AN INTRODUCTION TO GAME THEORY

talio, and Beil (1990). The game in Exercise 34.1 is taken from Moulin (1986b, 72). The game in Exercise 34.2 was first studied by Palfrey and Rosenthal (1983). Exercise 34.3 is based on Braess (1968); see also Murchland (1970). The game in Exercise 38.2 is taken from Brams, Kilgour, and Davis (1993).

3 Nash Equilibrium: Illustrations

IN THIS CHAPTER I discuss in detail a few key models that use the notion of Nash equilibrium to study economic, political, and biological phenomena. The discussion shows how the notion of Nash equilibrium improves our understanding of a wide variety of phenomena. It also illustrates some of the many forms strategic games and their Nash equilibria can take. The models in Sections 3.1 and 3.2 are related to each other, whereas those in each of the other sections are independent of each other.

3.1 Cournot's model of oligopoly

3.1.1 Introduction

How does the outcome of competition among the firms in an industry depend on the characteristics of the demand for the firms' output, the nature of the firms' cost functions, and the number of firms? Will the benefits of technological improvements be passed on to consumers? Will a reduction in the number of firms generate a less desirable outcome? To answer these questions we need a model of the interaction between firms competing for the business of consumers. In this section and the next I analyze two such models. Economists refer to them as models of "oligopoly" (competition between a small number of sellers), though they involve no restriction on the number of firms; the label reflects the strategic interaction they capture. Both models were studied first in the 19th century, before the notion of Nash equilibrium was formalized for a general strategic game. The first is due to the economist Cournot (1838).

3.1.2 General model

A single good is produced by n firms. The cost to firm i of producing q_i units of the good is $C_i(q_i)$, where C_i is an increasing function (more output is more costly to produce). All the output is sold at a single price, determined by the demand for the good and the firms' total output. Specifically, if the firms' total output is Q, then the market price is $P(Q)$; P is called the "inverse demand function". Assume that P is a decreasing function when it is positive: if the firms' total output increases, then the price decreases (unless it is already zero). If the output of each firm i is q_i, then the price is $P(q_1 + \cdots + q_n)$, so that firm i's revenue is $q_i P(q_1 + \cdots + q_n)$. Thus firm i's profit, equal to its revenue minus its cost, is

$$\pi_i(q_1, \ldots, q_n) = q_i P(q_1 + \cdots + q_n) - C_i(q_i). \tag{56.1}$$

Cournot suggested that the industry be modeled as the following strategic game, which I refer to as **Cournot's oligopoly game**.

Players The firms.

Actions Each firm's set of actions is the set of its possible outputs (nonnegative numbers).

Preferences Each firm's preferences are represented by its profit, given in (56.1).

3.1.3 *Example: duopoly with constant unit cost and linear inverse demand function*

For specific forms of the functions C_i and P we can compute a Nash equilibrium of Cournot's game. Suppose there are two firms (the industry is a "duopoly"), each firm's cost function is the same, given by $C_i(q_i) = cq_i$ for all q_i ("unit cost" is constant, equal to c), and the inverse demand function is linear where it is positive, given by

$$P(Q) = \begin{cases} \alpha - Q & \text{if } Q \leq \alpha \\ 0 & \text{if } Q > \alpha, \end{cases} \tag{56.2}$$

where $\alpha > 0$ and $c > 0$ are constants. This inverse demand function is shown in Figure 57.1. (Note that the price $P(Q)$ cannot be equal to $\alpha - Q$ for *all* values of Q, for then it would be negative for $Q > \alpha$.) Assume that $c < \alpha$, so that there is some value of total output Q for which the market price $P(Q)$ is greater than the firms' common unit cost c. (If c were to exceed α, there would be no output for the firms at which they could make any profit, because the market price never exceeds α.)

To find the Nash equilibria in this example, we can use the procedure based on the firms' best response functions (Section 2.8.3). First we need to find the firms' payoffs (profits). If the firms' outputs are q_1 and q_2, then the market price $P(q_1 + q_2)$ is $\alpha - q_1 - q_2$ if $q_1 + q_2 \leq \alpha$ and zero if $q_1 + q_2 > \alpha$. Thus firm 1's profit is

$$\begin{aligned} \pi_1(q_1, q_2) &= q_1(P(q_1 + q_2) - c) \\ &= \begin{cases} q_1(\alpha - c - q_1 - q_2) & \text{if } q_1 + q_2 \leq \alpha \\ -cq_1 & \text{if } q_1 + q_2 > \alpha. \end{cases} \end{aligned}$$

To find firm 1's best response to any given output q_2 of firm 2, we need to study firm 1's profit as a function of its output q_1 for given values of q_2. If $q_2 = 0$, then

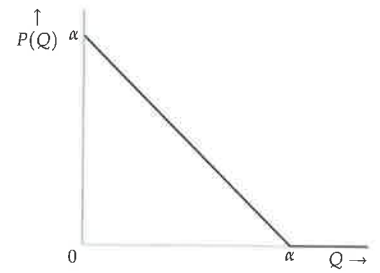

Figure 57.1 The inverse demand function in the example of Cournot's game studied in Section 3.1.3.

firm 1's profit is $\pi_1(q_1, 0) = q_1(\alpha - c - q_1)$ for $q_1 \leq \alpha$, a quadratic function that is zero when $q_1 = 0$ and when $q_1 = \alpha - c$. This function is the black curve in Figure 58.1. Given the symmetry of quadratic functions (Section 17.3), the output q_1 of firm 1 that maximizes its profit is $q_1 = \frac{1}{2}(\alpha - c)$. (If you know calculus, you can reach the same conclusion by setting the derivative of firm 1's profit with respect to q_1 equal to zero and solving for q_1.) Thus firm 1's best response to an output of zero for firm 2 is $b_1(0) = \frac{1}{2}(\alpha - c)$.

As the output q_2 of firm 2 increases, the profit firm 1 can obtain at any given output decreases, because more output of firm 2 means a lower price. The gray curve in Figure 58.1 is an example of $\pi_1(q_1, q_2)$ for $q_2 > 0$ and $q_2 < \alpha - c$. Again this function is a quadratic up to the output $q_1 = \alpha - q_2$ that leads to a price of zero. Specifically, the quadratic is $\pi_1(q_1, q_2) = q_1(\alpha - c - q_2 - q_1)$, which is zero when $q_1 = 0$ and when $q_1 = \alpha - c - q_2$. From the symmetry of quadratic functions (or some calculus), we conclude that the output that maximizes $\pi_1(q_1, q_2)$ is $q_1 = \frac{1}{2}(\alpha - c - q_2)$. (When $q_2 = 0$, this is equal to $\frac{1}{2}(\alpha - c)$, the best response to an output of zero that we found in the previous paragraph.)

When $q_2 > \alpha - c$, the value of $\alpha - c - q_2$ is negative. Thus for such a value of q_2, we have $q_1(\alpha - c - q_2 - q_1) < 0$ for all positive values of q_1: firm 1's profit is negative for any positive output, so that its best response is to produce no output.

We conclude that the best response of firm 1 to the output q_2 of firm 2 depends on the value of q_2: if $q_2 \leq \alpha - c$, then firm 1's best response is $\frac{1}{2}(\alpha - c - q_2)$, whereas if $q_2 > \alpha - c$, then firm 1's best response is 0. Or, more compactly,

$$b_1(q_2) = \begin{cases} \frac{1}{2}(\alpha - c - q_2) & \text{if } q_2 \leq \alpha - c \\ 0 & \text{if } q_2 > \alpha - c. \end{cases}$$

Because firm 2's cost function is the same as firm 1's, its best response function b_2 is also the same: for any number q, we have $b_2(q) = b_1(q)$. Of course, firm 2's best response function associates a value of firm 2's output with every output of firm 1, whereas firm 1's best response function associates a value of firm 1's output with every output of firm 2, so we plot them relative to different axes. They

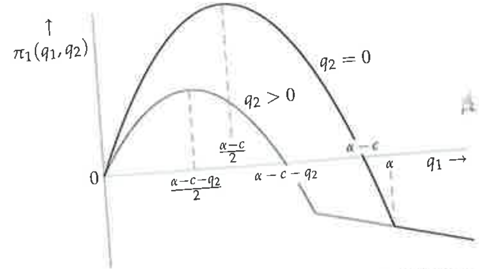

Figure 58.1 Firm 1's profit as a function of its output, given firm 2's output. The black curve shows the case $q_2 = 0$, whereas the gray curve shows a case in which $q_2 > 0$.

are shown in Figure 59.1 (b_1 is black; b_2 is gray). As for a general game (see Section 2.8.3), b_1 associates each point on the vertical axis with a point on the horizontal axis, and b_2 associates each point on the horizontal axis with a point on the vertical axis.

A Nash equilibrium is a pair (q_1^*, q_2^*) of outputs for which q_1^* is a best response to q_2^*, and q_2^* is a best response to q_1^*:

$$q_1^* = b_1(q_2^*) \quad \text{and} \quad q_2^* = b_2(q_1^*)$$

(see (36.3)). The set of such pairs is the set of points at which the best response functions in Figure 59.1 intersect. From the figure we see that there is exactly one such point, which is given by the solution of the two equations

$$\begin{aligned} q_1 &= \tfrac{1}{2}(\alpha - c - q_2) \\ q_2 &= \tfrac{1}{2}(\alpha - c - q_1). \end{aligned}$$

Solving these two equations (by substituting the second into the first and then isolating q_1, for example) we find that $q_1^* = q_2^* = \frac{1}{3}(\alpha - c)$.

In summary, when there are two firms, the inverse demand function is given by $P(Q) = \alpha - Q$ for $Q \leq \alpha$, and the cost function of each firm is $C_i(q_i) = cq_i$, Cournot's oligopoly game has a unique Nash equilibrium $(q_1^*, q_2^*) = (\frac{1}{3}(\alpha - c), \frac{1}{3}(\alpha - c))$. The total output in this equilibrium is $\frac{2}{3}(\alpha - c)$, so that the price at which output is sold is $P(\frac{2}{3}(\alpha - c)) = \frac{1}{3}(\alpha + 2c)$. As α increases (meaning that consumers are willing to pay more for the good), the equilibrium price and the output of each firm increase. As c (the unit cost of production) increases, the output of each firm falls and the price rises; each unit increase in c leads to a two-thirds of a unit increase in the price.

EXERCISE 58.1 (Cournot's duopoly game with linear inverse demand and different unit costs) Find the Nash equilibrium of Cournot's game when there are two firms, the inverse demand function is given by (56.2), the cost function of each firm i is

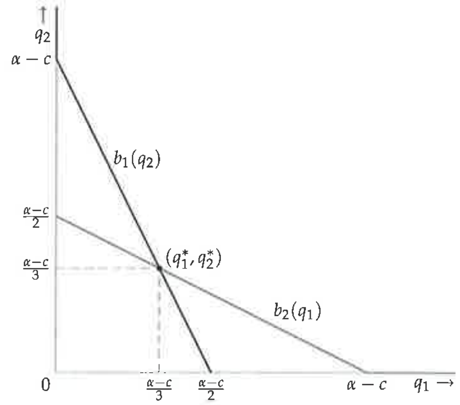

Figure 59.1 The best response functions in Cournot's duopoly game when the inverse demand function is given by (56.2) and the cost function of each firm is cq. The unique Nash equilibrium is $(q_1^*, q_2^*) = (\frac{1}{3}(\alpha - c), \frac{1}{3}(\alpha - c))$.

$C_i(q_i) = c_i q_i$, where $c_1 > c_2$, and $c_1 < \alpha$. (There are two cases, depending on the size of c_1 relative to c_2.) Which firm produces more output in an equilibrium? What is the effect of technical change that lowers firm 2's unit cost c_2 (while not affecting firm 1's unit cost c_1) on the firms' equilibrium outputs, the total output, and the price?

EXERCISE 59.1 (Cournot's duopoly game with linear inverse demand and a quadratic cost function) Find the Nash equilibrium of Cournot's game when there are two firms, the inverse demand function is given by (56.2), and the cost function of each firm i is $C_i(q_i) = q_i^2$.

In the next exercise each firm's cost function has a component that is independent of output. You will find in this case that Cournot's game may have more than one Nash equilibrium.

EXERCISE 59.2 (Cournot's duopoly game with linear inverse demand and a fixed cost) Find the Nash equilibria of Cournot's game when there are two firms, the inverse demand function is given by (56.2), and the cost function of each firm i is given by

$$C_i(q_i) = \begin{cases} 0 & \text{if } q_i = 0 \\ f + cq_i & \text{if } q_i > 0, \end{cases}$$

where $c > 0$, $f > 0$, and $c < \alpha$. (Note that the fixed cost f affects only the firm's decision of whether to operate; it does not affect the output a firm wishes to produce *if it wishes to operate*.)

So far we have assumed that each firm's objective is to maximize its profit. The next exercise asks you to consider a case in which one firm's objective is to maximize its market share.

EXERCISE 60.1 (Variant of Cournot's duopoly game with market-share-maximizing firms) Find the Nash equilibrium (equilibria?) of a variant of the example of Cournot's duopoly game that differs from the one in this section (linear inverse demand, constant unit cost) only in that one of the two firms chooses its output to maximize its market share subject to not making a loss, rather than to maximize its profit. What happens if *each* firm maximizes its market share?

3.1.4 Properties of Nash equilibrium

Two economically interesting properties of a Nash equilibrium of Cournot's game concern the relation between the firms' equilibrium profits and the profits they could obtain if they acted collusively, and the character of an equilibrium when the number of firms is large.

Comparison of Nash equilibrium with collusive outcomes In Cournot's game with two firms, is there any pair of outputs at which *both* firms' profits exceed their levels in a Nash equilibrium? The next exercise asks you to show that the answer is "yes" in the example considered in the previous section (3.1.3). Specifically, both firms can increase their profits relative to their equilibrium levels by reducing their outputs.

EXERCISE 60.2 (Nash equilibrium of Cournot's duopoly game and collusive outcomes) Find the total output (call it Q^*) that maximizes the firms' *total* profit in Cournot's game when there are two firms and the inverse demand function and cost functions take the forms assumed Section 3.1.3. Compare $\frac{1}{2}Q^*$ with each firm's output in the Nash equilibrium, and show that each firm's equilibrium profit is less than its profit in the "collusive" outcome in which each firm produces $\frac{1}{2}Q^*$. Why is this collusive outcome not a Nash equilibrium?

The same is true more generally. For nonlinear inverse demand functions and cost functions, the shapes of the firms' best response functions differ, in general, from those in the example studied in the previous section. But for many inverse demand functions and cost functions, the game has a Nash equilibrium and, for any equilibrium, there are pairs of outputs in which each firm's output is less than its equilibrium level and each firm's profit exceeds its equilibrium level.

To see why, suppose that (q_1^*, q_2^*) is a Nash equilibrium and consider the set of pairs (q_1, q_2) of outputs at which firm 1's profit is at least its equilibrium profit. The assumption that P is decreasing (higher total output leads to a lower price) implies that if (q_1, q_2) is in this set and $q_2' < q_2$, then (q_1, q_2') is also in the set. (We have $q_1 + q_2' < q_1 + q_2$, and hence $P(q_1 + q_2') > P(q_1 + q_2)$, so that firm 1's profit at (q_1, q_2') exceeds its profit at (q_1, q_2).) Thus in Figure 61.1 the set of pairs of outputs at which firm 1's profit is at least its equilibrium profit lies on or below the line $q_2 = q_2^*$; an example of such a set is shaded light gray. Similarly, the set of pairs of

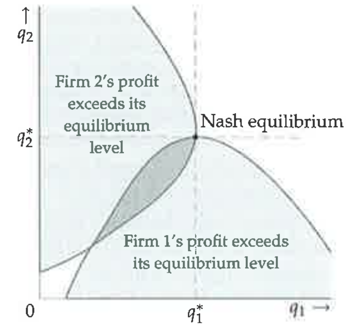

Figure 61.1 The pair (q_1^*, q_2^*) is a Nash equilibrium; along each gray curve one of the firm's profits is constant, equal to its profit at the equilibrium. The area shaded dark gray is the set of pairs of outputs at which both firms' profits exceed their equilibrium levels.

outputs at which firm 2's profit is at least its equilibrium profit lies on or to the left of the line $q_1 = q_1^*$, and an example is shaded light gray.

We see that if the parts of the boundaries of these sets indicated by the gray lines in Figure 61.1 are smooth, then the two sets must intersect; in the figure the intersection is shaded dark gray. At every pair of outputs in this area, each firm's output is less than its equilibrium level ($q_i < q_i^*$ for $i = 1, 2$) and each firm's profit is higher than its equilibrium profit. That is, *both* firms are better off by restricting their outputs.

Dependence of Nash equilibrium on number of firms How does the equilibrium outcome in Cournot's game depend on the number of firms? If each firm's cost function has the same constant unit cost c, the best outcome for consumers compatible with no firm's making a loss has a price of c and a total output of $\alpha - c$. The next exercise asks you to show that if, for this cost function, the inverse demand function is linear (as in Section 3.1.3), then the price in the Nash equilibrium of Cournot's game decreases as the number of firms increases, approaching c. That is, from the viewpoint of consumers, the outcome is better the larger the number of firms, and when the number of firms is very large, the outcome is close to the best one compatible with nonnegative profits for the firms.

EXERCISE 61.1 (Cournot's game with many firms) Consider Cournot's game in the case of an arbitrary number n of firms; retain the assumptions that the inverse demand function takes the form (56.2) and the cost function of each firm i is $C_i(q_i) = cq_i$ for all q_i, with $c < \alpha$. Find the best response function of each firm and set up the conditions for $(q_1^*, \ldots, q_n^*)$ to be a Nash equilibrium (see (36.3)), assuming that there is a Nash equilibrium in which all firms' outputs are positive. Solve these equations to find the Nash equilibrium. (For $n = 2$ your answer should be $(\frac{1}{3}(\alpha - c), \frac{1}{3}(\alpha - c))$, the equilibrium found in Section 3.1.3. First show that in an equilibrium all firms produce the same output, then solve for that output. If you cannot show that all firms produce the same output, simply assume that they do.) Find the price at which output is sold in a Nash equilibrium and show that this price decreases as n increases, approaching c as the number of firms increases without bound.

The main idea behind this result does not depend on the assumptions about the inverse demand function and the firms' cost functions. Suppose, more generally, that the inverse demand function is any decreasing function, that each firm's cost function is the same, denoted by C, and that there is a single output, say $\underline{q}$, at which the average cost of production $C(q)/q$ is minimal. In this case, any given total output is produced most efficiently by each firm's producing $\underline{q}$, and the lowest price compatible with the firms' not making losses is the minimal value of the average cost. The next exercise asks you to show that in a Nash equilibrium of Cournot's game in which the firms' total output is large relative to $\underline{q}$, this is the price at which the output is sold.

② EXERCISE 62.1 (Nash equilibrium of Cournot's game with small firms) Suppose that there are infinitely many firms, all of which have the same cost function C. Assume that $C(0) = 0$, and for $q > 0$ the function $C(q)/q$ has a unique minimizer $\underline{q}$; denote the minimum of $C(q)/q$ by $\underline{p}$. Assume that the inverse demand function P is decreasing. Show that in any Nash equilibrium the firms' total output Q^* satisfies

$$P(Q^* + \underline{q}) \leq \underline{p} \leq P(Q^*).$$

(That is, the price is at least the minimal value $\underline{p}$ of the average cost, but is close enough to this minimum that increasing the total output of the firms by $\underline{q}$ would reduce the price to at most $\underline{p}$.) To establish these inequalities, show that if $P(Q^*) < \underline{p}$ or $P(Q^* + \underline{q}) > \underline{p}$, then Q^* is not the total output of the firms in a Nash equilibrium, because in each case at least one firm can deviate and increase its profit.

3.1.5 *A generalization of Cournot's game: using common property*

In Cournot's game, the payoff function of each firm i is $q_i P(q_1 + \cdots + q_n) - C_i(q_i)$. In particular, each firm's payoff depends only on its output and the sum of all the firm's outputs, not on the distribution of the total output among the firms, and decreases when this sum increases (given that P is decreasing). That is, the payoff of each firm i may be written as $f_i(q_i, q_1 + \cdots + q_n)$, where the function f_i is decreasing in its second argument (given the value of its first argument, q_i).

This general payoff function captures many situations in which players compete in using a piece of common property whose value to any one player diminishes as total use increases. The property might, for example, be a village green that is less valuable to any given farmer the higher the total number of sheep that graze on it.

The first property of a Nash equilibrium in Cournot's model discussed in Section 3.1.3 applies to this general model: common property is "overused" in a Nash equilibrium in the sense that every player's payoff increases when every player reduces her use of the property from its equilibrium level. For example, all farmers' payoffs increase if each farmer reduces her use of the village green from its equilibrium level: in an equilibrium the green is "overgrazed". The argument is the same as the one illustrated in Figure 61.1 in the case of two players, because this argument depends only on the fact that each player's payoff function is smooth and is decreasing in the other player's action. (In Cournot's model, the "common property" that is overused is the demand for the good.)

⑦ EXERCISE 63.1 (Interaction among resource users) A group of n firms uses a common resource (a river or a forest, for example) to produce output. As the total amount of the resource used by all firms increases, any given firm can produce less output. Denote by x_i the amount of the resource used by firm i ($= 1, \ldots, n$). Assume specifically that firm i's output is $x_i(1 - (x_1 + \cdots + x_n))$ if $x_1 + \cdots + x_n \leq 1$, and zero otherwise. Each firm i chooses x_i to maximize its output. Formulate this situation as a strategic game. Find the Nash equilibria of the game. Find an action profile $(x_1, \ldots, x_n)$ at which each firm's output is higher than it is at any Nash equilibrium.

3.2 Bertrand's model of oligopoly

3.2.1 *General model*

In Cournot's game, each firm chooses an output; the price is determined by the demand for the good in relation to the total output produced. In an alternative model of oligopoly, associated with a review of Cournot's book by Bertrand (1883), each firm chooses a price, and produces enough output to meet the demand it faces, given the prices chosen by all the firms. The model is designed to shed light on the same questions that Cournot's game addresses; as we shall see, some of the answers it gives are different.

The economic setting for the model is similar to that for Cournot's game. A single good is produced by n firms; each firm can produce q_i units of the good at a cost of $C_i(q_i)$. It is convenient to specify demand by giving a "demand function" D, rather than an inverse demand function as we did for Cournot's game. The interpretation of D is that if the good is available at the price p, then the total amount demanded is $D(p)$.

Assume that if the firms set different prices, then all consumers purchase the good from the firm with the lowest price, which produces enough output to meet this demand. If more than one firm sets the lowest price, all the firms doing so share the demand at that price equally. A firm whose price is not the lowest price receives no demand and produces no output. (Note that a firm does not choose its output strategically; it simply produces enough to satisfy all the demand it faces, given the prices, even if its price is below its unit cost, in which case it makes a loss. This assumption can be modified at the price of complicating the model.)

In summary, **Bertrand's oligopoly game** is the following strategic game.

Players The firms.

Actions Each firm's set of actions is the set of possible prices (nonnegative numbers).

Preferences Firm i's preferences are represented by its profit, which is equal to $p_iD(p_i)/m - C_i(D(p_i)/m)$ if firm i is one of m firms setting the lowest price ($m = 1$ if firm i's price p_i is lower than every other price), and equal to zero if some firm's price is lower than p_i.

3.2.2 *Example: duopoly with constant unit cost and linear demand function*

Suppose, as in Section 3.1.3, that there are two firms, each of whose cost function has constant unit cost c (that is, $C_i(q_i) = cq_i$ for $i = 1, 2$). Assume that the demand function is $D(p) = \alpha - p$ for $p \leq \alpha$ and $D(p) = 0$ for $p > \alpha$, and that $c < \alpha$.

Because the cost of producing each unit is the same, equal to c, firm i makes the profit of $p_i - c$ on every unit it sells. Thus its profit is

$$\pi_i(p_1, p_2) = \begin{cases} (p_i - c)(\alpha - p_i) & \text{if } p_i < p_j \\ \frac{1}{2}(p_i - c)(\alpha - p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j, \end{cases}$$

where j is the other firm ($j = 2$ if $i = 1$, and $j = 1$ if $i = 2$).

As before, we can find the Nash equilibria of the game by finding the firms' best response functions. If firm j charges p_j, what is the best price for firm i to charge? We can reason informally as follows. If firm i charges p_j, it shares the market with firm j; if it charges slightly less, it sells to the entire market. Thus if p_j exceeds c, so that firm i makes a positive profit selling the good at a price slightly below p_j, firm i is definitely better off serving all the market at such a price than serving half of the market at the price p_j. If p_j is very high, however, firm i may be able to do even better: by reducing its price significantly below p_j it may increase its profit, because the extra demand engendered by the lower price may more than compensate for the lower revenue per unit sold. Finally, if p_j is less than c, then firm i's profit is negative if it charges a price less than or equal to p_j, whereas this profit is zero if it charges a higher price. Thus in this case firm i would like to charge any price greater than p_j, to make sure that it gets no customers. (Remember that if customers arrive at its door it is obliged to serve them, regardless of whether it makes a profit by so doing.)

We can make these arguments precise by studying firm i's payoff as a function of its price p_i for various values of the price p_j of firm j. Denote by p^m the value of p (price) that maximizes $(p - c)(\alpha - p)$. This price would be charged by a firm with a monopoly of the market (because $(p - c)(\alpha - p)$ is the profit of such a firm).

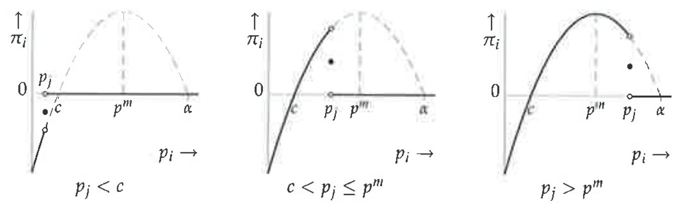

Figure 65.1 Three cross sections (in black) of firm i's payoff function in Bertrand's duopoly game. Where the payoff function jumps, its value is given by the small disk; the small circles indicate points that are excluded as values of the functions.

Three cross sections of firm i's payoff function, for different values of p_j, are shown in black in Figure 65.1. (The gray dashed line is the function $(p_i - c)(\alpha - p_i)$.)

- If $p_j < c$ (firm j's price is below the unit cost), then firm i's profit is negative if $p_i \leq p_j$ and zero if $p_i > p_j$ (see the left panel of Figure 65.1). Thus *any* price greater than p_j is a best response to p_j. That is, the set of firm i's best responses is $B_i(p_j) = \{p_i: p_i > p_j\}$.

- If $p_j = c$, then the analysis is similar to that of the previous case except that p_j, as well as any price greater than p_j, yields a profit of zero, and hence is a best response to p_j: $B_i(p_j) = \{p_i: p_i \geq p_j\}$.

- If $c < p_j \leq p^m$, then firm i's profit increases as p_i increases to p_j, then drops abruptly at p_j (see the middle panel of Figure 65.1). Thus there is no best response: firm i wants to choose a price less than p_j, but is better off the closer that price is to p_j. For any price less than p_j there is a higher price that is also less than p_j, so there is no best price. (I have assumed that a firm can choose *any* number as its price; in particular, it is not restricted to charge an integral number of cents.) Thus $B_i(p_j)$ is empty (has no members).

- If $p_j > p^m$, then p^m is the unique best response of firm i (see the right panel of Figure 65.1): $B_i(p_j) = \{p^m\}$.

In summary, firm i's best response function is given by

$$B_i(p_j) = \begin{cases} \{p_i: p_i > p_j\} & \text{if } p_j < c \\ \{p_i: p_i \geq p_j\} & \text{if } p_j = c \\ \varnothing & \text{if } c < p_j \leq p^m \\ \{p^m\} & \text{if } p^m < p_j, \end{cases}$$

where $\varnothing$ denotes the set with no members (the "empty set"). Note the respects in which this best response function differs qualitatively from a firm's best response function in Cournot's game: for some actions of its opponent, a firm has no best response, and for some actions it has multiple best responses.

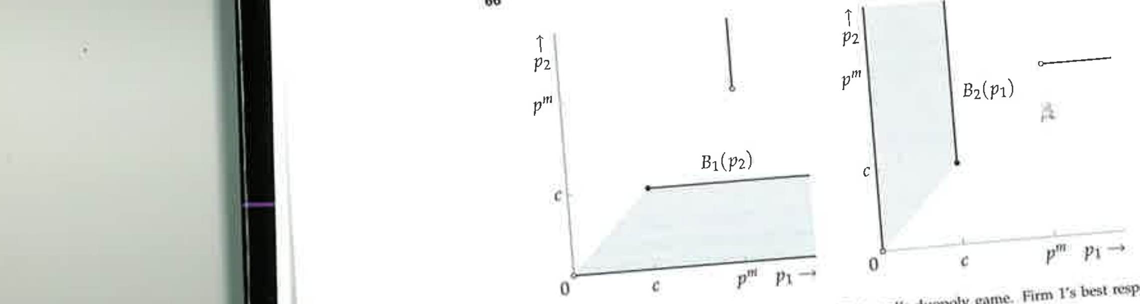

Figure 66.1 The firms' best response functions in Bertrand's duopoly game. Firm 1's best response function is in the left panel; firm 2's is in the right panel.

The fact that firm i has *no* best response when $c < p_j < p^m$ is an artifact of modeling price as a continuous variable (a firm can choose its price to be any nonnegative number). If instead we assume that each firm's price must be a multiple of some indivisible unit ϵ (e.g. price must be an integral number of cents), then firm i's optimal response to a price p_j with $c < p_j < p^m$ is $p_j - \epsilon$. I model price as a continuous variable because doing so simplifies some of the analysis; in Exercise 67.2 you are asked to study the case of discrete prices.

When $p_j < c$, firm i's set of best responses is the set of all prices greater than p_j. In particular, prices between p_j and c are best responses. You may object that setting a price less than c is not very sensible. Such a price exposes firm i to the risk of posting a loss (if firm j chooses a higher price) and has no advantage over the price of c, regardless of firm j's price. That is, such a price is *weakly dominated* (Definition 46.1) by the price c. Nevertheless, such a price *is* a best response! That is, it is optimal for firm i to choose such a price, *given* firm j's price: there is no price that yields firm i a higher profit, *given* firm j's price. The point is that when asking if a player's action is a best response to her opponent's action, we do not consider the "risk" that the opponent will take some other action.

Figure 66.1 shows the firms' best response functions (firm 1's on the left, firm 2's on the right). The shaded gray area in the left panel indicates that for a price p_2 less than c, *any* price greater than p_2 is a best response for firm 1. The absence of a black line along the sloping left boundary of this area indicates that only prices p_1 *greater than* (not equal to) p_2 are included. The black line along the top of the area indicates that for $p_2 = c$ any price greater than *or equal to* c is a best response. As before, the dot indicates a point that is included, whereas the small circle indicates a point that is excluded. Firm 2's best response function has a similar interpretation.

A Nash equilibrium is a pair (p_1^*, p_2^*) of prices such that p_1^* is a best response to p_2^*, and p_2^* is a best response to p_1^*—that is, p_1^* is in $B_1(p_2^*)$ and p_2^* is in $B_2(p_1^*)$ (see (36.2)). If we superimpose the two best response functions, any such pair is in the intersection of their graphs. If you do so, you will see that the graphs have a single

point of intersection, namely $(p_1^*, p_2^*) = (c, c)$. That is, the game has a single Nash equilibrium, in which each firm charges the price c.

The method of finding the Nash equilibria of a game by constructing the players' best response functions is systematic. As long as these functions may be computed, the method straightforwardly leads to the set of Nash equilibria. However, in some games we can make a direct argument that avoids the need to construct the entire best response functions. Using a combination of intuition and trial and error, we find the action profiles that seem to be equilibria; then we show precisely that any such profile is an equilibrium and every other profile is not an equilibrium. To show that a pair of actions is not a Nash equilibrium we need only find a *better* response for one of the players—not necessarily the *best* response.

In Bertrand's game we can argue as follows. (*i*) First we show that $(p_1, p_2) = (c, c)$ is a Nash equilibrium. If one firm charges the price c, then the other firm can do no better than charge the price c also, because if it raises its price it sells no output, and if it lowers its price it posts a loss. (*ii*) Next we show that no other pair (p_1, p_2) is a Nash equilibrium, as follows.

- If $p_i < c$ for either $i = 1$ or $i = 2$, then the profit of the firm whose price is lowest (or the profit of both firms, if the prices are the same) is negative, and this firm can increase its profit (to zero) by raising its price to c.
- If $p_i = c$ and $p_j > c$, then firm i is better off increasing its price slightly, making its profit positive rather than zero.
- If $p_i > c$ and $p_j > c$, suppose that $p_i \geq p_j$. Then firm i can increase its profit by lowering p_i to slightly below p_j if $D(p_j) > 0$ (i.e. if $p_j < \alpha$) and to p^m if $D(p_j) = 0$ (i.e. if $p_j \geq \alpha$).

In conclusion, both arguments show that when the unit cost of production is a constant c, the same for both firms, and demand is linear, Bertrand's game has a unique Nash equilibrium, in which each firm's price is equal to c.

⑦ EXERCISE 67.1 (Bertrand's duopoly game with constant unit cost) Consider the extent to which the analysis depends upon the demand function D taking the specific form $D(p) = \alpha - p$. Suppose that D is any function for which $D(p) \geq 0$ for all p and there exists $\overline{p} > c$ such that $D(p) > 0$ for all $p \leq \overline{p}$. Is (c, c) still a Nash equilibrium? Is it still the only Nash equilibrium?

? EXERCISE 67.2 (Bertrand's duopoly game with discrete prices) Consider the variant of the example of Bertrand's duopoly game in this section in which each firm is restricted to choose a price that is an integral number of cents. Take the monetary unit to be a cent, and assume that c is an integer and $\alpha > c + 1$. Is (c, c) a Nash equilibrium of this game? Is there any other Nash equilibrium?

3.2.3 *Discussion*

For a duopoly in which both firms have the same constant unit cost and the demand function is linear, the Nash equilibria of Cournot's and Bertrand's games generate different economic outcomes. The equilibrium price in Bertrand's game is equal to the common unit cost c, whereas the price associated with the equilibrium of Cournot's game is $\frac{1}{3}(\alpha + 2c)$, which exceeds c because $c < \alpha$. In particular, the equilibrium price in Bertrand's game is the lowest price compatible with the firms' not posting losses, whereas the price at the equilibrium of Cournot's game is higher. In Cournot's game, the price decreases toward c as the number of firms increases (Exercise 61.1), whereas in Bertrand's game it is c even if there are only two firms. In the next exercise you are asked to show that as the number of firms increases in Bertrand's game, the price remains c.

? EXERCISE 68.1 (Bertrand's oligopoly game) Consider Bertrand's oligopoly game when the cost and demand functions satisfy the conditions in Section 3.2.2 and there are n firms, with $n \geq 3$. Show that the set of Nash equilibria is the set of profiles $(p_1, \ldots, p_n)$ of prices for which $p_i \geq c$ for all i and at least two prices are equal to c. (Show that any such profile is a Nash equilibrium, and that every other profile is not a Nash equilibrium.)

What accounts for the difference between the Nash equilibria of Cournot's and Bertrand's games? The key point is that different strategic variables (output in Cournot's game, price in Bertrand's game) imply different types of strategic reasoning by the firms. In Cournot's game a firm changes its behavior if it can increase its profit by changing its output, on the assumption that the other firms' outputs will remain the same and the price will adjust to clear the market. In Bertrand's game a firm changes its behavior if it can increase its profit by changing its price, on the assumption that the other firms' prices will remain the same and their outputs will adjust to clear the market. Which assumption makes more sense depends on the context. For example, the wholesale market for agricultural produce may fit Cournot's game better, whereas the retail market for food may fit Bertrand's game better.

Under some variants of the assumptions in the previous section (3.2.2), Bertrand's game has no Nash equilibrium. In one case the firms' cost functions have constant unit costs, and these costs are different; in another case the cost functions have a fixed component. In both these cases, as well as in some other cases, an equilibrium is restored if we modify the way in which consumers are divided between the firms when the prices are the same, as the following exercises show. (We can think of the division of consumers between firms charging the same price as being determined as part of the equilibrium. Note that we retain the assumption that if the firms charge different prices, then the one charging the lower price receives all the demand.)

? EXERCISE 68.2 (Bertrand's duopoly game with different unit costs) Consider Bertrand's duopoly game under a variant of the assumptions of Section 3.2.2 in which the firms' unit costs are different, equal to c_1 and c_2, where $c_1 < c_2$. Denote by p_1^m the price that maximizes $(p - c_1)(\alpha - p)$, and assume that $c_2 < p_1^m$ and that the function $(p - c_1)(\alpha - p)$ is increasing in p up to p_1^m.

a. Suppose that the rule for splitting up consumers when the prices are equal assigns all consumers to firm 1 when both firms charge the price c_2. Show that $(p_1, p_2) = (c_2, c_2)$ is a Nash equilibrium and that no other pair of prices is a Nash equilibrium.

b. Show that no Nash equilibrium exists if the rule for splitting up consumers when the prices are equal assigns some consumers to firm 2 when both firms charge c_2.

? EXERCISE 69.1 (Bertrand's duopoly game with fixed costs) Consider Bertrand's game under a variant of the assumptions of Section 3.2.2 in which the cost function of each firm i is given by $C_i(q_i) = f + cq_i$ for $q_i > 0$, and $C_i(0) = 0$, where f is positive and less than the maximum of $(p - c)(\alpha - p)$ with respect to p. Denote by $\overline{p}$ the price p that satisfies $(p - c)(\alpha - p) = f$ and is less than the maximizer of $(p - c)(\alpha - p)$ (see Figure 69.1). Show that if firm 1 gets all the demand when both firms charge the same price, then $(\overline{p}, \overline{p})$ is a Nash equilibrium. Show also that no other pair of prices is a Nash equilibrium. (First consider cases in which the firms charge the same price, then cases in which they charge different prices.)

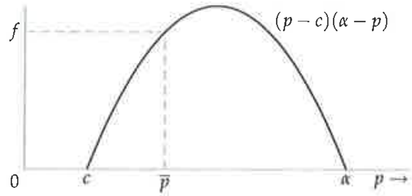

Figure 69.1 The determination of the price $\overline{p}$ in Exercise 69.1.

COURNOT, BERTRAND, AND NASH: SOME HISTORICAL NOTES

Associating the names of Cournot and Bertrand with the strategic games in Sections 3.1 and 3.2 invites two conclusions. First, that Cournot, writing in the first half of the 19th century, developed the concept of Nash equilibrium in the context of a model of oligopoly. Second, that Bertrand, dissatisfied with Cournot's game, proposed an alternative model in which price rather than output is the strategic variable. On both points the history is much less straightforward.

Cournot presented his "equilibrium" as the outcome of a dynamic adjustment process in which, in the case of two firms, the firms alternately choose best responses to each other's outputs. During such an adjustment process, each firm, when choosing an output, acts on the assumption that the other firm's output will remain the same, an assumption shown to be incorrect when the other firm subsequently adjusts its output. The fact that the adjustment process rests on the firms' acting on assumptions constantly shown to be false was the subject of criticism in a leading presentation of Cournot's model (Fellner 1949) available at the time Nash was developing his idea.

Certainly Nash did not literally generalize Cournot's idea: the evidence suggests that he was completely unaware of Cournot's work when developing the notion of Nash equilibrium (Leonard 1994, 502–503). In fact, only gradually, as Nash's work was absorbed into mainstream economic theory, was Cournot's solution interpreted as a Nash equilibrium (Leonard 1994, 507–509).

The association of the price-setting model with Bertrand (a mathematician) rests on a paragraph he wrote in 1883 in a review of Cournot's book. (The book, published in 1838, had previously been largely ignored.) The review is confused. Bertrand is under the impression that in Cournot's model the firms compete in prices, undercutting each other to attract more business! He argues that there is "no solution" because there is no limit to the fall in prices, a result he says that Cournot's formulation conceals (Bertrand 1883, 503). In brief, Bertrand's understanding of Cournot's work is flawed; he sees that price competition leads each firm to undercut the other, but his conclusion about the outcome is incorrect.

Through the lens of modern game theory we see that the models associated with Cournot and Bertrand are strategic games that differ only in the strategic variable, the solution in both cases being a Nash equilibrium. Until Nash's work, the picture was much murkier.

3.3 Electoral competition

What factors determine the number of political parties and the policies they propose? How is the outcome of an election affected by the electoral system and the voters' preferences among policies? A model that is the foundation for many theories of political phenomena addresses these questions. In the model, each of several candidates chooses a policy; each citizen has preferences over policies and votes for one of the candidates.

A simple version of this model is a strategic game in which the players are the candidates and a policy is a number, referred to as a "position". (The compression of all policy differences into one dimension is a major abstraction, though political positions are often categorized on a left–right axis.) After the candidates have chosen positions, each of a set of citizens votes (nonstrategically) for the candidate whose position she likes best. The candidate who obtains the most votes wins. Each candidate cares only about winning; no candidate has an ideological attachment to any position. Specifically, each candidate prefers to win than to tie for first place (in which case perhaps the winner is determined randomly), and prefers to tie for first place than to lose; if she ties for first place, she prefers to do so with as few other candidates as possible.

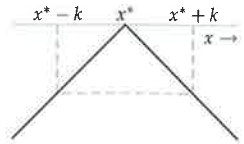

Figure 71.1 The payoff of a voter whose favorite position is x^*, as a function of the winning position, x.

There is a continuum of voters, each with a favorite position. The distribution of these favorite positions over the set of all possible positions is arbitrary. In particular, this distribution may not be uniform: a large fraction of the voters may have favorite positions close to one point, while few voters have favorite positions close to some other point. A position that turns out to have special significance is the *median* favorite position: the position m with the property that exactly half of the voters' favorite positions are at most m, and half of the voters' favorite positions are at least m. (I assume that the distribution of favorite positions is such that there is only one position with this property.)

Each voter's distaste for any position is given by the distance between that position and her favorite position. In particular, for any value of k, a voter whose favorite position is x^* is indifferent between the positions $x^* - k$ and $x^* + k$. (Refer to Figure 71.1.)

Under this assumption, each candidate attracts the votes of all citizens whose favorite positions are closer to her position than to the position of any other candidate. An example is shown in Figure 71.2. In this example there are three candidates, with positions x_1, x_2, and x_3. Candidate 1 attracts the votes of every citizen whose favorite position is in the interval, labeled "Votes for 1", up to the midpoint $\frac{1}{2}(x_1 + x_2)$ of the line segment from x_1 to x_2; candidate 2 attracts the votes of every citizen whose favorite position is in the interval from $\frac{1}{2}(x_1 + x_2)$ to $\frac{1}{2}(x_2 + x_3)$; and candidate 3 attracts the remaining votes. I assume that citizens whose favorite position is $\frac{1}{2}(x_1 + x_2)$ divide their votes equally between candidates 1 and 2, and those whose favorite position is $\frac{1}{2}(x_2 + x_3)$ divide their votes equally between candidates 2 and 3. If two or more candidates take the same position, then they share equally the votes that the position attracts.

In summary, I consider the following strategic game, which, in honor of its originator, I call **Hotelling's model of electoral competition**.

Players The candidates.

Actions Each candidate's set of actions is the set of positions (numbers).

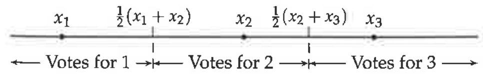

Figure 71.2 The allocation of votes between three candidates, with positions x_1, x_2, and x_3.

Preferences Each candidate's preferences are represented by a payoff function that assigns n to every action profile in which she wins outright, k to every action profile in which she ties for first place with $n-k$ other candidates (for $1 \le k \le n-1$), and 0 to every action profile in which she loses, where positions attract votes in the way described in the previous paragraph.

Suppose there are two candidates. We can find a Nash equilibrium of the game by studying the players' best response functions. Fix the position x_2 of candidate 2 and consider the best position for candidate 1. First suppose that $x_2 < m$. If candidate 1 takes a position to the left of x_2, then candidate 2 attracts the votes of all citizens whose favorite positions are to the right of $\frac{1}{2}(x_1+x_2)$, a set that includes the 50% of citizens whose favorite positions are to the right of m, and more. Thus candidate 2 wins, and candidate 1 loses. If candidate 1 takes a position to the right of x_2, then she wins as long as the dividing line between her supporters and those of candidate 2 is less than m (see Figure 72.1). If she is so far to the right that this dividing line lies to the right of m, then she loses. She prefers to win than to lose, and is indifferent between all the outcomes in which she wins, so her set of best responses to x_2 is the set of positions that causes the midpoint $\frac{1}{2}(x_1+x_2)$ of the line segment from x_2 to x_1 to be less than m. (If this midpoint is *equal* to m, then the candidates tie.) The condition $\frac{1}{2}(x_1+x_2) < m$ is equivalent to $x_1 < 2m - x_2$, so candidate 1's set of best responses to x_2 is the set of all positions between x_2 and $2m - x_2$ (excluding the points x_2 and $2m - x_2$).

A symmetric argument applies to the case in which $x_2 > m$. In this case candidate 1's set of best responses to x_2 is the set of all positions between $2m - x_2$ and x_2.

Finally consider the case in which $x_2 = m$. In this case candidate 1's unique best response is to choose the *same* position, m! If she chooses any other position, then she loses, whereas if she chooses m, then she ties for first place.

In summary, candidate 1's best response function is defined by

$$B_1(x_2) = \begin{cases} \{x_1 : x_2 < x_1 < 2m - x_2\} & \text{if } x_2 < m \\ \{m\} & \text{if } x_2 = m \\ \{x_1 : 2m - x_2 < x_1 < x_2\} & \text{if } x_2 > m. \end{cases}$$

Candidate 2 faces exactly the same incentives as candidate 1, and hence has the same best response function. The candidates' best response functions are shown in Figure 73.1.

If you superimpose the two best response functions, you will see that the game has a unique Nash equilibrium, in which both candidates choose the position m,

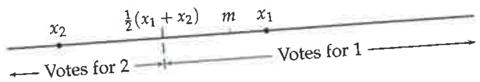

Figure 72.1 An action profile (x_1, x_2) for which candidate 1 wins.

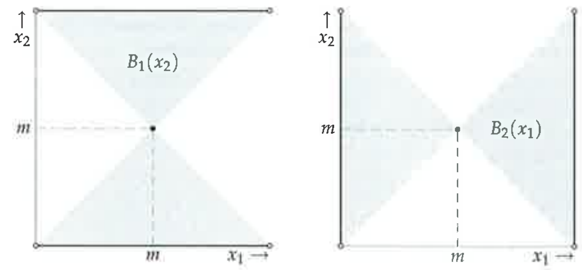

Figure 73.1 The candidates' best response functions in Hotelling's model of electoral competition with two candidates. Candidate 1's best response function is in the left panel; candidate 2's is in the right panel. (The edges of the shaded areas are excluded.)

the voters' median favorite position. (Remember that the edges of the shaded area, which correspond to pairs of positions that result in ties, are excluded from the best response functions.) The outcome is that the election is a tie.

As in the case of Bertrand's duopoly game in the previous section, we can make a direct argument that (m, m) is the unique Nash equilibrium of the game, without constructing the best response functions. First, (m, m) is an equilibrium: it results in a tie, and if either candidate chooses a position different from m, then she loses. Second, no other pair of positions is a Nash equilibrium, by the following argument.

- If one candidate loses, then she can do better by moving to m, where she either wins outright (if her opponent's position is different from m) or ties for first place (if her opponent's position is m).

- If the candidates tie (because their positions are either the same or symmetric about m), then either candidate can do better by moving to m, where she wins outright.

Our conclusion is that the competition between the candidates to secure a majority of the votes drives them to select the same position, equal to the median of the citizens' favorite positions. Hotelling (1929, 54), the originator of the model, writes that this outcome is "strikingly exemplified." He continues, "The competition for votes between the Republican and Democratic parties [in the United States] does not lead to a clear drawing of issues, an adoption of two strongly contrasted positions between which the voter may choose. Instead, each party strives to make its platform as much like the other's as possible."

? EXERCISE 73.1 (Electoral competition with asymmetric voters' preferences) Consider a variant of Hotelling's model in which voters' preferences are asymmetric. Specifically, suppose that each voter cares twice as much about policy differences to the left of her favorite position than about policy differences to the right of her favorite position. How does this affect the Nash equilibrium?

In the model considered so far, no candidate has the option of staying out of the race. Suppose that we give each candidate this option; assume that it is better than losing and worse than tying for first place. Then the Nash equilibrium remains as before: both players enter the race and choose the position m. The direct argument differs from the one before only in that in addition we need to check that there is no equilibrium in which one or both of the candidates stay out of the race. If one candidate stays out, then, given the other candidate's position, she can enter and either win outright or tie for first place. If both candidates stay out, then either candidate can enter and win outright.

The next exercise asks you to consider the Nash equilibria of this variant of the model when there are three candidates.

EXERCISE 74.1 (Electoral competition with three candidates) Consider a variant of Hotelling's model in which there are three candidates and each candidate has the option of staying out of the race, which she regards as better than losing and worse than tying for first place. Show that if less than one-third of the citizens' favorite positions are equal to the median favorite position (m), then the game has no Nash equilibrium. Argue as follows. First, show that the game has no Nash equilibrium in which a single candidate enters the race. Second, show that in any Nash equilibrium in which more than one candidate enters, all candidates that enter tie for first place. Third, show that there is no Nash equilibrium in which two candidates enter the race. Fourth, show that there is no Nash equilibrium in which all three candidates enter the race and choose the same position. Finally, show that there is no Nash equilibrium in which all three candidates enter the race and do not all choose the same position.

EXERCISE 74.2 (U.S. presidential election) Consider a variant of Hotelling's model that captures features of a U.S. presidential election. Voters are divided between two states. State 1 has more electoral college votes than does state 2. The winner is the candidate who obtains the most electoral college votes. Denote by m_i the median favorite position among the citizens of state i, for $i = 1, 2$; assume that $m_2 < m_1$. Each of two candidates chooses a single position. Each citizen votes (nonstrategically) for the candidate whose position is closest to her favorite position. The candidate who wins a majority of the votes in a state obtains all the electoral college votes of that state; if for some state the candidates obtain the same number of votes, they each obtain half of the electoral college votes of that state. Find the Nash equilibrium (equilibria?) of the strategic game that models this situation.

So far we have assumed that the candidates care only about winning; they are not at all concerned with the winner's position. The next exercise asks you to consider the case in which each candidate cares *only* about the winner's position, and not at all about winning. (You may be surprised by the equilibrium.)

EXERCISE 75.1 (Electoral competition between candidates who care only about the winning position) Consider the variant of Hotelling's model in which the candidates (like the citizens) care about the winner's position, and not at all about winning per se. There are two candidates. Each candidate has a favorite position; her dislike for other positions increases with their distance from her favorite position. Assume that the favorite position of one candidate is less than m and the favorite position of the other candidate is greater than m. Assume also that if the candidates tie when they take the positions x_1 and x_2, then the outcome is the compromise policy $\frac{1}{2}(x_1 + x_2)$. Find the set of Nash equilibria of the strategic game that models this situation. (First consider pairs (x_1, x_2) of positions for which either $x_1 < m$ and $x_2 < m$, or $x_1 > m$ and $x_2 > m$. Next consider pairs (x_1, x_2) for which either $x_1 < m < x_2$, or $x_2 < m < x_1$, then those for which $x_1 = m$ and $x_2 \neq m$, or $x_1 \neq m$ and $x_2 = m$. Finally consider the pair (m, m).)

The set of candidates in Hotelling's model is given. The next exercise asks you to analyze a model in which the set of candidates is generated as part of an equilibrium.

EXERCISE 75.2 (Citizen-candidates) Consider a game in which the players are the citizens. Any citizen may, at some cost $c > 0$, become a candidate. Assume that the only position a citizen can espouse is her favorite position, so that a citizen's only decision is whether to stand as a candidate. After all citizens have (simultaneously) decided whether to become candidates, each citizen votes for her favorite candidate, as in Hotelling's model. Citizens care about the position of the winning candidate; a citizen whose favorite position is x loses $|x - x^*|$ if the winning candidate's position is x^*. (For any number z, $|z|$ denotes the absolute value of z: $|z| = z$ if $z > 0$ and $|z| = -z$ if $z < 0$.) Winning confers the benefit b. Thus a citizen who becomes a candidate and ties with $k - 1$ other candidates for first place obtains the payoff $b/k - c$; a citizen with favorite position x who becomes a candidate and is not one of the candidates tied for first place obtains the payoff $-|x - x^*| - c$, where x^* is the winner's position; and a citizen with favorite position x who does not become a candidate obtains the payoff $-|x - x^*|$, where x^* is the winner's position. Assume that for every position x there is a citizen for whom x is the favorite position. Show that if $b \leq 2c$, then the game has a Nash equilibrium in which one citizen becomes a candidate. Is there an equilibrium (for any values of b and c) in which two citizens, each with favorite position m, become candidates? Is there an equilibrium in which two citizens with favorite positions different from m become candidates?

Hotelling's model assumes a basic agreement among the voters about the ordering of the positions. For example, if one voter prefers x to y to z and another voter prefers y to z to x, no voter prefers z to x to y. The next exercise asks you to study a model that does not so restrict the voters' preferences.

EXERCISE 75.3 (Electoral competition for more general preferences) Suppose that there is a finite number of positions and a finite, odd, number of voters. For any positions x and y, each voter either prefers x to y or prefers y to x. (That is, no voter regards any two positions as equally desirable.) We say that a position x^* is a *Condorcet winner* if for every position y different from x^*, a majority of voters prefer x^* to y.

a. Show that for any configuration of preferences there is at most one Condorcet winner.

b. Give an example in which no Condorcet winner exists. (Suppose that there are three positions (x, y, and z) and three voters. Assume that voter 1 prefers x to y to z. Construct preferences for the other two voters with the properties that one voter prefers x to y and the other prefers y to x, one prefers x to z and the other prefers z to x, and one prefers y to z and the other prefers z to y. The preferences you construct must, of course, satisfy the condition that a voter who prefers a to b and b to c also prefers a to c, where a, b, and c are any positions.)

c. Consider the strategic game in which two candidates simultaneously choose positions, as in Hotelling's model. If the candidates choose different positions, each voter endorses the candidate whose position she prefers, and the candidate who receives the most votes wins. If the candidates choose the same position, they tie. Show that this game has a unique Nash equilibrium if the voters' preferences are such that there is a Condorcet winner, and has no Nash equilibrium if the voters' preferences are such that there is no Condorcet winner.

A variant of Hotelling's model of electoral competition can be used to analyze the choices of product characteristics by competing firms in situations in which price is not a significant variable. (Think of radio stations that offer different styles of music, for example.) The set of positions is the range of possible characteristics for the product, and the citizens are consumers rather than voters. Consumers' tastes differ; each consumer buys (at a fixed price, possibly zero) one unit of the product she likes best. The model differs substantially from Hotelling's model of electoral competition in that each firm's objective is to maximize its market share, rather than to obtain a market share larger than that of any other firm. In the next exercise you are asked to show that the Nash equilibria of this game in the case of two or three firms are the same as those in Hotelling's model of electoral competition.

? EXERCISE 76.1 (Competition in product characteristics) In the variant of Hotelling's model that captures competing firms' choices of product characteristics, show that when there are two firms the unique Nash equilibrium is (m, m) (both firms offer the consumers' median favorite product), and when there are three firms there is no Nash equilibrium. (Start by arguing that when there are two firms whose products differ, either firm is better off making its product more similar to that of its rival.)

3.4 The *War of Attrition*

The game known as the *War of Attrition* elaborates on the ideas captured by the game *Hawk–Dove* (Exercise 31.2). It was originally posed as a model of a conflict between two animals fighting over prey. Each animal chooses the time at which it intends to give up. When an animal gives up, its opponent obtains all the prey (and the time at which the winner intended to give up is irrelevant). If both animals give up at the same time, then each has an equal chance of obtaining the prey. Fighting is costly: each animal prefers as short a fight as possible.

The game models not only such a conflict between animals, but also many other disputes. The "prey" can be any indivisible object, and "fighting" can be any costly activity—for example, simply waiting.

To define the game precisely, let time be a continuous variable that starts at 0 and runs indefinitely. Assume that the value party i attaches to the object in dispute is $v_i > 0$ and the value it attaches to a 50% chance of obtaining the object is $v_i/2$. Each unit of time that passes before the dispute is settled (i.e. one of the parties concedes) costs each party one unit of payoff. Thus if player i concedes first, at time t_i, her payoff is $-t_i$ (she spends t_i units of time and does not obtain the object). If the other player concedes first, at time t_j, player i's payoff is $v_i - t_j$ (she obtains the object after t_j units of time). If both players concede at the same time, player i's payoff is $\frac{1}{2}v_i - t_i$, where t_i is the common concession time. The **War of Attrition** is the following strategic game.

Players The two parties to a dispute.

Actions Each player's set of actions is the set of possible concession times (nonnegative numbers).

Preferences Player i's preferences are represented by the payoff function

$$u_i(t_1, t_2) = \begin{cases} -t_i & \text{if } t_i < t_j \\ \frac{1}{2}v_i - t_i & \text{if } t_i = t_j \\ v_i - t_j & \text{if } t_i > t_j, \end{cases}$$

where j is the other player.

To find the Nash equilibria of this game, we start, as before, by finding the players' best response functions. Intuitively, if player j's intended concession time is early enough (t_j is small), then it is optimal for player i to wait for player j to concede. That is, in this case player i should choose a concession time later than t_j; any such time is as good as any other. By contrast, if player j intends to hold out for a long time (t_j is large), then player i should concede immediately. Because player i values the object at v_i, the length of time it is worth her waiting is v_i.

To make these ideas precise, we can study player i's payoff function for various fixed values of t_j, the concession time of player j. The three cases that the intuitive argument suggests are qualitatively different are shown in Figure 78.1: $t_j < v_i$ in