3 Nash Equilibrium: Illustrations

3.1 Cournot's model of oligopoly 53
3.2 Bertrand's model of oligopoly 61
3.3 Electoral competition 68
3.4 The *War of Attrition* 75
3.5 Auctions 78
3.6 Accident law 89
Prerequisite: Chapter 2.

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 nineteenth 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

$$\pi_i(q_1, \ldots, q_n) = q_i P(q_1 + \cdots + q_n) - C_i(q_i). \tag{54.1}$$

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 (54.1).

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

$$P(Q) = \begin{cases} \alpha - Q & \text{if } Q \le \alpha \\ 0 & \text{if } Q > \alpha, \end{cases} \tag{54.2}$$

where $\alpha > 0$ and $c \ge 0$ are constants. This inverse demand function is shown in Figure 55.1. (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 $q_1 + q_2 \le \alpha$ and zero if $q_1 + q_2 > \alpha$. Thus firm 1's profit is

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

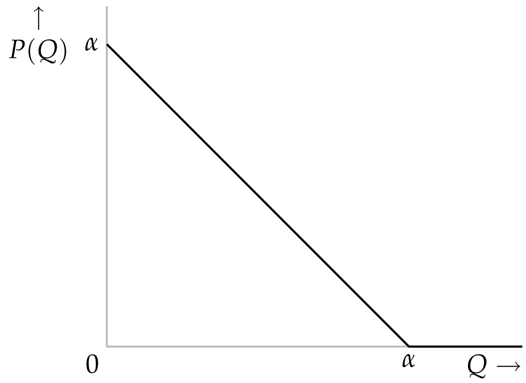

Figure 55.1 The inverse demand function in the example of Cournot's game studied in Section 3.1.3.

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 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 56.1. 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 56.1 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 the output of zero.

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

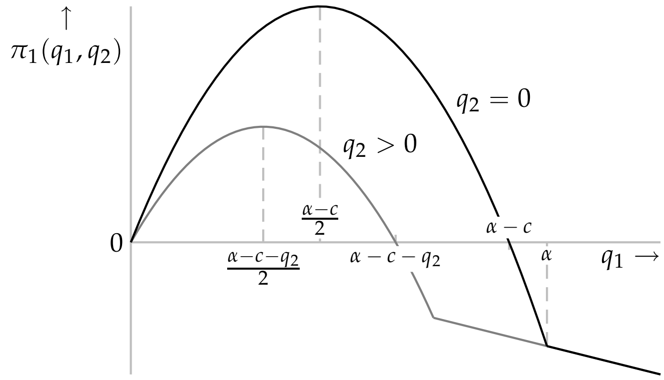

Figure 56.1 Firm 1's profit as a function of its output, given firm 2's output. The black curve shows the case $q_2 = 0$, whereas the gray curve shows a case in which $q_2 > 0$.

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 are shown in Figure 56.2 (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.

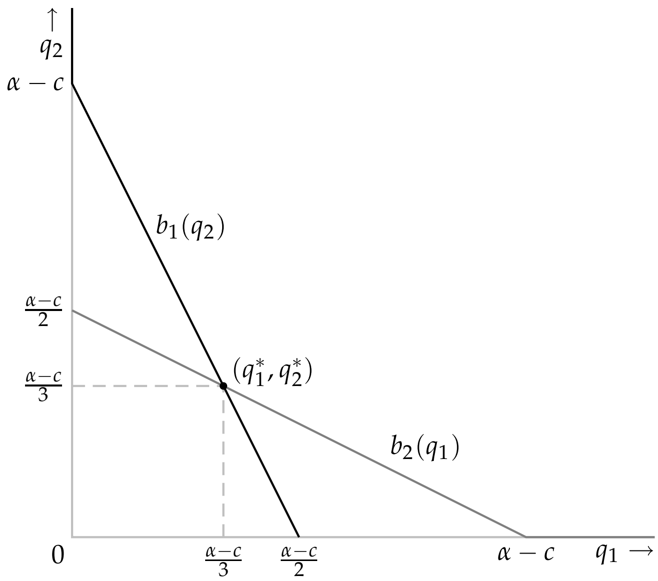

Figure 56.2 The best response functions in Cournot's duopoly game when the inverse demand function is given by (54.2) and the cost function of each firm is cq. The unique Nash equilibrium is $(q_1^*, q_2^*) = (\frac{1}{3}(\alpha - c), \frac{1}{3}(\alpha - c))$.

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 (34.3)). The set of such pairs is the set of points at which the best response functions in Figure 56.2 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 increases. 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 57.1 (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 (54.2), the cost function of each firm i is $C_i(q_i) = c_i q_i$, 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 57.2 (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 (54.2), 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 57.3 (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 (54.2), 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 $c \geq 0$, $f > 0$, and $c < \alpha$. (Note that the fixed cost f affects only the firm's decision of whether or not 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 58.1 (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. Specifically, both firms can increase their profits relative to their equilibrium levels by reducing their outputs.

? EXERCISE 58.2 (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 59.1 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 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.

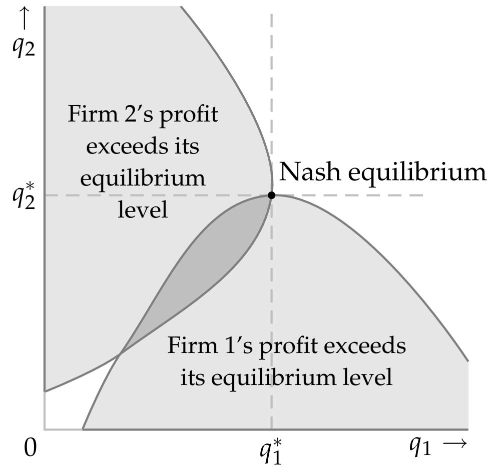

Figure 59.1 The pair (q_1^*, q_2^*) is a Nash equilibrium; along each gray curve one of the firm's profits is constant, equal to its profit at the equilibrium. The area shaded dark gray is the set of pairs of outputs at which both firms' profits exceed their equilibrium levels.

We see that if the parts of the boundaries of these sets indicated by the gray lines in the figure 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 59.1 (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 (54.2) 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 (34.3)), 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 the previous section. 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 on 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.

? EXERCISE 60.1 (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 be a village green, for example; the higher the total number of sheep grazed there, the less valuable the green is to any given farmer.

The first property of a Nash equilibrium in Cournot's model discussed in the previous section 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 59.1 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.)

? EXERCISE 61.1 (Interaction among resource-users) A group of n firms uses a common resource (a river or a forest, for example) to produce output. As more of the resource is used, 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 values of α and c such that the game is the same as the one studied in Exercise 59.1, and hence find its Nash equilibria. Find an action profile $(x_1, \ldots, x_n)$ at which each firm's output is higher than it is at the 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, equal to $p_i D(p_i)/m - C_i(D(p_i)/m)$ 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 functions 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, whether or not 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). Three cross-sections of firm i's payoff function, for different values of p_j, are shown in black in Figure 63.1. (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 63.1). 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 $B_i(p_j) = \{p_i\colon p_i > p_j\}$.

- 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: $B_i(p_j) = \{p_i\colon p_i \geq p_j\}$.

- 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 63.1). 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 63.1): $B_i(p_j) = \{p^m\}$.

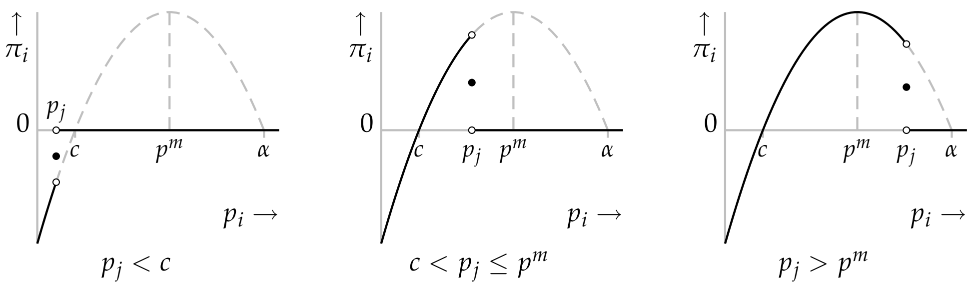

Figure 63.1 Three cross-sections (in black) of firm i's payoff function in Bertrand's duopoly game. Where the payoff function jumps, its value is given by the small disk; the small circles indicate points that are excluded as values of the functions.

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

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

where ∅ 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.

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 65.2 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 making 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 45.1) 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 64.1 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.

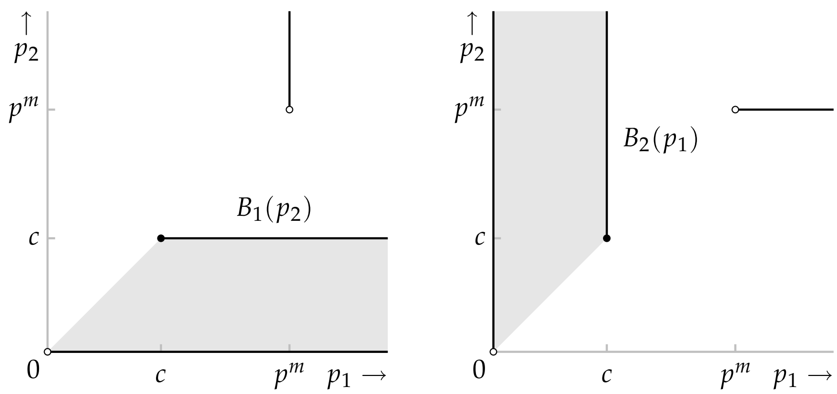

Figure 64.1 The firms' best response functions in Bertrand's duopoly game. Firm 1's best response function is in the left panel; firm 2's is in the right panel.

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 (34.2)). 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. So 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 makes 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.

? EXERCISE 65.1 (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 65.2 (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 making losses, whereas the price at the equilibrium of Cournot's game is higher. In Cournot's game, the price decreases towards c as the number of firms increases (Exercise 59.1), 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 66.1 (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 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, 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 66.2 (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 67.1 (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 67.1). 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.)

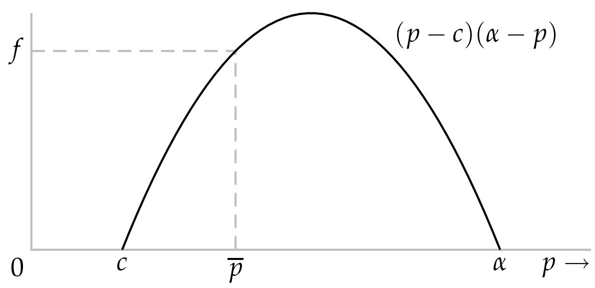

Figure 67.1 The determination of the price $\overline{p}$ in Exercise 67.1.

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 nineteenth 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 in a review of Cournot's book written by Bertrand in 1883. (Cournot's 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) than to lose, and if she ties for first place she prefers to do so with as few other candidates as possible.

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 there is only one such position.)

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 69.1.)

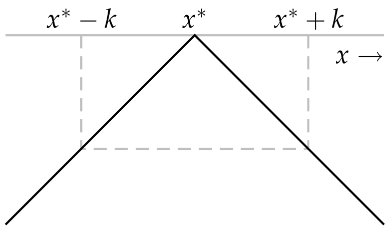

Figure 69.1 The payoff of a voter whose favorite position is x^*, as a function of the winning position, x.

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 69.2. 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.

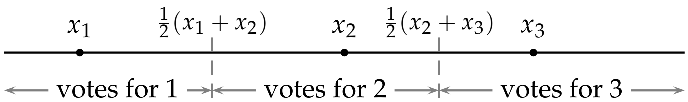

Figure 69.2 The allocation of votes between three candidates, with positions x_1, x_2, and x_3.

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).

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 $1 \leq k \leq n - 1$), 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 so long as the dividing line between her supporters and those of candidate 2 is less than m (see Figure 70.1). 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 $2m - x_2$).

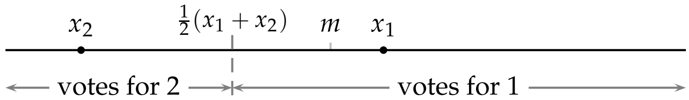

Figure 70.1 An action profile (x_1, x_2) for which candidate 1 wins.

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

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

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 71.1.

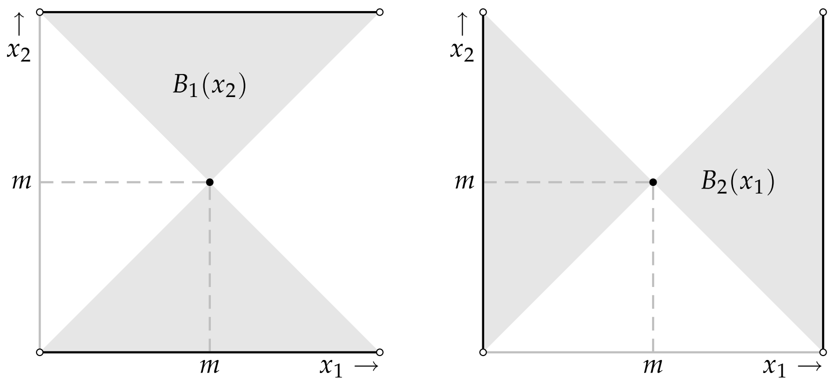

Figure 71.1 The candidates' best response functions in Hotelling's model of electoral competition with two candidates. Candidate 1's best response function is in the left panel; candidate 2's is in the right panel. (The edges of the shaded areas are excluded.)

If you superimpose the two best response functions, you see that the game has a unique Nash equilibrium, in which both candidates choose the position m, 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 USA] 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 72.1 (Electoral competition with asymmetric voters' preferences) Consider a variant of Hotelling's model in which voters's 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 stays 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 72.2 (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. Use the following arguments to show that the game has no Nash equilibrium. First, show that there is 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 72.3 (Electoral competition in two districts) Consider a variant of Hotelling's model that captures features of a US presidential election. Voters are divided between two districts. District 1 is worth more electoral college votes than is district 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 district 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 in closest to her favorite position. The candidate who wins a majority of the votes in a district obtains all the electoral college votes of that district; if the candidates obtain the same number of votes in a district, they each obtain half of the electoral college votes of that district. 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 73.1 (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 73.2 (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 74.1 (Electoral competition for more general preferences) 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. (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 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 such 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 74.2 (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 29.2). 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 they each have 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 equally good. 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 76.1: $t_j < v_i$ in the left panel, $t_j = v_i$ in the middle panel, and $t_j > v_i$ in the right panel. Player i's best responses in each case are her actions for which her payoff is highest: the set of times after t_j if $t_j < v_i$, 0 and the set of times after t_j if $t_j = v_i$, and 0 if $t_j > v_i$.

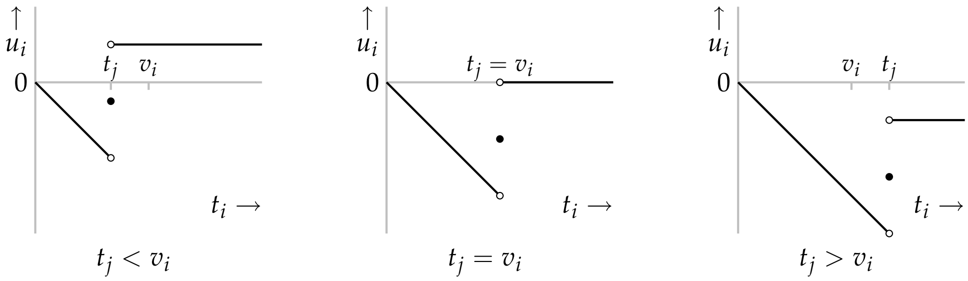

Figure 76.1 Three cross-sections of player i's payoff function in the *War of Attrition*.

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

$$B_i(t_j) = \begin{cases} \{t_i\colon t_i > t_j\} & \text{if } t_j < v_i \\ \{t_i\colon t_i = 0 \text{ or } t_i > t_j\} & \text{if } t_j = v_i \\ \{0\} & \text{if } t_j > v_i. \end{cases}$$

For a case in which $v_1 > v_2$, this function is shown in the left panel of Figure 76.2 for $i = 1$ and $j = 2$ (player 1's best response function), and in the right panel for $i = 2$ and $j = 1$ (player 2's best response function).

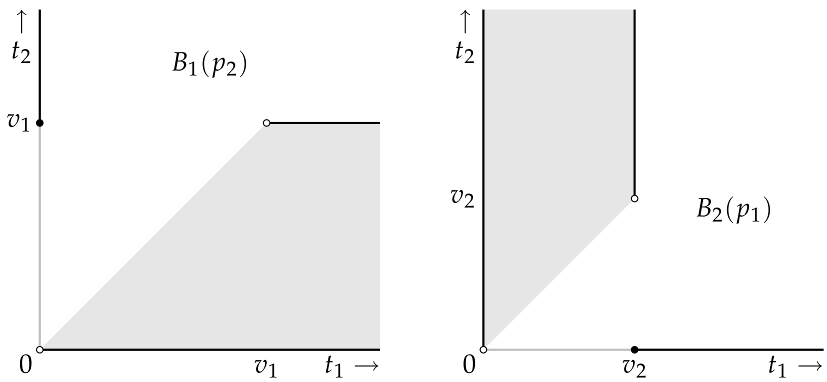

Figure 76.2 The players' best response functions in the *War of Attrition* (for a case in which $v_1 > v_2$). Player 1's best response function is in the left panel; player 2's is in the right panel. (The sloping edges are excluded.)

Superimposing the players' best response functions, we see that there are two areas of intersection: the vertical axis at and above v_1 and the horizontal axis at and to the right of v_2. Thus (t_1, t_2) is a Nash equilibrium of the game if and only if either

$$t_1 = 0 \text{ and } t_2 \geq v_1$$

or

$$t_2 = 0 \text{ and } t_1 \geq v_2.$$

In words, in every equilibrium either player 1 concedes immediately and player 2 concedes at time v_1 or later, or player 2 concedes immediately and player 1 concedes at time v_2 or later.

? EXERCISE 77.1 (Direct argument for Nash equilibria of *War of Attrition*) Give a direct argument, not using information about the entire best response functions, for the set of Nash equilibria of the *War of Attrition*. (Argue that if $t_1 = t_2$, $0 < t_i < t_j$, or $0 = t_i < t_j < v_i$ (for $i = 1$ and $j = 2$, or $i = 2$ and $j = 1$) then the pair (t_1, t_2) is not a Nash equilibrium. Then argue that any remaining pair is a Nash equilibrium.)

Three features of the equilibria are notable. First, in no equilibrium is there any fight: one player always concedes immediately. Second, either player may concede first, regardless of the players' valuations. In particular, there are always equilibria in which the player who values the object more highly concedes first. Third, the equilibria are *asymmetric* (the players' actions are different), even when $v_1 = v_2$, in which case the game is symmetric—the players' sets of actions are the same and player 1's payoff to (t_1, t_2) is the same as player 2's payoff to (t_2, t_1) (Definition 49.3). Given this asymmetry, the populations from which the two players are drawn must be distinct in order to interpret the Nash equilibria as action profiles compatible with steady states. One player might be the current owner of the object in dispute, and the other a challenger, for example. In this case the equilibria correspond to the two conventions that a challenger always gives up immediately, and that an owner always does so. (Some evidence is discussed in the box on page 412.) If all players—those in the role of player 1 as well as those in the role of player 2—are drawn from a single population, then only symmetric equilibria are relevant (see Section 2.10). The *War of Attrition* has no such equilibria, so the notion of Nash equilibrium makes no prediction about the outcome in such a situation.

? EXERCISE 77.2 (Variant of *War of Attrition*) Consider the variant of the *War of Attrition* in which each player attaches no value to the time spent waiting for the other player to concede, but the object in dispute loses value as time passes. (Think of a rotting animal carcass or a melting ice cream cone.) Assume that the value of the object to each player i after t units of time is $v_i - t$ (and the value of a 50% chance of obtaining the object is $\frac{1}{2}(v_i - t)$). Specify the strategic game that models this situation (take care with the payoff functions). Construct the analogue of Figure 76.1, find the players' best response functions, and hence find the Nash equilibria of the game.

The *War of Attrition* is an example of a "game of timing", in which each player's action is a number and each player's payoff depends sensitively on whether her action is greater or less than the other player's action. In many such games, each player's strategic variable is the time at which to act, hence the name "game of timing". The next two exercises are further examples of such games. (In the first the strategic variable is time, whereas in the second it is not.)

? EXERCISE 78.1 (Timing product release) Two firms are developing competing products for a market of fixed size. The longer a firm spends on development, the better its product. But the first firm to release its product has an advantage: the customers it obtains will not subsequently switch to its rival. (Once a person starts using a product, the cost of switching to an alternative, even one significantly better, is too high to make a switch worthwhile.) A firm that releases its product first, at time t, captures the share $h(t)$ of the market, where h is a function that increases from time 0 to time T, with $h(0) = 0$ and $h(T) = 1$. The remaining market share is left for the other firm. If the firms release their products at the same time, each obtains half of the market. Each firm wishes to obtain the highest possible market share. Model this situation as a strategic game and find its Nash equilibrium (equilibria?). (When finding firm i's best response to firm j's release time t_j, there are three cases: that in which $h(t_j) < \frac{1}{2}$ (firm j gets less than half of the market if it is the first to release its product), that in which $h(t_j) = \frac{1}{2}$, and that in which $h(t_j) > \frac{1}{2}$.)

? EXERCISE 78.2 (A fight) Each of two people has one unit of a resource. Each person chooses how much of the resource to use in fighting the other individual and how much to use productively. If each person i devotes y_i to fighting then the total output is $f(y_1, y_2) \geq 0$ and person i obtains the fraction $p_i(y_1, y_2)$ of the output, where

$$p_i(y_1, y_2) = \begin{cases} 1 & \text{if } y_i > y_j \\ \frac{1}{2} & \text{if } y_i = y_j \\ 0 & \text{if } y_i < y_j. \end{cases}$$

The function f is continuous (small changes in y_1 and y_2 cause small changes in $f(y_1, y_2)$), is decreasing in both y_1 and y_2 (the more each player devotes to fighting, the less output is produced), and satisfies $f(1, 1) = 0$ (if each player devotes all her resource to fighting then no output is produced). (If you prefer to deal with a specific function f, take $f(y_1, y_2) = 2 - y_1 - y_2$.) Each person cares only about the amount of output she receives, and prefers to receive as much as possible. Specify this situation as a strategic game and find its Nash equilibrium (equilibria?). (Use a direct argument: first consider pairs (y_1, y_2) with $y_1 \neq y_2$, then those with $y_1 = y_2 < 1$, then those with $y_1 = y_2 = 1$.)

3.5 Auctions

3.5.1 Introduction

In an "auction", a good is sold to the party who submits the highest bid. Auctions, broadly defined, are used to allocate significant economic resources, from works of art to short-term government bonds to offshore tracts for oil and gas exploration to the radio spectrum. They take many forms. For example, bids may be called out sequentially (as in auctions for works of art) or may be submitted in sealed envelopes; the price paid may be the highest bid, or some other price; if more than one unit of a good is being sold, bids may be taken on all units simultaneously, or the units may be sold sequentially. A game-theoretic analysis helps us to understand the consequences of various auction designs; it suggests, for example, the design likely to be the most effective at allocating resources, and the one likely to raise the most revenue. In this section I discuss auctions in which every buyer knows her own valuation and every other buyer's valuation of the item being sold. Chapter 9 develops tools that allow us to study, in Section 9.6, auctions in which buyers are not perfectly informed of each other's valuations.

AUCTIONS FROM BABYLONIA TO EBAY

Auctioning has a very long history. Herodotus, a Greek writer of the fifth century BC who, together with Thucydides, created the intellectual field of history, describes auctions in Babylonia. He writes that the Babylonians' "most sensible" custom was an annual auction in each village of the women of marriageable age. The women most attractive to the men were sold first; they commanded positive prices, whereas men were paid to be matched with the least desirable women. In each auction, bids appear to have been called out sequentially, the man who bid the most winning and paying the price he bid.

Auctions were also used in Athens in the fifth and fourth centuries BC to sell the rights to collect taxes, to dispose of confiscated property, and to lease land and mines. The evidence on the nature of the auctions is slim, but some interesting accounts survive. For example, the Athenian politician Andocides (c. 440–391 BC) reports collusive behavior in an auction of tax-collection rights (see Langdon 1994, 260).

Auctions were frequent in ancient Rome, and continued to be used in medieval Europe after the end of the Roman empire (tax-collection rights were annually auctioned by the towns of the medieval and early modern Low Countries, for example). The earliest use of the English word "auction" given by the *Oxford English Dictionary* dates from 1595, and concerns an auction "when will be sold Slaves, household goods, etc.". Rules surviving from the auctions of this era show that in some cases, at least, bids were called out sequentially, with the bidder remaining at the end obtaining the object at the price she bid (Cassady 1967, 30–31). A variant of this mechanism, in which a time limit is imposed on the bids, is reported by the English diarist and naval administrator Samuel Pepys (1633–1703). The auctioneer lit a short candle, and bids were valid only if made before the flame went out. Pepys reports that a flurry of bidding occurred at the last moment. At an auction on September 3, 1662, a bidder "cunninger than the rest" told him that just as the flame goes out, "the smoke descends", signaling the moment at which one should bid, an observation Pepys found "very pretty" (Pepys 1970, 185–186).

The auction houses of Sotheby's and Christie's were founded in the mid-18th century. At the beginning of the twenty-first century, they are being eclipsed, at least in the value of the goods they sell, by online auction companies. For example, eBay, founded in September 1995, sold US$1.3 billion of merchandise in 62 million auctions during the second quarter of 2000, roughly double the numbers for the second quarter of the previous year; Sotheby's and Christie's together sell around US$1 billion of art and antiques each quarter.

The mechanism used by eBay shares a feature with the ones Pepys observed: all bids must be received before some fixed time. The way in which the price is determined differs. In an eBay auction, a bidder submits a "proxy bid" that is not revealed; the prevailing price is a small increment above the second-highest proxy bid. As in the 17th century auctions Pepys observed, many bidders on eBay act at the last moment—a practice known as "sniping" in the argot of cyberspace. Other online auction houses use different termination rules. For example, Amazon waits ten minutes after a bid before closing an auction. The fact that last-minute bidding is much less common in Amazon auctions than it is in eBay auctions has attracted the attention of game theorists, who have begun to explore models that explain it in terms of the difference in the auctions' termination rules (see, for example, Ockenfels and Roth 2002).

In recent years, many countries have auctioned the rights to the radio spectrum, used for wireless communication. These auctions have been much studied by game theorists; they are discussed in the box on page 300.

3.5.2 Second-price sealed-bid auctions

In a common form of auction, people sequentially submit increasing bids for an object. (The word "auction" comes from the Latin *augere*, meaning "to increase".) When no one wishes to submit a bid higher than the current bid, the person making the current bid obtains the object at the price she bid.

Given that every person is certain of her valuation of the object before the bidding begins, during the bidding no one can learn anything relevant to her actions. Thus we can model the auction by assuming that each person decides, before bidding begins, the most she is willing to bid—her "maximal bid". When the players carry out their plans, the winner is the person whose maximal bid is highest. How much does she need to bid? Eventually only she and the person with the second highest maximal bid will be left competing against each other. In order to win, she therefore needs to bid slightly more than the *second highest* maximal bid. If the bidding increment is small, we can take the price the winner pays to be *equal* to the second highest maximal bid.

Thus we can model such an auction as a strategic game in which each player chooses an amount of money, interpreted as the *maximal* amount she is willing to bid, and the player who chooses the highest amount obtains the object and pays a price equal to the second highest amount.

This game models also a situation in which the people simultaneously put bids in sealed envelopes, and the person who submits the highest bid wins and pays a price equal to the *second* highest bid. For this reason the game is called a *second-price sealed-bid* auction.

To define the game precisely, denote by v_i the value player i attaches to the object; if she obtains the object at the price p her payoff is $v_i - p$. Assume that the players' valuations of the object are all different and all positive; number the players 1 through n in such a way that $v_1 > v_2 > \cdots > v_n > 0$. Each player i submits a (sealed) bid b_i. If player i's bid is higher than every other bid, she obtains the object at a price equal to the second-highest bid, say b_j, and hence receives the payoff $v_i - b_j$. If some other bid is higher than player i's bid, player i does not obtain the object, and receives the payoff of zero. If player i is in a tie for the highest bid, her payoff depends on the way in which ties are broken. A simple (though arbitrary) assumption is that the winner is the player among those submitting the highest bid whose number is smallest (i.e. whose valuation of the object is highest). (If the highest bid is submitted by players 2, 5, and 7, for example, the winner is player 2.) Under this assumption, player i's payoff when she bids b_i and is in a tie for the highest bid is $v_i - b_i$ if her number is lower than that of any other player submitting the bid b_i, and 0 otherwise.

In summary, a **second-price sealed-bid auction** (with complete information) is the following strategic game.

Players The n bidders, where $n \geq 2$.

Actions The set of actions of each player is the set of possible bids (nonnegative numbers).

Preferences The payoff of any player i is $v_i - b_j$, where b_j is the highest bid submitted by a player other than i if either b_i is higher than every other bid, or b_i is at least as high as every other bid and the number of every other player who bids b_i is greater than i. Otherwise player i's payoff is 0.

This game has many Nash equilibria. One equilibrium is $(b_1, \ldots, b_n) = (v_1, \ldots, v_n)$: each player's bid is equal to her valuation of the object. Because $v_1 > v_2 > \cdots > v_n$, the outcome is that player 1 obtains the object at the price b_2; her payoff is $v_1 - b_2$ and every other player's payoff is zero. This profile is a Nash equilibrium by the following argument.

- If player 1 changes her bid to some other price at least equal to b_2 then the outcome does not change (recall that she pays the *second* highest bid, not the highest bid). If she changes her bid to a price less than b_2 then she loses and obtains the payoff of zero.
- If some other player lowers her bid or raises it to some price at most equal to b_1 then she remains a loser; if she raises her bid above b_1 then she wins but, in paying the price b_1, makes a loss (because her valuation is less than b_1).

Another equilibrium is $(b_1, \ldots, b_n) = (v_1, 0, \ldots, 0)$. In this equilibrium, player 1 obtains the object and pays the price of zero. The profile is an equilibrium because if player 1 changes her bid then the outcome remains the same, and if any of the remaining players raises her bid then either the outcome remains the same (if her new bid is at most v_1) or causes her to obtain the object at a price that exceeds her valuation (if her bid exceeds v_1). (The auctioneer obviously has an incentive for the price to be bid up, but she is not a player in the game!)

In both of these equilibria, player 1 obtains the object. But there are also equilibria in which player 1 does not obtain the object. Consider, for example, the action profile $(v_2, v_1, 0, \ldots, 0)$, in which player 2 obtains the object at the price v_2 and every player (including player 2) receives the payoff of zero. This action profile is a Nash equilibrium by the following argument.

- If player 1 raises her bid to v_1 or more, she wins the object but her payoff remains zero (she pays the price v_1, bid by player 2). Any other change in her bid has no effect on the outcome.
- If player 2 changes her bid to some other price greater than v_2, the outcome does not change. If she changes her bid to v_2 or less she loses, and her payoff remains zero.
- If any other player raises her bid to at most v_1, the outcome does not change. If she raises her bid above v_1 then she wins, but in paying the price v_1 (bid by player 2) she obtains a negative payoff.

? EXERCISE 82.1 (Nash equilibrium of second-price sealed-bid auction) Find a Nash equilibrium of a second-price sealed-bid auction in which player n obtains the object.

Player 2's bid in this equilibrium exceeds her valuation, and thus may seem a little rash: if player 1 were to increase her bid to any value less than v_1, player 2's payoff would be negative (she would obtain the object at a price greater than her valuation). This property of the action profile does not affect its status as an equilibrium, because in a Nash equilibrium a player does not consider the "risk" that another player will take an action different from her equilibrium action; each player simply chooses an action that is optimal, *given* the other players' actions. But the property does suggest that the equilibrium is less plausible as the outcome of the auction than the equilibrium in which every player bids her valuation.

The same point takes a different form when we interpret the strategic game as a model of events that unfold over time. Under this interpretation, player 2's action v_1 means that she will continue bidding until the price reaches v_1. If player 1 is *sure* that player 2 will continue bidding until the price is v_1, then player 1 rationally stops bidding when the price reaches v_2 (or, indeed, when it reaches any other level at most equal to v_1). But there is little reason for player 1 to believe that player 2 will in fact stay in the bidding if the price exceeds v_2: player 2's action is not credible, because if the bidding were to go above v_2, player 2 would rationally withdraw.

The weakness of the equilibrium is reflected in the fact that player 2's bid v_1 is weakly dominated by the bid v_2. More generally,

> *in a second-price sealed-bid auction (with complete information), a player's bid equal to her valuation weakly dominates all her other bids.*

That is, for any bid $b_i \neq v_i$, player i's bid v_i is at least as good as b_i, *no matter what the other players bid*, and is better than b_i for some actions of the other players. (See Definition 45.1.) A player who bids less than her valuation stands not to win in some cases in which she could profit by winning (when the highest of the other bids is between her bid and her valuation), and never stands to gain relative to the situation in which she bids her valuation; a player who bids more than her valuation stands to win in some cases in which she obtains a negative payoff by doing so (when the highest of the remaining bids is between her valuation and her bid), and never stands to gain relative to the situation in which she bids her valuation. The key point is that in a second-price auction, a player who changes her bid does not lower the price she pays, but only possibly changes her status from that of a winner into that of a loser, or vice versa.

A precise argument is shown in Figure 83.1, which compares player i's payoffs to the bid v_i with her payoffs to a bid $b_i < v_i$ (top table), and to a bid $b_i > v_i$ (bottom table), as a function of the highest of the other players' bids, denoted $\overline{b}$. In each case, for all bids of the other players, player i's payoffs to v_i are at least as large as her payoffs to the other bid, and for bids of the other players such that $\overline{b}$ is in the middle column of each table, player i's payoffs to v_i are greater than her payoffs to the other bid. Thus player i's bid v_i weakly dominates all her other bids.

		Highest of other players' bids, $\overline{b}$		
		$\overline{b} < b_i$ or $\overline{b} = b_i$ & b_i wins	$b_i < \overline{b} < v_i$ or $\overline{b} = b_i$ & b_i loses	$\overline{b} > v_i$
i's bid	$b_i < v_i$	$v_i - \overline{b}$	0	0
	v_i	$v_i - \overline{b}$	$v_i - \overline{b}$	0

		$\overline{b} \leq v_i$	$v_i < \overline{b} < b_i$ or $\overline{b} = b_i$ & b_i wins	$\overline{b} > b_i$ or $\overline{b} = b_i$ & b_i loses
i's bid	v_i	$v_i - \overline{b}$	0	0
	$b_i > v_i$	$v_i - \overline{b}$	$v_i - \overline{b}$ (< 0)	0

Figure 83.1 Player i's payoffs in a second-price sealed-bid auction, as a function of the highest of the other player's bids, denoted $\overline{b}$. The top table gives her payoffs to the bids $b_i < v_i$ and v_i, and the bottom table gives her payoffs to the bids v_i and $b_i > v_i$.

In summary, a second-price auction has many Nash equilibria, but the equilibrium $(b_1, \ldots, b_n) = (v_1, \ldots, v_n)$ in which every player's bid is equal to her valuation of the object is distinguished by the fact that every player's action weakly dominates all her other actions.

? EXERCISE 84.1 (Second-price sealed-bid auction with two bidders) Find *all* the Nash equilibria of a second-price sealed-bid auction with two bidders. (Construct the players' best response functions. Apart from a difference in the tie-breaking rule, the game is the same as the one in Exercise 77.2.)

? EXERCISE 84.2 (Auctioning the right to choose) An action affects each of two people. The right to choose the action is sold in a second-price auction. That is, the two people simultaneously submit bids, and the one who submits the higher bid chooses her favorite action and pays (to a third party) the amount bid by the *other* person, who pays nothing. (Assume that if the bids are the same, person 1 is the winner.) For $i = 1, 2$, the payoff of person i when the action is a and person i pays m is $u_i(a) - m$. In the game that models this situation, find for each player a bid that weakly dominates all the player's other bids (and thus find a Nash equilibrium in which each player's equilibrium action weakly dominates all her other actions).

3.5.3 First-price sealed-bid auctions

A first-price auction differs from a second-price auction only in that the winner pays the price she bids, not the second highest bid. Precisely, a **first-price sealed-bid auction** (with complete information) is defined as follows.

Players The n bidders, where $n \geq 2$.

Actions The set of actions of each player is the set of possible bids (nonnegative numbers).

Preferences The payoff of any player i is $v_i - b_i$ if either b_i is higher than every other bid, or b_i is at least as high as every other bid and the number of every other player who bids b_i is greater than i. Otherwise player i's payoff is 0.

This game models an auction in which people submit sealed bids and the highest bid wins. (You conduct such an auction when you solicit offers for a car you wish to sell, or, as a buyer, get estimates from contractors to fix your leaky basement, assuming in both cases that you do not inform potential bidders of existing bids.) The game models also a dynamic auction in which the auctioneer begins by announcing a high price, which she gradually lowers until someone indicates her willingness to buy the object. (Flowers in the Netherlands are sold in this way.) A bid in the strategic game is interpreted as the price at which the bidder will indicate her willingness to buy the object in the dynamic auction.

One Nash equilibrium of a first-price sealed-bid auction is $(b_1, \ldots, b_n) = (v_2, v_2, v_3, \ldots, v_n)$, in which player 1's bid is player 2's valuation v_2 and every other player's bid is her own valuation. The outcome of this equilibrium is that player 1 obtains the object at the price v_2.

? EXERCISE 84.3 (Nash equilibrium of first-price sealed-bid auction) Show that $(b_1, \ldots, b_n) = (v_2, v_2, v_3, \ldots, v_n)$ is a Nash equilibrium of a first-price sealed-bid auction.

A first-price sealed-bid auction has many other equilibria, but in all equilibria the winner is the player who values the object most highly (player 1), by the following argument. In any action profile $(b_1, \ldots, b_n)$ in which some player $i \neq 1$ wins, we have $b_i > b_1$. If $b_i > v_2$ then i's payoff is negative, so that she can do better by reducing her bid to 0; if $b_i \leq v_2$ then player 1 can increase her payoff from 0 to $v_1 - b_i$ by bidding b_i, in which case she wins. Thus no such action profile is a Nash equilibrium.

? EXERCISE 85.1 (First-price sealed-bid auction) Show that in a Nash equilibrium of a first-price sealed-bid auction the two highest bids are the same, one of these bids is submitted by player 1, and the highest bid is at least v_2 and at most v_1. Show also that any action profile satisfying these conditions is a Nash equilibrium.

In any equilibrium in which the winning bid exceeds v_2, at least one player's bid exceeds her valuation. As in a second-price sealed-bid auction, such a bid seems "risky", because it would yield the bidder a negative payoff if it were to win. In the equilibrium there is no risk, because the bid does not win; but, as before, the fact that the bid has this property reduces the plausibility of the equilibrium.

As in a second-price sealed-bid auction, the potential "riskiness" to player i of a bid $b_i > v_i$ is reflected in the fact that it is weakly dominated by the bid v_i, as shown by the following argument.

- If the other players' bids are such that player i loses when she bids b_i, then the outcome is the same whether she bids b_i or v_i.
- If the other players' bids are such that player i wins when she bids b_i, then her payoff is negative when she bids b_i and zero when she bids v_i (whether or not this bid wins).

However, in a first-price auction, unlike a second-price auction, a bid $b_i < v_i$ of player i is *not* weakly dominated by the bid v_i. If fact, such a bid is not weakly dominated by *any* bid. It is not weakly dominated by a bid $b_i' < b_i$, because if the other players' highest bid is between b_i' and b_i then b_i' loses whereas b_i wins and yields player i a positive payoff. And it is not weakly dominated by a bid $b_i' > b_i$, because if the other players' highest bid is less than b_i then both b_i and b_i' win and b_i yields a lower price.

Further, even though the bid v_i weakly dominates higher bids, this bid is itself weakly dominated, by a lower bid! If player i bids v_i her payoff is 0 regardless of the other players' bids, whereas if she bids less than v_i her payoff is either 0 (if she loses) or positive (if she wins).

In summary,

> *in a first-price sealed-bid auction (with complete information), a player's bid of at least her valuation is weakly dominated, and a bid of less than her valuation is not weakly dominated.*

An implication of this result is that in *every* Nash equilibrium of a first-price sealed-bid auction at least one player's action is weakly dominated. However, this property of the equilibria depends on the assumption that a bid may be any number. In the variant of the game in which bids and valuations are restricted to be multiples of some discrete monetary unit ϵ (e.g. a cent), an action profile $(v_2 - \epsilon, v_2 - \epsilon, b_3, \ldots, b_n)$ for any $b_j \leq v_j - \epsilon$ for $j = 3, \ldots, n$ is a Nash equilibrium in which no player's bid is weakly dominated. Further, every equilibrium in which no player's bid is weakly dominated takes this form. When ϵ is small, each such equilibrium is close to an equilibrium $(v_2, v_2, b_3, \ldots, b_n)$ (with $b_j \leq v_j$ for $j = 3, \ldots, n$) of the game with unrestricted bids. On this (somewhat *ad hoc*) basis, I select action profiles $(v_2, v_2, b_3, \ldots, b_n)$ with $b_j \leq v_j$ for $j = 3, \ldots, n$ as "distinguished" equilibria of a first-price sealed-bid auction.

One conclusion of this analysis is that while both second-price and first-price auctions have many Nash equilibria, yielding a variety of outcomes, their distinguished equilibria yield the *same* outcome. (Recall that the distinguished equilibrium of a second-price sealed-bid auction is the action profile in which every player bids her valuation.) In every distinguished equilibrium of each game, the object is sold to player 1 at the price v_2. In particular, the auctioneer's revenue is the same in both cases. Thus if we restrict attention to the distinguished equilibria, the two auction forms are "revenue equivalent". The rules are different, but the players' equilibrium bids adjust to the difference and lead to the same outcome:

> *the single Nash equilibrium in which no player's bid is weakly dominated in a second-price auction yields the same outcome as the distinguished equilibria of a first-price auction.*

? EXERCISE 86.1 (Third-price auction) Consider a *third*-price sealed-bid auction, which differs from a first- and a second-price auction only in that the winner (the person who submits the highest bid) pays the third highest price. (Assume that there are at least three bidders.)

a. Show that for any player i the bid of v_i weakly dominates any lower bid, but does not weakly dominate any higher bid. (To show the latter, for any bid $b_i > v_i$ find bids for the other players such that player i is better off bidding b_i than bidding v_i.)

b. Show that the action profile in which each player bids her valuation is not a Nash equilibrium.

c. Find a Nash equilibrium. (There are ones in which every player submits the same bid.)

3.5.4 Variants

Uncertain valuations One respect in which the models in this section depart from reality is in the assumption that each bidder is certain of both her own valuation and every other bidder's valuation. In most, if not all, actual auctions, information is surely less complete. The case in which the players are uncertain about each other's valuations has been thoroughly explored, and is discussed in Section 9.6. The result that a player's bidding her valuation weakly dominates all her other actions in a second-price auction survives when players are uncertain about each other's valuations, as does the revenue-equivalence of first- and second-price auctions under some conditions on the players' preferences.

Common valuations In some auctions the main difference between the bidders is not that the value the object differently but that they have different information about its value. For example, the bidders for an oil tract may put similar values on any given amount of oil, but have different information about how much oil is in the tract. Such auctions involve informational considerations that do not arise in the model we have studied in this section; they are studied in Section 9.6.3.

Multi-unit auctions In some auctions, like those for Treasury Bills (short-term government bonds) in the USA, many units of an object are available, and each bidder may value positively more than one unit. In each of the types of auction described below, each bidder submits a bid for each unit of the good. That is, an action is a list of bids $(b^1, \ldots, b^k)$, where b^1 is the player's bid for the first unit of the good, b^2 is her bid for the second unit, and so on. The player who submits the highest bid for any given unit obtains that unit. The auctions differ in the prices paid by the winners. (The first type of auction generalizes a first-price auction, whereas the next two generalize a second-price auction.)

Discriminatory auction The price paid for each unit is the winning bid for that unit.

Uniform-price auction The price paid for each unit is the same, equal to the highest rejected bid among all the bids for all units.

Vickrey auction A bidder who wins k objects pays the sum of the k highest rejected bids submitted by the *other* bidders.

The next exercise asks you to study these auctions when two units of an object are available.

? EXERCISE 87.1 (Multi-unit auctions) Two units of an object are available. There are n bidders. Bidder i values the first unit that she obtains at v_i and the second unit at w_i, where $v_i > w_i > 0$. Each bidder submits two bids; the two highest bids win. Retain the tie-breaking rule in the text. Show that in discriminatory and uniform-price auctions, player i's action of bidding v_i and w_i does not dominate all her other actions, whereas in a Vickrey auction it does. (In the case of a Vickrey auction, consider separately the cases in which the other players' bids are such that player i wins no units, one unit, and two units when her bids are v_i and w_i.)

Goods for which the demand exceeds the supply at the going price are sometimes sold to the people who are willing to wait longest in line. We can model such situations as multi-unit auctions in which each person's bid is the amount of time she is willing to wait.

⁇ EXERCISE 88.1 (Waiting in line) Two hundred people are willing to wait in line to see a movie at a theater whose capacity is one hundred. Denote person i's valuation of the movie in excess of the price of admission, expressed in terms of the amount of time she is willing to wait, by v_i. That is, person i's payoff if she waits for t_i units of time is $v_i - t_i$. Each person attaches no value to a second ticket, and cannot buy tickets for other people. Assume that $v_1 > v_2 > \cdots > v_{200}$. Each person chooses an arrival time. If several people arrive at the same time then their order in line is determined by their index (lower-numbered people go first). If a person arrives to find 100 or more people already in line, her payoff is zero. Model the situation as a variant of a discriminatory multi-unit auction, in which each person submits a bid for only one unit, and find its Nash equilibria. (Look at your answer to Exercise 85.1 before seeking the Nash equilibria.) Arrival times for people at movies do not in general seem to conform with a Nash equilibrium. What feature missing from the model could explain the pattern of arrivals?

The next exercise is another application of a multi-unit auction. As in the previous exercise each person wants to buy only one unit, but in this case the price paid by the winners is the highest losing bid.

? EXERCISE 88.2 (Internet pricing) A proposal to deal with congestion on electronic message pathways is that each message should include a field stating an amount of money the sender is willing to pay for the message to be sent. Suppose that during some time interval, each of n people wants to send one message and the capacity of the pathway is k messages, with $k < n$. The k messages whose bids are highest are the ones sent, and each of the persons sending these messages pays a price equal to the $(k+1)$st highest bid. Model this situation as a multi-unit auction. (Use the same tie-breaking rule as the one in the text.) Does a person's action of bidding the value of her message weakly dominate all her other actions? (Note that the auction differs from those considered in Exercise 87.1 because each person submits only one bid. Look at the argument in the text that in a second-price sealed-bid auction a player's action of bidding her value weakly dominates all her other actions.)

Lobbying as an auction Variants of the models in this section can be used to understand some situations that are not explicitly auctions. An example, illustrated in the next exercise, is the competition between groups pressuring a government to follow policies they favor. This exercise shows also that the outcome of an auction may depend significantly (and perhaps counterintuitively) on the form the auction takes.

? EXERCISE 88.3 (Lobbying as an auction) A government can pursue three policies, x, y, and z. The monetary values attached to these policies by two interest groups, A and B, are given in Figure 89.1. The government chooses a policy in response to the payments the interest groups make to it. Consider the following two mechanisms.

First-price auction Each interest group chooses a policy and an amount of money it is willing to pay. The government chooses the policy proposed by the group willing to pay the most. This group makes its payment to the government, and the losing group makes no payment.

Menu auction Each interest group states, for each policy, the amount it is willing to pay to have the government implement that policy. The government chooses the policy for which the sum of the payments the groups are willing to make is the highest, and *each* group pays the government the amount of money it is willing to pay for that policy.

In each case each interest group's payoff is the value it attaches to the policy implemented minus the payment it makes. Assume that a tie is broken by the government's choosing the policy, among those tied, whose name is first in the alphabet.

	x	y	z
Interest group A	0	3	-100
Interest group B	0	-100	3

Figure 89.1 The values of the interest groups for the policies x, y, and z in Exercise 88.3.

Show that the first-price auction has a Nash equilibrium in which lobby A says it will pay 103 for y, lobby B says it will pay 103 for z, and the government's revenue is 103. Show that the menu auction has a Nash equilibrium in which lobby A announces that it will pay 3 for x, 6 for y, and 0 for z, and lobby B announces that it will pay 3 for x, 0 for y, and 6 for z, and the government chooses x, obtaining a revenue of 6. (In each case the pair of actions given is in fact the unique equilibrium.)

3.6 Accident law

3.6.1 Introduction

In some situations, laws influence the participants' payoffs and hence their actions. For example, a law may provide for the victim of an accident to be compensated by a party who was at fault, and the size of the compensation may affect the care that each party takes. What laws can we expect to produce socially desirable outcomes? A game-theoretic analysis is useful in addressing this question.

3.6.2 The game

Consider the interaction between an *injurer* (player 1) and a *victim* (player 2). The victim suffers a loss that depends on the amounts of care taken by both her and the injurer. (How badly you hurt yourself when you fall down on the sidewalk in front of my house depends on both how well I have cleared the ice and how carefully you tread.) Denote by a_i the amount of care player i takes, measured in monetary terms, and by $L(a_1, a_2)$ the loss, also measured in monetary terms, suffered by the victim, as a function of the amounts of care. (In many cases the victim does not suffer a loss with certainty, but only with probability less than one. In such cases we can interpret $L(a_1, a_2)$ as the expected loss—the average loss suffered over many occurrences.) Assume that $L(a_1, a_2) > 0$ for all values of (a_1, a_2), and that more care taken by either player reduces the loss: L is decreasing in a_1 for any fixed value of a_2, and decreasing in a_2 for any fixed value of a_1.

A legal rule determines the fraction of the loss borne by the injurer, as a function of the amounts of care taken. Denote this fraction by $\rho(a_1, a_2)$. If $\rho(a_1, a_2) = 0$ for all (a_1, a_2), for example, the victim bears the entire loss, regardless of how much care she takes or how little care the injurer takes. At the other extreme, $\rho(a_1, a_2) = 1$ for all (a_1, a_2) means that the victim is fully compensated by the injurer no matter how careless she is or how careful the injurer is.

If the amounts of care are (a_1, a_2) then the injurer bears the cost a_1 of taking care and suffers the loss of $L(a_1, a_2)$, of which she bears the fraction $\rho(a_1, a_2)$. Thus the injurer's payoff is

$$-a_1 - \rho(a_1, a_2)L(a_1, a_2).$$

Similarly, the victim's payoff is

$$-a_2 - (1 - \rho(a_1, a_2))L(a_1, a_2).$$

For any given legal rule, embodied in ρ, we can model the interaction between the injurer and victim as the following strategic game.

Players The injurer and the victim.

Actions The set of actions of each player is the set of possible levels of care (nonnegative numbers).

Preferences The injurer's preferences are represented by the payoff function $-a_1 - \rho(a_1, a_2)L(a_1, a_2)$ and the victim's preferences are represented by the payoff function $-a_2 - (1 - \rho(a_1, a_2))L(a_1, a_2)$, where a_1 is the injurer's level of care and a_2 is the victim's level of care.

How do the equilibria of this game depend upon the legal rule? Do any legal rules lead to socially desirable equilibrium outcomes?

I restrict attention to a class of legal rules known as *negligence with contributory negligence*. (This class was established in the USA in the mid-nineteenth century, and prevailed until the mid-1970s.) Each rule in this class requires the injurer to compensate the victim for a loss if and only if *both* the victim is sufficiently careful *and* the injurer is sufficiently careless; the required compensation is the total loss. Rules in the class differ in the standards of care they specify for each party. The rule that specifies the standards of care X_1 for the injurer and X_2 for the victim requires the injurer to pay the victim the entire loss $L(a_1, a_2)$ when $a_1 < X_1$ (the injurer is insufficiently careful) and $a_2 \geq X_2$ (the victim is sufficiently careful), and nothing otherwise. That is, under this rule the fraction $\rho(a_1, a_2)$ of the loss borne by the injurer is

$$\rho(a_1, a_2) = \begin{cases} 1 & \text{if } a_1 < X_1 \text{ and } a_2 \geq X_2 \\ 0 & \text{if } a_1 \geq X_1 \text{ or } a_2 < X_2. \end{cases}$$

Included in this class of rules are those for which X_1 is a positive finite number and $X_2 = 0$ (the injurer has to pay if she is not sufficiently careful, even if the victim takes no care at all), known as rules of *pure negligence*, and that for which X_1 is infinite and $X_2 = 0$ (the injurer has to pay regardless of how careful she is and how careless the victim is), known as the rule of *strict liability*.

3.6.3 Nash equilibrium

Suppose we decide that the pair $(\hat{a}_1, \hat{a}_2)$ of actions is socially desirable. We wish to answer the question: are there values of X_1 and X_2 such that the game generated by the rule of negligence with contributory negligence for (X_1, X_2) has $(\hat{a}_1, \hat{a}_2)$ as its unique Nash equilibrium? If the answer is affirmative, then, assuming the solution concept of Nash equilibrium is appropriate for the situation we are considering, we have found a legal rule that induces the socially desirable outcome.

Specifically, suppose that we select as socially desirable the pair $(\hat{a}_1, \hat{a}_2)$ of actions that maximizes the sum of the players' payoffs. That is,

$$(\hat{a}_1, \hat{a}_2) \text{ maximizes } -a_1 - a_2 - L(a_1, a_2).$$

(For some functions L, this pair $(\hat{a}_1, \hat{a}_2)$ may be a reasonable candidate for a socially desirable outcome; in other cases it may induce a very inequitable distribution of payoff between the players, and thus be an unlikely candidate.)

I claim that the unique Nash equilibrium of the game induced by the legal rule of negligence with contributory negligence for $(X_1, X_2) = (\hat{a}_1, \hat{a}_2)$ is $(\hat{a}_1, \hat{a}_2)$. That is, if the standards of care are equal to their socially desirable levels, then these are the levels chosen by an injurer and a victim in the only equilibrium of the game. The outcome is that the injurer pays no compensation: her level of care is $\hat{a}_1$, just high enough that $\rho(a_1, a_2) = 0$. At the same time the victim's level of care is $\hat{a}_2$, high enough that if the injurer reduces her level of care even slightly then she has to pay full compensation.

I first argue that $(\hat{a}_1, \hat{a}_2)$ is *a* Nash equilibrium of the game, then show that it is the *only* equilibrium. To show that $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium, I need to show that the injurer's action $\hat{a}_1$ is a best response to the victim's action $\hat{a}_2$ and *vice versa*.

Injurer's action Given that the victim's action is $\hat{a}_2$, the injurer has to pay compensation if and only if $a_1 < \hat{a}_1$. Thus the injurer's payoff is

$$u_1(a_1, \hat{a}_2) = \begin{cases} -a_1 - L(a_1, \hat{a}_2) & \text{if } a_1 < \hat{a}_1 \\ -a_1 & \text{if } a_1 \geq \hat{a}_1. \end{cases} \tag{91.1}$$

For $a_1 = \hat{a}_1$, this payoff is $-\hat{a}_1$. If she takes more care than $\hat{a}_1$, she is worse off, because care is costly and, beyond $\hat{a}_1$, does not reduce her liability for compensation. If she takes less care, then, given the victim's level of care, she has to pay compensation, and we need to compare the money saved by taking less care with the size of the compensation. The argument is a little tricky. First, by definition,

$$(\hat{a}_1, \hat{a}_2) \text{ maximizes } -a_1 - a_2 - L(a_1, a_2).$$

Hence

$$\hat{a}_1 \text{ maximizes } -a_1 - \hat{a}_2 - L(a_1, \hat{a}_2)$$

(given $\hat{a}_2$). Because $\hat{a}_2$ is a constant, it follows that

$$\hat{a}_1 \text{ maximizes } -a_1 - L(a_1, \hat{a}_2).$$

But from (91.1) we see that $-a_1 - L(a_1, \hat{a}_2)$ is the injurer's payoff $u_1(a_1, \hat{a}_2)$ when her action is $a_1 < \hat{a}_1$ and the victim's action is $\hat{a}_2$. We conclude that the injurer's payoff takes a form like that in the left panel of Figure 92.1. In particular, $\hat{a}_1$ maximizes $u_1(a_1, \hat{a}_2)$, so that $\hat{a}_1$ is a best response to $\hat{a}_2$.

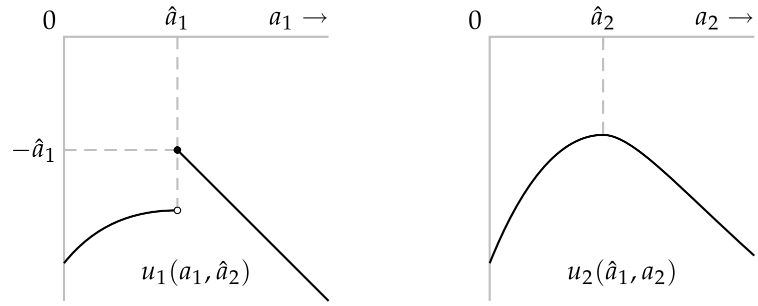

Figure 92.1 Left panel: the injurer's payoff as a function of her level of care a_1 when the victim's level of care is $a_2 = \hat{a}_2$ (see (91.1)). Right panel: the victim's payoff as a function of her level of care a_2 when the injurer's level of care is $a_1 = \hat{a}_1$ (see (92.1)).

Victim's action Given that the injurer's action is $\hat{a}_1$, the victim never receives compensation. Thus her payoff is

$$u_2(\hat{a}_1, a_2) = -a_2 - L(\hat{a}_1, a_2). \tag{92.1}$$

We can argue as we did for the injurer. By definition, $(\hat{a}_1, \hat{a}_2)$ maximizes $-a_1 - a_2 - L(a_1, a_2)$, so

$$\hat{a}_2 \text{ maximizes } -\hat{a}_1 - a_2 - L(\hat{a}_1, a_2)$$

(given $\hat{a}_1$). Because $\hat{a}_1$ is a constant, it follows that

$$\hat{a}_2 \text{ maximizes } -a_2 - L(\hat{a}_1, a_2), \tag{92.2}$$

which is the victim's payoff (see (92.1) and the right panel of Figure 92.1). That is, $\hat{a}_2$ maximizes $u_2(\hat{a}_1, a_2)$, so that $\hat{a}_2$ is a best response to $\hat{a}_1$.

We conclude that $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium of the game induced by the legal rule of negligence with contributory negligence when the standards of care are $\hat{a}_1$ for the injurer and $\hat{a}_2$ for the victim.

To show that $(\hat{a}_1, \hat{a}_2)$ is the *only* Nash equilibrium of the game, first consider the injurer's best response function. Her payoff function is

$$u_1(a_1, a_2) = \begin{cases} -a_1 - L(a_1, a_2) & \text{if } a_1 < \hat{a}_1 \text{ and } a_2 \geq \hat{a}_2 \\ -a_1 & \text{if } a_1 \geq \hat{a}_1 \text{ or } a_2 < \hat{a}_2. \end{cases}$$

We can split the analysis into three cases, according to the victim's level of care.

$a_2 < \hat{a}_2$: In this case the injurer does not have to pay any compensation, regardless of her level of care; her payoff is $-a_1$, so that her best response is $a_1 = 0$.

$a_2 = \hat{a}_2$: In this case the injurer's best response is $\hat{a}_1$, as argued when showing that $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium.

$a_2 > \hat{a}_2$: In this case the injurer's best response is at most $\hat{a}_1$, because her payoff for larger values of a_1 is equal to $-a_1$, a decreasing function of a_1.

We conclude that the injurer's best response function takes a form like that shown in the left panel of Figure 93.1.

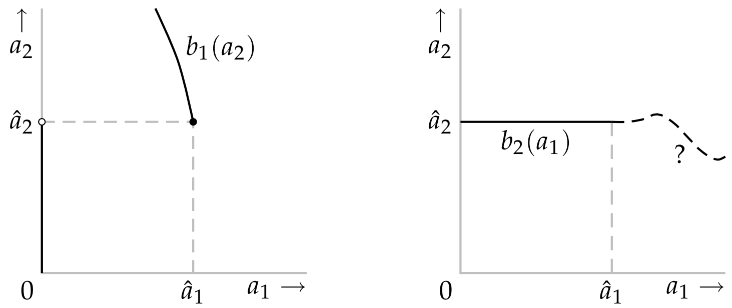

Figure 93.1 The players' best response functions under the rule of negligence with contributory negligence when $(X_1, X_2) = (\hat{a}_1, \hat{a}_2)$. Left panel: the injurer's best response function b_1. Right panel: the victim's best response function b_2. (The position of the victim's best response function for $a_1 > \hat{a}_1$ is not significant, and is not determined in the text.)

Now, given that the injurer's best response to any value of a_2 is never greater than $\hat{a}_1$, in any equilibrium we have $a_1 \leq \hat{a}_1$: any point (a_1, a_2) at which the victim's best response function crosses the injurer's best response function must have $a_1 \leq \hat{a}_1$. (Draw a few possible best response functions for the victim in the left panel of Figure 93.1.) We know that the victim's best response to $\hat{a}_1$ is $\hat{a}_2$ (because $(\hat{a}_1, \hat{a}_2)$ is a Nash equilibrium), so we need to worry only about the victim's best responses to values of a_1 with $a_1 < \hat{a}_1$ (i.e. for cases in which the injurer takes insufficient care).

Let $a_1 < \hat{a}_1$. Then if the victim takes insufficient care she bears the loss; otherwise she is compensated for the loss, and hence bears only the cost a_2 of her taking care. Thus the victim's payoff is

$$u_2(a_1, a_2) = \begin{cases} -a_2 - L(a_1, a_2) & \text{if } a_2 < \hat{a}_2 \\ -a_2 & \text{if } a_2 \geq \hat{a}_2. \end{cases} \tag{94.1}$$

Now, by (92.2) the level of care $\hat{a}_2$ maximizes $-a_2 - L(\hat{a}_1, a_2)$, so that

$$-a_2 - L(\hat{a}_1, a_2) \leq -\hat{a}_2 - L(\hat{a}_1, \hat{a}_2) \text{ for all } a_2.$$

Further, the loss is nonnegative, so $-\hat{a}_2 - L(\hat{a}_1, \hat{a}_2) \leq -\hat{a}_2$. We conclude that

$$-a_2 - L(\hat{a}_1, a_2) \leq -\hat{a}_2 \text{ for all } a_2. \tag{94.2}$$

Finally, the loss increases as the injurer takes less care, so that given $a_1 < \hat{a}_1$ we have $L(a_1, a_2) > L(\hat{a}_1, a_2)$ for all a_2. Thus $-a_2 - L(a_1, a_2) < -a_2 - L(\hat{a}_1, a_2)$ for all a_2, and hence, using (94.2),

$$-a_2 - L(a_1, a_2) < -\hat{a}_2 \text{ for all } a_2.$$

From (94.1) it follows that the victim's best response to any $a_1 < \hat{a}_1$ is $\hat{a}_2$, as shown in the right panel of Figure 93.1.

Combining the two best response functions we see that $(\hat{a}_1, \hat{a}_2)$, the pair of levels of care that maximizes the sum of the players' payoffs, is the unique Nash equilibrium of the game. That is, the rule of negligence with contributory negligence for standards of care equal to $\hat{a}_1$ and $\hat{a}_2$ induces the players to choose these levels of care. If legislators can determine the values of $\hat{a}_1$ and $\hat{a}_2$ then by writing these levels into law they will induce a game that has as its unique Nash equilibrium the socially optimal actions.

Other standards also induce a pair of levels of care equal to $(\hat{a}_1, \hat{a}_2)$, as you are asked to show in the following exercise.

?? EXERCISE 94.3 (Alternative standards of care under negligence with contributory negligence) Show that $(\hat{a}_1, \hat{a}_2)$ is the unique Nash equilibrium for the rule of negligence with contributory negligence for any value of (X_1, X_2) for which *either* $X_1 = \hat{a}_1$ and $X_2 \leq \hat{a}_2$ (including the pure negligence case of $X_2 = 0$), *or* $X_1 \geq M$ and $X_2 = \hat{a}_2$ for sufficiently large M. (Use the lines of argument in the text.)

? EXERCISE 94.4 (Equilibrium under strict liability) Study the Nash equilibrium (equilibria?) of the game studied in the text under the rule of strict liability, in which X_1 is infinite and $X_2 = 0$ (i.e. the injurer is liable for the loss no matter how careful she is and how careless the victim is). How are the equilibrium actions related to $\hat{a}_1$ and $\hat{a}_2$?

Notes

The model in Section 3.1 was developed by Cournot (1838). The model in Section 3.2 is widely credited to Bertrand (1883). The box on page 67 is based on

Leonard (1994) and Magnan de Bornier (1992). The models are discussed in more detail by Shapiro (1989).

The model in Section 3.3 is due to Hotelling (1929) (though the focus of his paper is a model in which the players are firms that choose not only locations, but also prices). Downs (1957, especially Ch. 8) popularized Hotelling's model, using it to gain insights about electoral competition. Shepsle (1991) and Osborne (1995) survey work in the field.

The *War of Attrition* studied in Section 3.4 is due to Maynard Smith (1974); it is a variant of the *Dollar Auction* presented by Shubik (1971) (see Example 173.2).

Vickrey (1961) initiated the formal modeling of auctions, as studied in Section 3.5. The literature is surveyed by Wilson (1992). The box on page 79 draws on Herodotus' *Histories* (Book 1, paragraph 196; see, for example, Herodotus 1998, 86), Langdon (1994), Cassady (1967, Ch. 3), Shubik (1983), Andreau (1999, 38–39), the website `www.eBay.com`, Ockenfels and Roth (2000), and personal correspondence with Robin G. Osborne (on ancient Greece and Rome) and John H. Munro (on medieval Europe).

The model of accident law discussed in Section 3.6 originated with Brown (1973) and Diamond (1974); the result about negligence with contributory negligence is due to Brown (1973, 340–341). The literature is surveyed by Benoît and Kornhauser (2002).

Novshek and Sonnenschein (1978) study, in a general setting, the issue addressed in Exercise 60.1. A brief summary of the early work on common property is given in the Notes to Chapter 2. The idea of the tie-breaking rule being determined by the equilibrium, used in Exercises 66.2 and 67.1, is due to Simon and Zame (1990). The result in Exercise 73.1 is due to Wittman (1977). Exercise 73.2 is based on Osborne and Slivinski (1996). The notion of a Condorcet winner defined in Exercise 74.1 is associated with Marie-Jean-Antoine-Nicolas de Caritat, marquis de Condorcet (1743–1794), an early student of voting procedures. The game in Exercise 78.1 is a variant of a game studied by Blackwell and Girschick (1954, Example 5 in Ch. 2). It is an example of a *noisy duel* (which models the situation of duelists, each of whom chooses when to fire a single bullet, which her opponent hears, as she gradually approaches her rival). Duels were first modeled as games in the late 1940s by members of the RAND Corporation in the USA; see Karlin (1959b, Ch. 5). Exercise 88.3 is based on Boylan (1997). The situation considered in Exercise 88.1, in which people decide when to join a queue, is studied by Holt and Sherman (1982). Exercise 88.2 is based on MacKie-Mason and Varian (1995).