6 Extensive Games with Perfect Information: Illustrations

Prerequisite: Chapter 5

THE first three sections of this chapter illustrate the notion of subgame perfect equilibrium in games in which the longest history has length two or three. The last section studies a game with an arbitrary finite horizon. Some games with infinite horizons are studied in Chapters 14, 15, and 16.

6.1 The ultimatum game, the holdup game, and agenda control

6.1.1 The ultimatum game

Bargaining over the division of a pie may naturally be modeled as an extensive game. Here I analyze a very simple game that is the basis of a richer model studied in Chapter 16. The game is so simple, in fact, that you may not initially think of it as a model of "bargaining".

Two people use the following procedure to split \$$c$. Person 1 offers person 2 an amount of money up to \$$c$. If 2 accepts this offer, then 1 receives the remainder of the \$$c$. If 2 rejects the offer, then *neither* person receives any payoff. Each person cares *only* about the amount of money she receives, and (naturally!) prefers to receive as much as possible.

Assume that the amount person 1 offers can be any number, not necessarily an integral number of cents. Then the following extensive game, known as the **ultimatum game**, models the procedure.

Players The two people.

Terminal histories The set of sequences (x, Z), where x is a number with $0 \leq x \leq c$ (the amount of money that person 1 offers to person 2) and Z is either Y ("yes, I accept") or N ("no, I reject").

Player function $P(\varnothing) = 1$ and $P(x) = 2$ for all x.

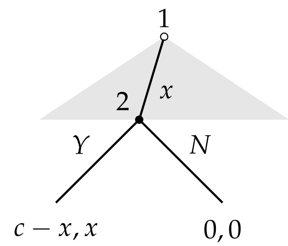

Figure 182.1 An illustration of the ultimatum game. The gray triangle represents the continuum of possible offers of player 1; the black lines indicate the terminal histories that start with the offer x.

Preferences Each person's preferences are represented by payoffs equal to the amounts of money she receives. For the terminal history (x, Y) person 1 receives $c - x$ and person 2 receives x; for the terminal history (x, N) each person receives 0.

This game is illustrated in Figure 182.1, in which the continuum of offers of player 1 is represented by the gray triangle, and the black lines indicate the terminal histories that start with the offer x. The game has a finite horizon, so we can use backward induction to find its subgame perfect equilibria. First consider the subgames of length 1, in which person 2 either accepts or rejects an offer of person 1. For every possible offer of person 1, there is such a subgame. In the subgame that follows an offer x of person 1 for which $x > 0$, person 2's optimal action is to accept (if she rejects, she gets nothing). In the subgame that follows the offer $x = 0$, person 2 is indifferent between accepting and rejecting. Thus in a subgame perfect equilibrium person 2's strategy either accepts all offers (including 0), or accepts all offers $x > 0$ and rejects the offer $x = 0$.

Now consider the whole game. For each possible subgame perfect equilibrium strategy of person 2, we need to find the optimal strategy of person 1.

- If person 2 accepts all offers (including 0), then person 1's optimal offer is 0 (which yields her the payoff $\$c$).
- If person 2 accepts all offers except zero, then *no* offer of person 1 is optimal! No offer $x > 0$ is optimal, because the offer $x/2$ (for example) is better, given that person 2 accept both offers. And an offer of 0 is not optimal because person 2 rejects it, leading to a payoff of 0 for person 1, who is thus better off offering any positive amount less than $\$c$.

We conclude that the only subgame perfect equilibrium of the game is the strategy pair in which person 1 offers 0 and person 2 accepts all offers. In this equilibrium, person 1's payoff is $\$c$ and person 2's payoff is zero.

This one-sided outcome is a consequence of the one-sided structure of the game. If we allow person 2 to make a counteroffer after rejecting person 1's opening offer (and possibly allow further responses by both players), so that the model corresponds more closely to a "bargaining" situation, then under some circumstances the outcome is less one-sided. (An extension of this type is explored in Chapter 16.)

? EXERCISE 183.1 (Nash equilibria of the ultimatum game) Find the values of x for which there is a Nash equilibrium of the ultimatum game in which person 1 offers x.

? EXERCISE 183.2 (Subgame perfect equilibria of the ultimatum game with indivisible units) Find the subgame perfect equilibria of the variant of the ultimatum game in which the amount of money is available only in multiples of a cent.

? EXERCISE 183.3 (Dictator game and impunity game) The "dictator game" differs from the ultimatum game only in that person 2 does not have the option to reject person 1's offer (and thus has no strategic role in the game). The "impunity game" differs from the ultimatum game only in that person 1's payoff when person 2 rejects any offer x is $c - x$, rather than 0. (The game is named for the fact that person 2 is unable to "punish" person 1 for making a low offer.) Find the subgame perfect equilibria of each game.

?? EXERCISE 183.4 (Variants of ultimatum game and impunity game with equity-conscious players) Consider variants of the ultimatum game and impunity game in which each person cares not only about the amount of money she receives, but also about the equity of the allocation. Specifically, suppose that person i's preferences are represented by the payoff function given by $u_i(x_1, x_2) = x_i - \beta_i |x_1 - x_2|$, where x_i is the amount of money person i receives, $\beta_i > 0$, and, for any number z, $|z|$ denotes the absolute value of z (i.e. $|z| = z$ if $z > 0$ and $|z| = -z$ if $z < 0$). Assume $c = 1$. Find the set of subgame perfect equilibria of each game and compare them. Are there any values of β_1 and β_2 for which an offer is rejected in equilibrium? (An interesting further variant of the ultimatum game in which person 1 is uncertain about the value of β_2 is considered in Exercise 227.1.)

EXPERIMENTS ON THE ULTIMATUM GAME

The sharp prediction of the notion of subgame perfect equilibrium in the ultimatum game lends itself to experimental testing. The first test was conducted in the late 1970s among graduate students of economics in a class at the University of Cologne (in what was then West Germany). The amount c available varied among the games played; it ranged from 4 DM to 10 DM (around U.S.$2 to U.S.$5 at the time). A group of 42 students was split into two groups and seated on different sides of a room. Each member of one subgroup played the role of player 1 in an ultimatum game. She wrote down on a form the amount (up to c) that she demanded. Her form was then given to a randomly determined member of the other group, who, playing the role of player 2, either accepted what remained of the amount c or rejected it (in which case neither player received any payoff). Each player had 10 minutes to make her decision. The entire experiment was repeated a week later. (Güth, Schmittberger, and Schwarze 1982.)

In the first experiment the average demand by people playing the role of player 1 was $0.65c$, and in the second experiment it was $0.69c$, much less than the amount c or $c - 0.01$ predicted by the notion of subgame perfect equilibrium (0.01 DM was the smallest monetary unit; see Exercise 183.2). Almost 20% of offers were rejected over the two experiments, including one of 3 DM (out of a pie of 7 DM) and five of around 1 DM (out of pies of between 4 DM and 6 DM). Many other experiments, including one in which the amount of money to be divided was much larger (Hoffman, McCabe, and Smith 1996), have produced similar results. In brief, the results do not accord well with the predictions of subgame perfect equilibrium.

Or do they? Each player in the ultimatum game cares only about the amount of money she receives. But an experimental subject may care also about the amount of money her opponent receives. Further, a variant of the ultimatum game in which the players are equity conscious has subgame perfect equilibria in which offers are significant (as you will have discovered if you did Exercise 183.4).

However, if people are equity conscious in the strategic environment of the ultimatum game, they are presumably equity conscious also in related environments; an explanation of the experimental results in the ultimatum game based on the players' preferences' exhibiting equity conscience is not convincing if it applies only to that environment. Several related games have been studied, among them the dictator game and the impunity game (Exercise 183.3). In the subgame perfect equilibria of these games, player 1 offers 0; in a variant in which the players are equity conscious, player 1's offers are no higher than they are in the analogous variant of the ultimatum game, and, for moderate degrees of equity conscience, are lower (see Exercise 183.4). These features of the equilibria are broadly consistent with the experimental evidence on dictator, impunity, and ultimatum games (see, for example, Forsythe, Horowitz, Savin, and Sefton 1994, Bolton and Zwick 1995, and Güth and Huck 1997).

One feature of the experimental results is inconsistent with subgame perfect equilibrium even when players are equity conscious (at least given the form of the payoff functions in Exercise 183.4): positive offers are sometimes rejected. The equilibrium strategy of an equity-conscious player 2 in the ultimatum game rejects inequitable offers, but, knowing this, player 1 does not, in equilibrium, make such an offer. To generate rejections in equilibrium we need to further modify the model by assuming that people differ in their degree of equity conscience, and that player 1 does not know the degree of equity conscience of player 2 (see Exercise 227.1).

An alternative explanation of the experimental results focuses on player 2's behavior. The evidence is consistent with player 1's significant offers in the ultimatum game being driven by a fear that player 2 will reject small offers—a fear that is rational, because small offers are often rejected. Why does player 2 behave in this way? One argument is that in our daily lives, we use "rules of thumb" that work well in the situations in which we are typically involved; we do not calculate our rational actions in each situation. Further, we are not typically involved in one-shot situations with the structure of the ultimatum game. Instead, we usually engage in repeated interactions, where it is advantageous to "punish" a player who makes a paltry offer, and to build a reputation for not accepting such offers. Experimental subjects may apply such rules of thumb rather than carefully thinking through the logic of the game, and thus reject low offers in an ultimatum game but accept them in an impunity game, where rejection does not affect the proposer. The experimental evidence so far collected is broadly consistent with both this explanation and the explanation based on the nature of players' preferences.

? EXERCISE 185.1 (Bargaining over two indivisible objects) Consider a variant of the ultimatum game, with indivisible units. Two people use the following procedure to allocate two desirable identical indivisible objects. One person proposes an allocation (both objects go to person 1, both go to person 2, one goes to each person), which the other person then either accepts or rejects. In the event of rejection, neither person receives either object. Each person cares only about the number of objects she obtains. Construct an extensive game that models this situation and find its subgame perfect equilibria. Does the game have any Nash equilibrium that is not a subgame perfect equilibrium? Is there any outcome that is generated by a Nash equilibrium but not by any subgame perfect equilibrium?

?? EXERCISE 185.2 (Dividing a cake fairly) Two players use the following procedure to divide a cake. Player 1 divides the cake into two pieces, and then player 2 chooses one of the pieces; player 1 obtains the remaining piece. The cake is continuously divisible (no lumps!), and each player likes all parts of it.

a. Suppose that the cake is perfectly homogeneous, so that each player cares only about the size of the piece of cake she obtains. How is the cake divided in a subgame perfect equilibrium?

b. Suppose that the cake is not homogeneous: the players evaluate different parts of it differently. Represent the cake by the set C, so that a piece of the cake is a subset P of C. Assume that if P is a subset of P' not equal to P' (smaller than P'), then each player prefers P' to P. Assume also that the players' preferences are continuous: if a player prefers P to its complement (the remainder of the cake), then there is a subset P' of P not equal to P such that the player prefers P' to its complement. Let (P_1, P_2) (where P_1 and P_2 together constitute the whole cake C) be the division chosen by player 1 in a subgame perfect equilibrium of the divide-and-choose game, where P_2 is the piece chosen by player 2. Show that player 2 is indifferent between P_1 and P_2, and player 1 likes P_1 at least as much as P_2. Give an example in which player 1 prefers P_1 to P_2.

6.1.2 The holdup game

Before engaging in an ultimatum game in which she may accept or reject an offer of person 1, person 2 takes an action that affects the size c of the pie to be divided. She may exert little effort, resulting in a small pie, of size c_L, or great effort, resulting in a large pie, of size c_H. She dislikes exerting effort. Specifically, assume that her payoff is $x - E$ if her share of the pie is x, where $E = L$ if she exerts little effort and $E = H > L$ if she exerts great effort. The extensive game that models this situation is known as the **holdup game**.

? EXERCISE 186.1 (Holdup game) Formulate the holdup game precisely. (Write down the set of players, the set of terminal histories, the player function, and the players' preferences.)

What is the subgame perfect equilibrium of the holdup game? Each subgame that follows person 2's choice of effort is an ultimatum game, and thus has a unique subgame perfect equilibrium, in which person 1 offers 0 and person 2 accepts all offers. Now consider person 2's choice of effort at the start of the game. If she chooses L, then her payoff, given the outcome in the following subgame, is $-L$, whereas if she chooses H, then her payoff is $-H$. Consequently she chooses L. Thus the game has a unique subgame perfect equilibrium, in which person 2 exerts little effort and person 1 obtains all of the resulting small pie.

This equilibrium does not depend on the values of c_L, c_H, L, and H (given that $H > L$). In particular, even if c_H is much larger than c_L, but H is only slightly larger than L, person 2 exerts little effort in the equilibrium, although both players could be much better off if person 2 were to exert great effort (which, in this case, is not very great) and person 2 were to obtain some of the extra pie. No such superior outcome is sustainable in an equilibrium because person 2, having exerted great effort, may be "held up" for the entire pie by person 1.

This result does not depend sensitively on the extreme subgame perfect equilibrium outcome of the ultimatum game. A similar result emerges when the bargaining following person 2's choice of effort generates a more equal division of the pie. By exerting great effort, player 2 increases the size of the pie. The point is that if the negotiation results in some (not necessarily all) of this extra pie going to player 1, then for some values of player 2's cost of exerting great effort less than the value of the extra pie, player 2 prefers to exert little effort. In these circumstances, player 2's exerting great effort generates outcomes in which both players are better off than they are when player 2 exerts little effort, but because the bargaining puts some of the extra pie in the hands of player 1, player 2's incentive is to exert little effort.

6.1.3 Agenda control

In some legislatures, proposals for modifications of the law are formulated by committees. Under a "closed rule", the legislature may either accept or reject a proposed modification, but may not propose an alternative; in the event of rejection, the existing law is unchanged. That is, the committee controls the "agenda". (In Section 10.9 I consider a reason why a legislature might cede such power to a committee.)

Model an outcome as a number y. Assume that the legislature and committee have favorite outcomes that may differ, and that the preferences of each body are represented by a single-peaked payoff function symmetric about its favorite outcome, like the voters' preferences in Hotelling's model of electoral competition (see Figure 71.1). Assign numbers to outcomes so that the legislature's favorite outcome is 0; denote the committee's favorite outcome by $y_c > 0$. Then the following variant of the ultimatum game models the procedure. The players are the committee and the legislature. The committee proposes an outcome y, which the legislature either accepts or rejects. In the event of rejection the outcome is y_0, the "status quo". Note that the main respect in which this game differs from the ultimatum game is that the players' preferences are diametrically opposed only with regard to outcomes between 0 and y_c; if $y' < y'' < 0$ or $y_c < y'' < y'$, then both players prefer y'' to y'.

? EXERCISE 187.1 (Agenda control) Find the subgame perfect equilibrium of this game as a function of the status quo outcome y_0. Show, in particular, that for a range of values of y_0, an increase in the value of y_0 leads to a *decrease* in the value of the equilibrium outcome.

6.2 Stackelberg's model of duopoly

6.2.1 General model

In the models of oligopoly in Sections 3.1 and 3.2, each firm chooses its action not knowing the other firms' actions. How do the conclusions change when the firms move sequentially? Is a firm better off moving before or after the other firms?

In this section I consider a market in which there are two firms, both producing the same good. Firm i's cost of producing q_i units of the good is $C_i(q_i)$; the price at which output is sold when the total output is Q is $P_d(Q)$. (In Section 3.1 I denote this function P; here I add a d subscript to avoid a conflict with the player function of the extensive game.) Each firm's strategic variable is output, as in Cournot's model (Section 3.1), but the firms make their decisions sequentially, rather than simultaneously: one firm chooses its output, then the other firm does so, knowing the output chosen by the first firm.

We can model this situation by the following extensive game, known as **Stackelberg's duopoly game** (after an early analyst of duopoly with asynchronous actions).

Players The two firms.

Terminal histories The set of all sequences (q_1, q_2) of outputs for the firms (where each q_i, the output of firm i, is a nonnegative number).

Player function $P(\varnothing) = 1$ and $P(q_1) = 2$ for all q_1.

Preferences The payoff of firm i to the terminal history (q_1, q_2) is its profit $q_i P_d(q_1 + q_2) - C_i(q_i)$, for $i = 1, 2$.

Firm 1 moves at the start of the game. Thus a strategy of firm 1 is simply an output. Firm 2 moves after every history in which firm 1 chooses an output. Thus a strategy of firm 2 is a *function* that associates an output for firm 2 with each possible output of firm 1.

The game has a finite horizon, so we may use backward induction to find its subgame perfect equilibria.

- First, for any output of firm 1, we find the outputs of firm 2 that maximize its profit. Suppose that for each output q_1 of firm 1 there is one such output of firm 2; denote it $b_2(q_1)$. Then in any subgame perfect equilibrium, firm 2's strategy is b_2.
- Next, we find the outputs of firm 1 that maximize its profit, *given the strategy of firm 2*. When firm 1 chooses the output q_1, firm 2 chooses the output $b_2(q_1)$, resulting in a total output of $q_1 + b_2(q_1)$, and hence a price of $P_d(q_1 + b_2(q_1))$. Thus firm 1's output in a subgame perfect equilibrium is a value of q_1 that maximizes
$$q_1 P_d(q_1 + b_2(q_1)) - C_1(q_1). \tag{188.1}$$
Suppose that there is one such value of q_1; denote it q_1^*.

We conclude that if firm 2 has a unique best response $b_2(q_1)$ to each output q_1 of firm 1, and firm 1 has a unique best action q_1^*, given firm 2's best responses, then the subgame perfect equilibrium of the game is (q_1^*, b_2): firm 1's equilibrium strategy is q_1^* and firm 2's equilibrium strategy is the function b_2. The output chosen by firm 2, given firm 1's equilibrium strategy, is $b_2(q_1^*)$; denote this output q_2^*.

When firm 1 chooses any output q_1, the outcome, given that firm 2 uses its equilibrium strategy, is the pair of outputs $(q_1, b_2(q_1))$. That is, as firm 1 varies its output, the outcome varies along firm 2's best response function b_2. Thus we can characterize the subgame perfect equilibrium outcome (q_1^*, q_2^*) as the point on firm 2's best response function that maximizes firm 1's profit.

6.2.2 Example: constant unit cost and linear inverse demand

Suppose that $C_i(q_i) = cq_i$ for $i = 1, 2$, and
$$P_d(Q) = \begin{cases} \alpha - Q & \text{if } Q \le \alpha \\ 0 & \text{if } Q > \alpha, \end{cases} \tag{188.2}$$
where $c > 0$ and $c < \alpha$ (as in the example of Cournot's duopoly game in Section 3.1.3). We found that under these assumptions firm 2 has a unique best response to each output q_1 of firm 1, given by
$$b_2(q_1) = \begin{cases} \frac{1}{2}(\alpha - c - q_1) & \text{if } q_1 \le \alpha - c \\ 0 & \text{if } q_1 > \alpha - c. \end{cases}$$

Thus in a subgame perfect equilibrium of Stackelberg's game firm 2's strategy is this function b_2 and firm 1's strategy is the output q_1 that maximizes

$$q_1(\alpha - c - (q_1 + \tfrac{1}{2}(\alpha - c - q_1))) = \tfrac{1}{2}q_1(\alpha - c - q_1)$$

(refer to (188.1)). This function is a quadratic in q_1 that is zero when $q_1 = 0$ and when $q_1 = \alpha - c$. Thus its maximizer is $q_1 = \frac{1}{2}(\alpha - c)$.

We conclude that the game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output $\frac{1}{2}(\alpha - c)$ and firm 2's strategy is b_2. The outcome of the equilibrium is that firm 1 produces the output $q_1^* = \frac{1}{2}(\alpha - c)$ and firm 2 produces the output $q_2^* = b_2(q_1^*) = b_2(\frac{1}{2}(\alpha - c)) = \frac{1}{2}(\alpha - c - \frac{1}{2}(\alpha - c)) = \frac{1}{4}(\alpha - c)$. Firm 1's profit is $q_1^*(P_d(q_1^* + q_2^*) - c) = \frac{1}{8}(\alpha - c)^2$, and firm 2's profit is $q_2^*(P_d(q_1^* + q_2^*) - c) = \frac{1}{16}(\alpha - c)^2$. By contrast, in the unique Nash equilibrium of Cournot's (simultaneous-move) game under the same assumptions, each firm produces $\frac{1}{3}(\alpha - c)$ units of output and obtains the profit $\frac{1}{9}(\alpha - c)^2$. Thus under our assumptions firm 1 produces more output and obtains more profit in the subgame perfect equilibrium of the sequential game in which it moves first than it does in the Nash equilibrium of Cournot's game, and firm 2 produces less output and obtains less profit.

? EXERCISE 189.1 (Stackelberg's duopoly game with quadratic costs) Find the subgame perfect equilibrium of Stackelberg's duopoly game when $C_i(q_i) = q_i^2$ for $i = 1, 2$, and $P_d(Q) = \alpha - Q$ for all $Q \leq \alpha$ (with $P_d(Q) = 0$ for $Q > \alpha$). Compare the equilibrium outcome with the Nash equilibrium of Cournot's game under the same assumptions (Exercise 59.1).

6.2.3 Properties of subgame perfect equilibrium

First-mover's equilibrium profit In the example just studied, the first-mover is better off in the subgame perfect equilibrium of Stackelberg's game than it is in the Nash equilibrium of Cournot's game. A weak version of this result holds under very general conditions: for any cost and inverse demand functions for which firm 2 has a unique best response to each output of firm 1, firm 1 is at least as well off in any subgame perfect equilibrium of Stackelberg's game as it is in any Nash equilibrium of Cournot's game. This result follows from the general result in Exercise 177.3*a*. The argument is simple. One of firm 1's options in Stackelberg's game is to choose its output in some Nash equilibrium of Cournot's game. If it chooses such an output, then firm 2's best action is to choose its output in the same Nash equilibrium, given the assumption that it has a unique best response to each output of firm 1. Thus by choosing such an output, firm 1 obtains its profit at a Nash equilibrium of Cournot's game; by choosing a different output it may possibly obtain a higher payoff.

Equilibrium outputs In the example in the previous section (6.2.2), firm 1 produces more output in the subgame perfect equilibrium of Stackelberg's game than it does

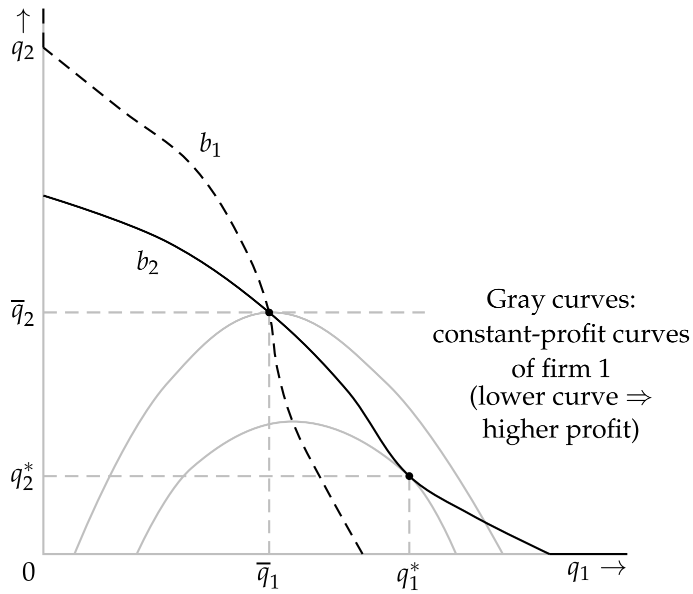

Figure 190.1 The subgame perfect equilibrium outcome (q_1^*, q_2^*) of Stackelberg's game and the Nash equilibrium $(\overline{q}_1, \overline{q}_2)$ of Cournot's game. Along each gray curve, firm 1's profit is constant; the lower curve corresponds to higher profit than does the upper curve. Each curve has a slope of zero where it crosses firm 1's best response function b_1.

in the Nash equilibrium of Cournot's game, and firm 2 produces less. A weak form of this result holds whenever firm 2's best response function is decreasing where it is positive (i.e. a higher output for firm 1 implies a lower optimal output for firm 2).

The argument is illustrated in Figure 190.1. The firms' best response functions are the curves labeled b_1 (dashed) and b_2. The Nash equilibrium of Cournot's game is the intersection $(\overline{q}_1, \overline{q}_2)$ of these curves. Along each gray curve, firm 1's profit is constant; the *lower* curve corresponds to a *higher* profit. (For any given value of firm 1's output, a reduction in the output of firm 2 increases the price and thus increases firm 1's profit.) Each constant-profit curve of firm 1 is horizontal where it crosses firm 1's best response function, because the best response is precisely the output that maximizes firm 1's profit, given firm 2's output. (Cf. Figure 61.1.) Thus the subgame perfect equilibrium outcome—the point on firm 2's best response function that yields the highest profit for firm 1—is the point (q_1^*, q_2^*) in Figure 190.1. In particular, given that the best response function of firm 2 is downward sloping, firm 1 produces at least as much, and firm 2 produces at most as much, in the subgame perfect equilibrium of Stackelberg's game as in the Nash equilibrium of Cournot's game.

For some cost and demand functions, firm 2's output in a subgame perfect equilibrium of Stackelberg's game is zero. An example is shown in Figure 191.1. The discontinuity in firm 2's best response function at q_1^* in this example may arise because firm 2 incurs a "fixed" cost—a cost independent of its output—when it produces a positive output (see Exercise 59.2). When firm 1's output is q_1^*, firm 2's

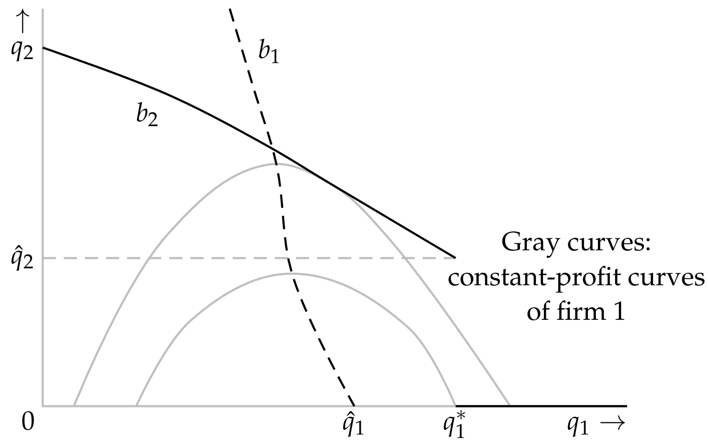

Figure 191.1 The subgame perfect equilibrium output q_1^* of firm 1 in Stackelberg's sequential game when firm 2 incurs a fixed cost. Along each gray curve, firm 1's profit is constant; the lower curve corresponds to higher profit than does the upper curve.

maximal profit is zero, which it obtains both when it produces no output (and does not pay the fixed cost) and when it produces the output $\hat{q}_2$. When firm 1 produces less than q_1^*, firm 2's maximal profit is positive, and firm 2 optimally produces a positive output; when firm 1 produces more than q_1^*, firm 2 optimally produces no output. Given this form of firm 2's best response function and the form of firm 1's constant-profit curves in Figure 190.1, the point on firm 2's best response function that yields firm 1 the highest profit is $(q_1^*, 0)$.

I claim that this example has a unique subgame perfect equilibrium, in which firm 1 produces q_1^* and firm 2's strategy coincides with its best response function except at q_1^*, where the strategy specifies the output 0. The output firm 2's equilibrium strategy specifies after each history must be a best response to firm 1's output, so the only question regarding firm 2's strategy is whether it specifies an output of 0 or $\hat{q}_2$ when firm 1's output is q_1^*. The argument that there is no subgame perfect equilibrium in which firm 2's strategy specifies the output $\hat{q}_2$ is similar to the argument that there is no subgame perfect equilibrium in the ultimatum game in which person 2 rejects the offer 0. If firm 2 produces the output $\hat{q}_2$ in response to firm 1's output q_1^*, then firm 1 has no optimal output: it would like to produce a little more than q_1^*, inducing firm 2 to produce zero, but is better off the closer its output is to q_1^*. Because there is no smallest output greater than q_1^*, no output is *optimal* for firm 1 in this case. Thus the game has no subgame perfect equilibrium in which firm 2's strategy specifies the output $\hat{q}_2$ in response to firm 1's output q_1^*.

Note that if firm 2 were entirely absent from the market, firm 1 would produce $\hat{q}_1$, less than q_1^*. Thus firm 2's presence affects the outcome, even though it produces no output.

? EXERCISE 191.1 (Stackelberg's duopoly game with fixed costs) Suppose that the inverse demand function is given by (188.2) 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 $c < \alpha$, as in Exercise 59.2. Show that if $c = 0$, $\alpha = 12$, and $f = 4$, Stackelberg's game has a unique subgame perfect equilibrium, in which firm 1's output is 8 and firm 2's output is zero. (Use your results from Exercise 59.2).

The value of commitment Firm 1's output in a subgame perfect equilibrium of Stackelberg's game is *not* in general a best response to firm 2's output: if firm 1 could adjust its output after firm 2 has chosen its output, then it would do so! (In the case shown in Figure 190.1, it would reduce its output.) However, if firm 1 had this opportunity, and firm 2 knew that it had the opportunity, then firm 2 would choose a different output. Indeed, if we simply add a third stage to the game, in which firm 1 chooses an output, then the first stage is irrelevant, and *firm 2* is effectively the first-mover; in the subgame perfect equilibrium firm 1 is worse off than it is in the Nash equilibrium of the simultaneous-move game. (In the example in Section 6.2.2, the unique subgame perfect equilibrium has firm 2 choose the output $(\alpha - c)/2$ and firm 1 choose the output $(\alpha - c)/4$.) In summary, even though firm 1 can increase its profit by changing its output after firm 2 has chosen its output, in the game in which it has this opportunity it is worse off than it is in the game in which it must choose its output before firm 2 and cannot subsequently modify this output. That is, firm 1 prefers to be *committed* not to change its mind.

? EXERCISE 192.1 (Sequential variant of Bertrand's duopoly game) Consider the variant of Bertrand's duopoly game (Section 3.2) in which first firm 1 chooses a price, then firm 2 chooses a price. Assume that each firm is restricted to choose a price that is an integral number of cents (as in Exercise 67.2), that each firm's unit cost is constant and equal to c (an integral number of cents), and that the monopoly profit is positive.

a. Specify an extensive game with perfect information that models this situation.

b. Give an example of a strategy of firm 1 and an example of a strategy of firm 2.

c. Find the subgame perfect equilibria of the game.

6.3 Buying votes

A legislature has k members, where k is an odd number. Two rival bills, X and Y, are being considered. The bill that attracts the votes of a majority of legislators will pass. Interest group X favors bill X, whereas interest group Y favors bill Y. Each group wishes to entice a majority of legislators to vote for its favorite bill. First interest group X gives an amount of money (possibly zero) to each legislator, then interest group Y does so. Each interest group wishes to spend as little as possible. Group X values the passing of bill X at $\$V_X > 0$ and the passing of bill Y

at zero, and group Y values the passing of bill Y at $\$V_Y > 0$ and the passing of bill X at zero. (For example, group X is indifferent between an outcome in which it spends V_X and bill X is passed and one in which it spends nothing and bill Y is passed.) Each legislator votes for the favored bill of the interest group that offers her the most money; a legislator to whom both groups offer the same amount of money votes for bill Y (an arbitrary assumption that simplifies the analysis without qualitatively changing the outcome). For example, if $k = 3$, the amounts offered to the legislators by group X are $x = (100, 50, 0)$, and the amounts offered by group Y are $y = (100, 0, 50)$, then legislators 1 and 3 vote for Y and legislator 2 votes for X, so that Y passes. (In some actual legislatures the inducements offered to legislators are more subtle than cash transfers.)

We can model this situation as the following extensive game.

Players The two interest groups, X and Y.

Terminal histories The set of all sequences (x, y), where x is a list of payments to legislators made by interest group X and y is a list of payments to legislators made by interest group Y. (That is, both x and y are lists of k nonnegative integers.)

Player function $P(\varnothing) = X$ and $P(x) = Y$ for all x.

Preferences The preferences of interest group X are represented by the payoff function

$$\begin{cases} V_X - (x_1 + \cdots + x_k) & \text{if bill } X \text{ passes} \\ -(x_1 + \cdots + x_k) & \text{if bill } Y \text{ passes,} \end{cases}$$

where bill Y passes after the terminal history (x, y) if and only if the number of components of y that are at least equal to the corresponding components of x is at least $\frac{1}{2}(k + 1)$ (a bare majority of the k legislators). The preferences of interest group Y are represented by the analogous function (where V_Y replaces V_X, y replaces x, and Y replaces X).

Before studying the subgame perfect equilibria of this game for arbitrary values of the parameters, consider two examples. First suppose that $k = 3$ and $V_X = V_Y = 300$. Under these assumptions, the most group X is willing to pay to get bill X passed is 300. For any payments it makes to the three legislators that sum to at most 300, two of the payments sum to at most 200, so that if group Y matches these payments it spends less than V_Y (= 300) and gets bill Y passed. Thus in any subgame perfect equilibrium group X makes no payments, group Y makes no payments, and (given the tie-breaking rule) bill Y is passed.

Now suppose that $k = 3$, $V_X = 300$, and $V_Y = 100$. In this case by paying each legislator more than 50, group X makes matching payments by group Y unprofitable: only by spending more than V_Y (= 100) can group Y cause bill Y to be passed. However, there is no subgame perfect equilibrium in which group X pays each legislator more than 50 because it can always pay a little less (as long as the payments still exceed 50) and still prevent group Y from profitably matching. In the only subgame perfect equilibrium group X pays each legislator exactly 50 and group Y makes no payments. Given group X's action, group Y is indifferent between matching X's payments (so that bill Y is passed) and making no payments. However, there is no subgame perfect equilibrium in which group Y matches group X's payments because if this were group Y's response, then group X could increase its payments a little, making matching payments by group Y unprofitable.

For arbitrary values of the parameters, the subgame perfect equilibrium outcome takes one of the forms in these two examples: either no payments are made and bill Y is passed, or group X makes payments that group Y does not wish to match, group Y makes no payments, and bill X is passed.

To find the subgame perfect equilibria in general, we may use backward induction. First consider group Y's best response to an arbitrary strategy x of group X. Let $\mu = \frac{1}{2}(k+1)$, a bare majority of k legislators, and denote by m_x the sum of the smallest μ components of x—the total payments Y needs to make to buy off a bare majority of legislators.

- If $m_x < V_Y$, then group Y can buy off a bare majority of legislators for less than V_Y, so that its best response to x is to match group X's payments to the μ legislators to whom group X's payments are smallest; the outcome is that bill Y is passed.

- If $m_x > V_Y$, then the cost to group Y of buying off any majority of legislators exceeds V_Y, so that group Y's best response to x is to make no payments; the outcome is that bill X is passed.

- If $m_x = V_Y$, then both the actions in the previous two cases are best responses by group Y to x.

We conclude that group Y's strategy in a subgame perfect equilibrium has the following properties.

- After a history x for which $m_x < V_Y$, group Y matches group X's payments to the μ legislators to whom X's payments are smallest.

- After a history x for which $m_x > V_Y$, group Y makes no payments.

- After a history x for which $m_x = V_Y$, group Y either makes no payments or matches group X's payments to the μ legislators to whom X's payments are smallest.

Given that group Y's subgame perfect equilibrium strategy has these properties, what should group X do? If it chooses a list of payments x for which $m_x < V_Y$, then group Y matches its payments to a bare majority of legislators, and bill Y passes. If it reduces all its payments, the same bill is passed. Thus the only list of payments x with $m_x < V_Y$ that may be optimal is $(0, \ldots, 0)$. If it chooses a list of payments x with $m_x > V_Y$, then group Y makes no payments, and bill X passes.

If it reduces all its payments a little (keeping the payments to every bare majority greater than V_Y), the outcome is the same. Thus no list of payments x for which $m_x > V_Y$ is optimal.

We conclude that in any subgame perfect equilibrium we have either $x = (0, \ldots, 0)$ (group X makes no payments) or $m_x = V_Y$ (the smallest sum of group X's payments to a bare majority of legislators is V_Y). Under what conditions does each case occur? If group X needs to spend more than V_X to deter group Y from matching its payments to a bare majority of legislators, then its best strategy is to make no payments ($x = (0, \ldots, 0)$). How much does it need to spend to deter group Y? It needs to pay more than V_Y to every bare majority of μ legislators, so that its total payment is more than kV_Y/μ. Thus if $V_X < kV_Y/\mu$, group X is better off making no payments than getting bill X passed by making payments large enough to deter group Y from matching its payments to a bare majority of legislators.

If $V_X > kV_Y/\mu$, on the other hand, group X can afford to make payments large enough to deter group Y from matching. In this case its best strategy is to pay each legislator V_Y/μ, so that its total payment to every bare majority of legislators is V_Y. Given this strategy, group Y is indifferent between matching group X's payments to a bare majority of legislators and making no payments. I claim that the game has no subgame perfect equilibrium in which group Y matches. The argument is similar to the argument that the ultimatum game has no subgame perfect equilibrium in which person 2 rejects the offer 0. Suppose that group Y matches. Then group X can increase its payoff by increasing its payments a little (keeping the total less than V_X), thereby deterring group Y from matching, and ensuring that bill X passes. Thus in any subgame perfect equilibrium group Y makes no payments in response to group X's strategy.

In conclusion, if $V_X \neq kV_Y/\mu$, then the game has a unique subgame perfect equilibrium, in which group Y's strategy is to

- match group X's payments to the μ legislators to whom X's payments are smallest after a history x for which $m_x < V_Y$, and
- make no payments after a history x for which $m_x \geq V_Y$

and group X's strategy depends on the relative sizes of V_X and V_Y:

- if $V_X < kV_Y/\mu$, then group X makes no payments;
- if $V_X > kV_Y/\mu$, then group X pays each legislator V_Y/μ.

If $V_X < kV_Y/\mu$, then the outcome is that neither group makes any payment, and bill Y is passed; if $V_X > kV_Y/\mu$, then the outcome is that group X pays each legislator V_Y/μ, group Y makes no payments, and bill X is passed. (If $V_X = kV_Y/\mu$, then the analysis is more complex.)

Three features of the subgame perfect equilibrium are significant. First, the outcome favors the second-mover in the game (group Y): only if $V_X > kV_Y/\mu$, which is close to $2V_Y$ when k is large, does group X manage to get bill X passed. Second, group Y never makes any payments! According to its equilibrium strategy it is prepared to make payments in response to certain strategies of group X, but given group X's *equilibrium* strategy, it spends not a cent. Third, if group X makes any payments (as it does in the equilibrium for $V_X > kV_Y/\mu$), then it makes a payment to *every* legislator. If there were no competing interest group but nonetheless each legislator would vote for bill X only if she were paid at least some amount, then group X would make payments to only a bare majority of legislators; if it were to act in this way in the presence of group Y, it would supply group Y with almost a majority of legislators who could be induced to vote for bill Y at no cost.

? EXERCISE 196.1 (Three interest groups buying votes) Consider a variant of the model in which there are *three* bills, X, Y, and Z, and *three* interest groups, X, Y, and Z, who choose lists of payments sequentially. Ties are broken in favor of the group moving later. Assume that if each bill obtains the vote of one legislator, then bill X passes. Find the bill passed in any subgame perfect equilibrium when $k = 3$ and (*a*) $V_X = V_Y = V_Z = 300$, (*b*) $V_X = 300$, $V_Y = V_Z = 100$, and (*c*) $V_X = 300$, $V_Y = 202$, $V_Z = 100$. (You may assume that in each case a subgame perfect equilibrium exists; note that you are not asked to find the subgame perfect equilibria themselves.)

? EXERCISE 196.2 (Interest groups buying votes under supermajority rule) Consider another variant of the model in which a supermajority is required to pass a bill. There are two bills, X and Y, and a "default outcome". A bill passes if and only if it receives at least $k^* > \frac{1}{2}(k+1)$ votes; if neither bill passes, the default outcome occurs. There are two interest groups. Both groups attach value 0 to the default outcome. Find the bill that is passed in any subgame perfect equilibrium when $k = 7$, $k^* = 5$, and (*a*) $V_X = V_Y = 700$ and (*b*) $V_X = 750$, $V_Y = 400$. In each case, would the legislators be better off or worse off if a simple majority of votes were required to pass a bill?

? EXERCISE 196.3 (Sequential positioning by two political candidates) Consider the variant of Hotelling's model of electoral competition in Section 3.3 in which the n candidates choose their positions sequentially, rather than simultaneously. Model this situation as an extensive game. Find the subgame perfect equilibrium (equilibria?) when $n = 2$.

?? EXERCISE 196.4 (Sequential positioning by three political candidates) Consider a further variant of Hotelling's model of electoral competition in which the n candidates choose their positions sequentially and each candidate has the option of staying out of the race. Assume that each candidate prefers to stay out than to enter and lose, prefers to enter and tie with any number of candidates than to stay out, and prefers to tie with as few other candidates as possible. Model the situation as an extensive game and find the subgame perfect equilibrium outcomes when $n = 2$ (easy) and when $n = 3$ and the voters' favorite positions are distributed uniformly from 0 to 1 (i.e. the fraction of the voters' favorite positions less than x is x) (hard).

6.4 A race

6.4.1 General model

Firms compete with each other to develop new technologies; authors compete with each other to write books and film scripts about momentous current events; scientists compete with each other to make discoveries. In each case the winner enjoys a significant advantage over the losers, and each competitor can, at a cost, increase her pace of activity. How do the presence of competitors and size of the prize affect the pace of activity? How does the identity of the winner of the race depend on each competitor's initial distance from the finish line?

We can model a race as an extensive game with perfect information in which the players alternately choose how many "steps" to take. Here I study a simple example of such a game, with two players.

Player i is initially $k_i > 0$ steps from the finish line, for $i = 1, 2$. On each of her turns, a player can either not take any steps (at a cost of 0), or can take one step, at a cost of $c(1)$, or two steps, at a cost of $c(2)$. The first player to reach the finish line wins a prize, worth $v_i > 0$ to player i; the losing player's payoff is 0. To make the game finite, I assume that if, on successive turns, neither player takes any step, the game ends and neither player obtains the prize.

I denote the game in which player i moves first by $G_i(k_1, k_2)$. The game $G_1(k_1, k_2)$ is defined precisely as follows.

Players The two parties.

Terminal histories The set of sequences of the form $(x^1, y^1, x^2, y^2, \ldots, x^T)$ or $(x^1, y^1, x^2, y^2, \ldots, y^T)$ for some integer T, where each x^t (the number of steps taken by player 1 on her tth turn) and each y^t (the number of steps taken by player 2 on her tth turn) is 0, 1, or 2, there are never two successive 0's except possibly at the end of a sequence, and either $y^1 + \cdots + y^{T-1} < k_2$ and $x^1 + \cdots + x^T = k_1$ (player 1 reaches the finish line first), or $x^1 + \cdots + x^T < k_1$ and $y^1 + \cdots + y^T = k_2$ (player 2 reaches the finish line first).

Player function $P(\varnothing) = 1$, $P(x^1) = 2$ for all x^1, $P(x^1, y^1) = 1$ for all (x^1, y^1), $P(x^1, y^1, x^2) = 2$ for all (x^1, y^1, x^2), and so on.

Preferences For a terminal history in which player i loses, her payoff is the negative of the sum of the costs of all her moves; for a terminal history in which she wins it is v_i minus the sum of these costs.

6.4.2 Subgame perfect equilibria of an example

A simple example illustrates the features of the subgame perfect equilibria of this game. Suppose that both v_1 and v_2 are between 6 and 7 (their exact values do not affect the equilibria), the cost $c(1)$ of a single step is 1, and the cost $c(2)$ of two steps is 4. (Given that $c(2) > 2c(1)$, each player, in the absence of a competitor, would like to take one step at a time.)

The game has a finite horizon, so we may use backward induction to find its subgame perfect equilibria. Each of its subgames is either a game $G_i(m_1, m_2)$ with $i = 1$ or $i = 2$ and $0 < m_1 \leq k_1$ and $0 < m_2 \leq k_2$, or, if the last player to move before the subgame took no steps, a game that differs from $G_i(m_1, m_2)$ only in that it ends if player i initially takes no steps (i.e. the only terminal history starting with 0 consists only of 0).

First consider the very simplest game, $G_1(1, 1)$, in which each player is initially one step from the finish line. If player 1 takes one step, she wins; if she does not move, then player 2 optimally takes one step (if she does not, the game ends) and wins. We conclude that the game has a unique subgame perfect equilibrium, in which player 1 initially takes one step and wins.

A similar argument applies to the game $G_1(1, 2)$. If player 1 does not move, then player 2 has the option of taking one or two steps. If she takes one step, then play moves to a subgame identical $G_1(1, 1)$, in which we have just concluded that player 1 wins. Thus player 2 takes two steps, and wins, if player 1 does not move at the start of $G_1(1, 2)$. We conclude that the game has a unique subgame perfect equilibrium, in which player 1 initially takes one step and wins.

Now consider player 1's options in the game $G_1(2, 1)$.

- Player 1 takes two steps: she wins, and obtains a payoff of at least $6 - 4 = 2$ (her valuation is more than 6, and the cost of two steps is 4).

- Player 1 takes one step: play moves to a subgame identical to $G_2(1, 1)$; we know that in the equilibrium of this subgame player 2 initially takes one step and wins.

- Player 1 does not move: play moves to a subgame in which player 2 is the first-mover and is one step from the finish line, and, if player 2 does not move, the game ends. In an equilibrium of this subgame, player 2 takes one step and wins.

We conclude that the game $G_1(2, 1)$ has a unique subgame perfect equilibrium, in which player 1 initially takes two steps and wins.

I have spelled out the details of the analysis of these cases to show how we use the result for the game $G_1(1, 1)$ to find the equilibria of the games $G_1(1, 2)$ and $G_1(2, 1)$. In general, the equilibria of the games $G_i(k_1, k_2)$ for all values of k_1 and k_2 up to $\overline{k}$ tell us the consequences of player 1's taking one or two steps in the game $G_1(\overline{k} + 1, \overline{k})$.

? EXERCISE 198.1 (The race $G_1(2, 2)$) Show that the game $G_1(2, 2)$ has a unique subgame perfect equilibrium outcome, in which player 1 initially takes two steps, and wins.

So far we have concluded that in any game in which each player is initially at most two steps from the finish line, the first-mover takes enough steps to reach the finish line, and wins.

Now suppose that player 1 is at most two steps from the finish line, but player 2 is three steps away. Suppose that player 1 takes only *one* step (even if she is initially two steps from the finish line). Then if player 2 takes either one or two steps, play moves to a subgame in which player 1 (the first-mover) wins. Thus player 2 is better off not moving (and not incurring any cost), in which case player 1 takes one step on her next turn, and wins. (Player 1 prefers to move one step at a time than to move two steps initially because the former costs her 2 whereas the latter costs her 4.) We conclude that the outcome of a subgame perfect equilibrium in the game $G_1(2,3)$ is that player 1 takes one step on her first turn, then player 2 does not move, and then player 1 takes another step, and wins.

By a similar argument, in a subgame perfect equilibrium of any game in which player 1 is at most two steps from the finish line and player 2 is three or more steps away, player 1 moves one step at a time, and player 2 does not move; player 1 wins. Symmetrically, in a subgame perfect equilibrium of any game in which player 1 is three or more steps from the finish line and player 2 is at most two steps away, player 1 does not move, and player 2 moves one step at a time, and wins.

Our conclusions so far are illustrated in Figure 200.1, where player 1 moves to the left and player 2 moves down. The values of (k_1, k_2) for which the subgame perfect equilibrium outcome has been determined so far are labeled. The label "1" means that, regardless of who moves first, in a subgame perfect equilibrium player 1 moves one step on each turn, and player 2 does not move; player 1 wins. Similarly, the label "2" means that, regardless of who moves first, player 2 moves one step on each turn, and player 1 does not move; player 2 wins. The label "f" means that the first player to move takes enough steps to reach the finish line, and wins.

Now consider the game $G_1(3,3)$. If player 1 takes one step, we reach the game $G_2(2,3)$. From Figure 200.1 we see that in the subgame perfect equilibrium of this game player 1 wins, and does so by taking one step at a time (the point $(2,3)$ is labeled "1"). If player 1 takes two steps, we reach the game $G_2(1,3)$, in which player 1 also wins. Player 1 prefers not to take two steps unless she has to, so in the subgame perfect equilibrium of $G_1(3,3)$ she takes one step at a time, and wins, and player 2 does not move. Similarly, in a subgame perfect equilibrium of $G_2(3,3)$, player 2 takes one step at a time, and wins, and player 1 does not move.

A similar argument applies to each of the games $G_i(3,4)$, $G_i(4,3)$, and $G_i(4,4)$ for $i = 1$, 2. The argument differs only if the first-mover is four steps from the finish line, in which case she initially takes two steps to reach a game in which she wins. (If she initially takes only one step, the other player wins.)

Now consider the game $G_i(3,5)$ for $i = 1$, 2. By taking one step in $G_1(3,5)$, player 1 reaches a game in which she wins by taking one step at a time. The cost of her taking three steps is less than v_1, so in a subgame perfect equilibrium of $G_1(3,5)$

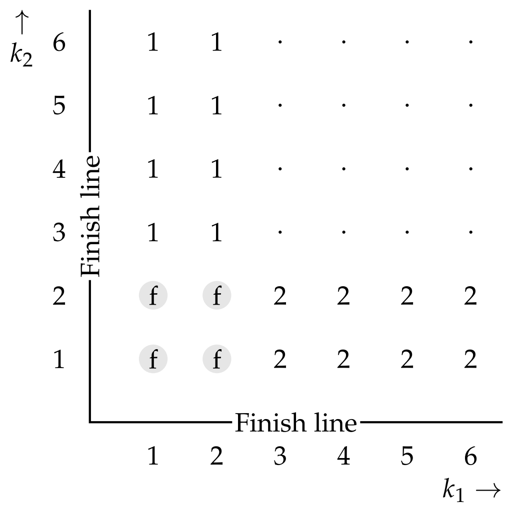

Figure 200.1 The subgame perfect equilibrium outcomes of the race $G_i(k_1, k_2)$. Player 1 moves to the left, and player 2 moves down. The values of (k_1, k_2) for which the subgame perfect equilibrium outcome has been determined so far are labeled; dots represent cases that have not yet been studied. The labels are explained in the text.

she takes one step at a time, and wins, and player 2 does not move. If player 2 takes either one or two steps in $G_2(3,5)$, she reaches a game (either $G_1(3,4)$ or $G_1(3,3)$) in which player 1 wins. Thus whatever she does, she loses, so that in a subgame perfect equilibrium she does not move and player 1 moves one step at a time. We conclude that in a subgame perfect equilibrium of both $G_1(3,5)$ and $G_2(3,5)$, player 1 takes one step on each turn and player 2 does not move; player 1 wins.

A similar argument applies to any game in which one player is initially three or four steps from the finish line and the other player is five or more steps from the finish line. We have now made arguments to justify the labeling in Figure 201.1, where the labels have the same meaning as in Figure 200.1, except that "f" means that the first player to move takes enough steps to reach the finish line *or* to reach the closest point labeled with her name, whichever is closer.

A feature of the subgame perfect equilibrium of the game $G_1(4,4)$ is noteworthy. Suppose that, as planned, player 1 takes two steps, but then player 2 deviates from her equilibrium strategy and takes two steps (rather than not moving). According to our analysis, player 1 should take two steps, to reach the finish line. If she does so, her payoff is negative (less than $7 - 4 - 4 = -1$). Nevertheless she should definitely take the two steps: if she does not, her payoff is even smaller (-4), because player 2 wins. The point is that the cost of her first move is "sunk"; her decision after player 2 deviates must be based on her options from that point on.

The analysis of the games in which each player is initially either five or six steps from the finish line involves arguments similar to those used in the previous cases, with one amendment. A player who is initially six steps from the finish line is

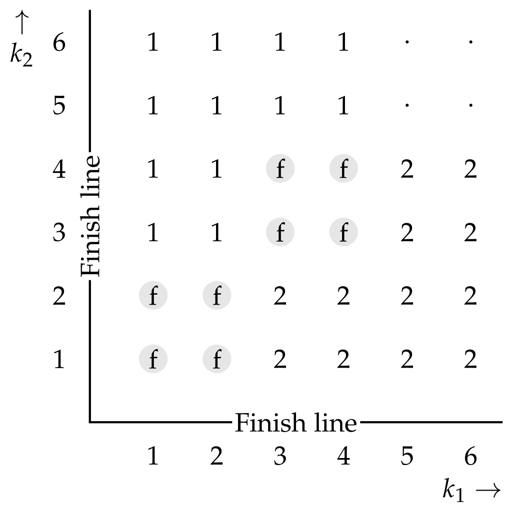

Figure 201.1 The subgame perfect equilibrium outcomes of the race $G_i(k_1, k_2)$. Player 1 moves to the left, and player 2 moves down. The values of (k_1, k_2) for which the subgame perfect equilibrium outcome has been determined so far are labeled; dots represent cases that have not yet been studied. The labels are explained in the text.

better off not moving at all (and obtaining the payoff 0) than she is moving two steps on any turn (and obtaining a negative payoff). An implication is that in the game $G_1(6,5)$, for example, player 1 does not move: if she takes only one step, then player 2 becomes the first-mover and, by taking a single step, moves the play to a game that she wins. We conclude that the first-mover wins in the games $G_i(5,5)$ and $G_i(6,6)$, whereas player 2 wins in $G_i(6,5)$ and player 1 wins in $G_i(5,6)$, for $i = 1, 2$.

A player who is initially more than six steps from the finish line obtains a negative payoff if she moves, even if she wins, so in any subgame perfect equilibrium she does not move. Thus our analysis of the game is complete. The subgame perfect equilibrium outcomes are indicated in Figure 202.1, which shows also the steps taken in the equilibrium of each game when player 1 is the first-mover.

? EXERCISE 201.1 (A race in which the players' valuations of the prize differ) Find the subgame perfect equilibrium outcome of the game in which player 1's valuation of the prize is between 6 and 7, and player 2's valuation is between 4 and 5.

In both of the following exercises, inductive arguments on the length of the game, like the one for $G_i(k_1, k_2)$, can be used.

? EXERCISE 201.2 (Removing stones) Two people take turns removing stones from a pile of n stones. Each person may, on each of her turns, remove either one or two stones. The person who takes the last stone is the winner; she gets \$1 from her opponent. Find the subgame perfect equilibria of the games that model this situation for $n = 1$ and $n = 2$. Find the winner in each subgame perfect equilibrium for

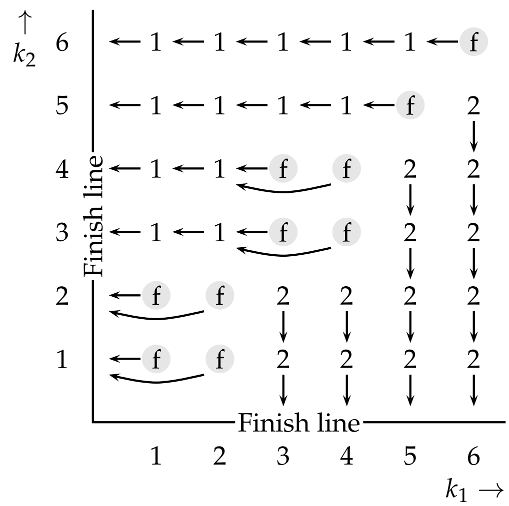

Figure 202.1 The subgame perfect equilibrium outcomes of the race $G_i(k_1, k_2)$. Player 1 moves to the left, and player 2 moves down. The arrows indicate the steps taken in the subgame perfect equilibrium outcome of the games in which player 1 moves first. The labels are explained in the text.

$n = 3$, using the fact that the subgame following player 1's removal of one stone is the game for $n = 2$ in which player 2 is the first-mover, and the subgame following player 1's removal of two stones is the game for $n = 1$ in which player 2 is the first-mover. Use the same technique to find the winner in each subgame perfect equilibrium for $n = 4$, and, if you can, for an arbitrary value of n.

? EXERCISE 202.1 (Hungry lions) The members of a hierarchical group of hungry lions face a piece of prey. If lion 1 does not eat the prey, the prey escapes and the game ends. If it eats the prey, it becomes fat and slow, and lion 2 can eat it. If lion 2 does not eat lion 1, the game ends; if it eats lion 1, then it may be eaten by lion 3, and so on. Each lion prefers to eat than to be hungry, but prefers to be hungry than to be eaten. Find the subgame perfect equilibrium (equilibria?) of the extensive game that models this situation for any number n of lions.

6.4.3 General lessons

Each player's equilibrium strategy involves a "threat" to speed up if the other player deviates. Consider, for example, the game $G_1(3, 3)$. Player 1's equilibrium strategy calls for her to take one step at a time, and player 2's equilibrium strategy calls for her not to move. Thus in the equilibrium outcome, player 1's debt climbs to 3 (the cost of her three single steps) before she reaches the finish line.

Now suppose that after player 1 takes her first step, player 2 deviates and takes a step. Then player 1's strategy calls for her to take two steps, raising her debt to 5. If at no stage can her debt exceed 3 (its maximal level if both players adhere to their equilibrium strategies), then her strategy cannot embody such threats.

The general point is that a limit on the debt a player can accumulate may affect the outcome even if it exceeds the player's debt in the equilibrium outcome in the absence of any limits. You are asked to study an example in the next exercise.

? EXERCISE 203.1 (A race with a liquidity constraint) Find the subgame perfect equilibrium of the variant of the game $G_1(3,3)$ in which player 1's debt may never exceed 3.

In the subgame perfect equilibrium of every game $G_i(k_1, k_2)$, only one player moves; her opponent "gives up". This property of equilibrium holds in more general games. What added ingredient might lead to an equilibrium in which both players are active? A player's uncertainty about the other's characteristics would seem to be such an ingredient: if a player does not know the cost of its opponent's moves, it may assign a positive probability less than one to its winning, at least until it has accumulated some evidence of its opponent's behavior, and while it is optimistic it may be active even though its rival is also active. To build such considerations into the model we need to generalize the model of an extensive game to encompass imperfect information, as we do in Chapter 10.

Another robust feature of the subgame perfect equilibrium of $G_i(k_1, k_2)$ is that the presence of a competitor has little effect on the speed of the player who moves. A lone player would move one step at a time. When there are two players, for most starting points the one that moves does so at the same leisurely pace. Only for a small number of starting points, in all of which the players' initial distances from the starting line are similar, does the presence of a competitor induce the active player to hasten its progress, and then only in the first period.

Notes

The first experiment on the ultimatum game is reported in Güth, Schmittberger, and Schwarze (1982). Grout (1984) is an early analysis of a holdup game. The model of agenda control in legislatures is based on Denzau and Mackay (1983); Romer and Rosenthal (1978) earlier explored a similar idea. The model in Section 6.2 derives its name from the analysis in von Stackelberg (1934, Chapter 4). The vote-buying game in Section 6.3 is taken from Groseclose and Snyder (1996). The model of a race in Section 6.4 is a simplification suggested by Vijay Krishna of a model of Harris and Vickers (1985).

For more discussion of the experimental evidence on the ultimatum game (discussed in the box on page 183), see Roth (1995). Bolton and Ockenfels (2000) study the implications of assuming that players are equity conscious, and relate these implications to the experimental outcomes in various games. The explanation of the experimental results in terms of rules of thumb is discussed by Aumann (1997, 7–8). The problem of fair division, an example of which is given in Exercise 185.2, is studied in detail by Brams and Taylor (1996), who trace the idea of divide-and-choose back to antiquity (p. 10). I have been unable to find the origin of the idea in Exercise 202.1; Barton Lipman suggested the formulation in the exercise.