump, and P is increasing smoothly and continuously in h . (In calculus terms, let P' and P'' both exist, and P' be positive) Mary’s probability is $P(h) - 0.1$ — that is, it is 10 percent less for equal heights.

Let us define as $h=0$ the maximum height that the lesser skaters can achieve, and assume that $P(0) = 0$.

- (a) Show that it cannot be an equilibrium for both Mary and Elena to choose the same value for h (Call them M and E).
- (b) Show for any pair of values (M, E) that it cannot be an equilibrium for Mary and Elena to choose those values.
- (c) Describe the optimal strategies to the best of your ability. (Do not get hung up on trying to answer this question; my expectations are not high here.)
- (d) What is a business analogy? Find some situation in business or economics that could use this same model.

3.13. The Kosovo War (easy)

Senator Robert Smith of New Hampshire said of the US policy in Serbia of bombing but promising not to use ground forces, “It’s like saying we’ll pass on you but we won’t run the football.” (*Human Events*, p. 1, April 16, 1999.) Explain what he meant, and why this is a strong criticism

of U.S. policy, using the concept of a mixed strategy equilibrium. (Foreign students: in American football, a team can choose to throw the football (to pass it) or to hold it and run with it to move towards the goal.) Construct a numerical example to compare the U.S. expected payoff in (a) a mixed strategy equilibrium in which it ends up not using ground forces, and (b) a pure strategy equilibrium in which the U.S. has committed not to use ground forces.

3.14. IMF Aid (easy)

Consider the game of Table 11.

Table 11: IMF Aid

		Debtor	
		Reform	Waste
IMF	Aid	3,2	-1,3
	No Aid	-1,1	0,0

Payoffs to: (IMF, Debtor).

- (a) What is the exact form of every Nash equilibrium?
- (b) For what story would this matrix be a good model?

3.15. Coupon Competition (hard)

Two marketing executives are arguing. Smith says that reducing our use of coupons will make us a less aggressive competitor, and that will hurt our sales. Jones says that reducing our use of coupons will make us a less aggressive competitor, but that will end up helping our sales.

Discuss, using the effect of reduced coupon use on your firm's reaction curve, under what circumstance each executive could be correct.

The War of Attrition: A Classroom Game for Chapter 3

Each firm consists of 3 students. Each year a firm must decide whether to stay in the industry or to exit. If it stays in, it incurs a fixed cost of 300 and a marginal cost of 2, and it chooses an integer price at which to sell. The firms can lose unlimited amounts of money; they are backed by large corporations who will keep supplying them with capital indefinitely.

Demand is inelastic at 60 up to a threshold price of \$10/unit, above which the quantity demanded falls to zero.

Each firm writes down its price (or the word "EXIT") on a piece of paper and gives it to the instructor. The instructor then writes the strategies of each firm on the blackboard (EXIT or price). The firm charging the lowest price sells to all 60 consumers. If there is a tie for the lowest price, the firms charging that price split the consumers evenly.

The game then starts with a new year, but any firm that has exited is out permanently and cannot re-enter. The game continues until only one firm is active, in which case it is awarded a prize of \$2,000, the capitalized value of being a monopolist. This means the game can continue forever, in theory. The instructor may wish to cut it off at some point, however.

The game can then be restarted and continued for as long as class time permits.