

# Game Theory

**Drew Fudenberg**  
**Jean Tirole**

*Published in India by*



**Ane Books Pvt. Ltd.**

New Delhi ♦ Chennai ♦ Mumbai  
Kolkata ♦ Thiruvananthapuram ♦ Pune ♦ Bengaluru

3.1 Introduction<sup>†</sup>

In the examples we examined in part I, such as the stag hunt, the prisoner's dilemma, and the battle of the sexes, the players choose their actions simultaneously. Much of the recent interest in the economic applications of game theory has been in situations with an important dynamic structure, such as entry and entry deterrence in industrial organization and the "time-consistency" problem in macroeconomics. Game theorists use the concept of a *game in extensive form* to model such dynamic situations. The extensive form makes explicit the order in which players move, and what each player knows when making each of his decisions. In this setting, strategies correspond to contingent plans instead of uncontingent actions. As we will see, the extensive form can be viewed as a multi-player generalization of a decision tree. Not surprisingly, many results and intuitions from decision theory have game-theoretic analogs. We will also see how to build up the strategic-form representation of a game from its extensive form. Thus, we will be able to apply the concepts and results of part I to dynamic games.

As a simple example of an extensive-form game, consider the idea of a "Stackelberg equilibrium" in a duopoly. As in the Cournot model, the actions of the firms are choices of output levels,  $q_1$  for player 1 and  $q_2$  for player 2. The difference is that we now suppose that player 1, the "Stackelberg leader," chooses her output level  $q_1$  first, and that player 2 observes  $q_1$  before choosing his own output level. To make things concrete, we suppose that production is costless, and that demand is linear, with  $p(q) = 12 - q$ , so that player  $i$ 's payoff is  $u_i(q_1, q_2) = [12 - (q_1 + q_2)]q_i$ . How should we extend the idea of Nash equilibrium to this setting? And how should we expect the players to play?

Since player 2 observes player 1's choice of output  $q_1$  before choosing  $q_2$ , in principle player 2 could condition his choice of  $q_2$  on the observed level of  $q_1$ . And since player 1 moves first, she cannot condition her output on player 2's. Thus, it is natural that player 2's *strategies* in this game should be maps of the form  $s_2: Q_1 \rightarrow Q_2$  (where  $Q_1$  is the space of feasible  $q_1$ 's and  $Q_2$  is the space of feasible  $q_2$ 's), while player 1's strategies are simply choices of  $q_1$ . Given a (pure) strategy profile of this form, the outcome is the output vector  $(q_1, s_2(q_1))$ , with payoffs  $u_i(q_1, s_2(q_1))$ .

Now that we have identified strategy spaces and the payoff functions, we can define a Nash equilibrium of this game in the obvious way: as a strategy profile such that neither player can gain by switching to a different strategy. Let's consider two particular Nash equilibria of this game.

The first equilibrium gives rise to the Stackelberg output levels normally associated with this game. In this equilibrium, player 2's strategy  $s_2$  is to choose, for each  $q_1$ , the level of  $q_2$  that solves  $\max_{q_2} u_2(q_1, q_2)$ , so that  $s_2$  is

identically equal to the Cournot reaction function  $r_2$  defined in chapter 1. With the payoffs we have specified,  $r_2(q_1) = 6 - q_1/2$ .

Nash equilibrium requires that player 1's strategy maximize her payoff given that  $s_2 = r_2$ , so that player 1's output level  $q_1^*$  is the solution to  $\max_{q_1} u_1(q_1, r_2(q_1))$ , which with the payoffs we specified gives  $q_1^* = 6$ .

The output levels  $(q_1^*, r_2(q_1^*))$  (here equal to  $(6, 3)$ ) are called the *Stackelberg outcome* of the game; this is the outcome economics students are taught to expect. In the usual case,  $r_2$  is a decreasing function, and so player 1 can decrease player 2's output by increasing her own. As a result, player 1's Stackelberg output level and payoff are typically higher than in the Cournot equilibrium where both players move simultaneously, and player 2's output and payoff are typically lower. (In our case the unique Cournot equilibrium is  $q_1^C = q_2^C = 4$ , with payoffs of 16 each; in the Stackelberg equilibrium the leader's payoff is 18 and the follower's is 9.)

Though the Stackelberg outcome may seem the natural prediction in this game, there are many other Nash equilibria. One of them is the profile " $q_1 = q_1^C; s_2(q_1) = q_2^C$  for all  $q_1$ ." These strategies really are a Nash equilibrium: Given that player 2's output will be  $q_2^C$  independent of  $q_1$ , player 1's problem is to maximize  $u_1(q_1, q_2^C)$ , and by definition this maximization is solved by the Cournot output  $q_1^C$ . And given that  $q_1 = q_1^C$ , player 2's payoff will be  $u_2(q_1^C, s_2(q_1^C))$ , which is maximized by any strategy  $s_2$  such that  $s_2(q_1^C) = q_2^C$ , including the constant strategy  $s_2(\cdot) \equiv q_2^C$ . Note, though, that this strategy is *not* a best response to other output levels that player 1 might have chosen but did not; i.e.,  $q_2^C$  is not in general a best response to  $q_1$  for  $q_1 \neq q_1^C$ .

So we have identified two Nash equilibria for the game where player 1 chooses her output first: one equilibrium with the "Stackelberg outputs" and one where the output levels are the same as if the players moved simultaneously. Why is the first equilibrium more reasonable, and what is wrong with the second one? Most game theorists would answer that the second equilibrium is "not credible," as it relies on an "empty threat" by player 2 to hold his output at  $q_2^C$  regardless of player 1's choice. This threat is empty because if player 1 were to present player 2 with the *fait accompli* of choosing the Stackelberg output  $q_1^*$ , player 2 would do better to choose a different level of  $q_2$ —in particular,  $q_2 = r_2(q_1^*)$ . Thus, if player 1 knows player 2's payoffs, the argument goes, she should not believe that player 2 would play  $q_2^C$  no matter what player 1's output. Rather, player 1 should predict that player 2 will play an optimal response to whatever  $q_1$  player 1 actually chooses, so that player 1 should predict that whatever level of  $q_1$  she chooses, player 2 will choose the optimal response  $r_2(q_1)$ . This argument picks out the "Stackelberg equilibrium" as the unique credible outcome. A more formal way of putting this is that the Stackelberg equilibrium is consistent with *backward induction*, so called because the idea is to start by solving for the optimal choice of the last mover for each

possible situation he might face, and then work backward to compute the optimal choice for the player before. The ideas of credibility and backward induction are clearly present in the textbook analysis of the Stackelberg game; they were informally applied by Schelling (1960) to the analysis of commitment in a number of settings. Selten (1965) formalized the intuition with his concept of a *subgame-perfect equilibrium*, which extends the idea of backward induction to extensive games where players move simultaneously in several periods, so the backward-induction algorithm is not applicable because there are several "last movers" and each of them must know the moves of the others to compute his own optimal choice.

This chapter will develop the formalism for modeling extensive games and develop the solution concepts of backward induction and subgame perfection. Although the extensive form is a fundamental concept in game theory, its definition may be a bit detailed for readers who are more interested in applications of games than in mastering the general theory. With such readers in mind, section 3.2 presents a first look at dynamic games by treating a class of games with a particularly simple structure: the class of "multi-stage games with observed actions." These games have "stages" such that (1) in each stage every player knows all the actions taken by any player, including "Nature," at any previous stage, and (2) players move "simultaneously" within each stage.

Though very special, this class of games includes the Stackelberg example we have just discussed, as well as many other examples from the economics literature. We use multi-stage games to illustrate the idea that strategies can be contingent plans, and to give a first definition of subgame perfection. As an illustration of the concepts, subsection 3.2.3 discusses how to model the idea of commitment, and addresses the particular example called the "time-consistency problem" in macroeconomics. Readers who lack the time or interest for the general extensive-game model are advised to skip from the end of section 3.2 to section 3.6, which gives a few cautions about the potential drawbacks of the ideas of backward induction and subgame perfection.

Section 3.3 introduces the concepts involved in defining an extensive form. Section 3.4 discusses strategies in the extensive form, called "behavior strategies," and shows how to relate them to the strategic-form strategies discussed in chapters 1 and 2. Section 3.5 gives the general definition of subgame perfection. We postpone discussion of more powerful equilibrium refinements to chapters 8 and 11 in order to first study several interesting classes of games which can be fruitfully analyzed with the tools we develop in this chapter.

Readers who already have some informal understanding of dynamic games and subgame perfection probably already know the material of section 3.2, and are invited to skip directly to section 3.3. (Teaching note: When planning to cover all of this chapter, it is probably not worth taking

the time to teach section 3.2 in class; you may or may not want to ask the students to read it on their own.)

### 3.2 Commitment and Perfection in Multi-Stage Games with Observed Actions<sup>1</sup>

#### 3.2.1 What Is a Multi-Stage Game?

Our first step is to give a more precise definition of a "multi-stage game with observed actions." Recall that we said that this meant that (1) all players knew the actions chosen at all previous stages  $0, 1, 2, \dots, k-1$  when choosing their actions at stage  $k$ , and that (2) all players move "simultaneously" in each stage  $k$ . (We adopt the convention that the first stage is "stage 0" in order to simplify the notation concerning discounting when stages are interpreted as periods.) Players move simultaneously in stage  $k$  if each player chooses his or her action at stage  $k$  without knowing the stage- $k$  action of any other player. Common usage to the contrary, "simultaneous moves" does not exclude games where players move in alternation, as we allow for the possibility that some of the players have the one-element choice set "do nothing." For example, the Stackelberg game has two stages: In the first stage, the leader chooses an output level (and the follower "does nothing"). In the second stage, the follower knows the leader's output and chooses an output level of his own (and the leader "does nothing"). Cournot and Bertrand games are one-stage games: All players choose their actions at once and the game ends. Dixit's (1979) model of entry and entry deterrence (based on work by Spence (1977)) is a more complex example: In the first stage of this game, an incumbent invests in capacity; in the second stage, an entrant observes the capacity choice and decides whether to enter. If there is no entry, the incumbent chooses output as a monopolist in the third stage; if entry occurs, the two firms choose output simultaneously as in Cournot competition.

Often it is natural to identify the "stages" of the game with time periods, but this is not always the case. A counterexample is the Rubinstein-Ståhl model of bargaining (discussed in chapter 4), where each "time period" has two stages. In the first stage of each period, one player proposes an agreement; in the second stage, the other player either accepts or rejects the proposal. The distinction is that time periods refer to some physical measure of the passing of time, such as the accumulation of delay costs in the bargaining model, whereas the stages need not have a direct temporal interpretation.

In the first stage of a multi-stage game (stage 0), all players  $i \in \mathcal{I}$  simultaneously choose actions from choice sets  $A_i(h^0)$ . (Remember that some of the choice sets may be the singleton "do nothing." We let  $h^0 = \emptyset$  be the "history" at the start of play.) At the end of each stage, all players observe

the stage's action profile. Let  $a^0 = (a_1^0, \dots, a_n^0)$  be the stage-0 action profile. At the beginning of stage 1, players know history  $h^1$ , which can be identified with  $a^0$  given that  $h^0$  is trivial. In general, the actions player  $i$  has available in stage 1 may depend on what has happened previously, so we let  $A_i(h^1)$  denote the possible second-stage actions when the history is  $h^1$ . Continuing iteratively, we define  $h^{k+1}$ , the history at the end of stage  $k$ , to be the sequence of actions in the previous periods,

$$h^{k+1} = (a^0, a^1, \dots, a^k),$$

and we let  $A_i(h^{k+1})$  denote player  $i$ 's feasible actions in stage  $k+1$  when the history is  $h^{k+1}$ . We let  $K+1$  denote the total number of stages in the game, with the understanding that in some applications  $K = +\infty$ , corresponding to an infinite number of stages; in this case the "outcome" when the game is played will be an infinite history,  $h^\infty$ . Since each  $h^{k+1}$  by definition describes an entire sequence of actions from the beginning of the game on, the set  $H^{K+1}$  of all "terminal histories" is the same as the set of possible outcomes when the game is played.

In this setting, a *pure strategy for player  $i$*  is simply a contingent plan of how to play in each stage  $k$  for possible history  $h^k$ . (We will postpone discussion of mixed strategies until section 3.3, as they will not be used in the examples we discuss here.) If we let  $H^k$  denote the set of all stage- $k$  histories, and let

$$A_i(H^k) = \bigcup_{h^k \in H^k} A_i(h^k),$$

a pure strategy for player  $i$  is a sequence of maps  $\{s_i^k\}_{k=0}^K$ , where each  $s_i^k$  maps  $H^k$  to the set of player  $i$ 's feasible actions  $A_i(H^k)$  (i.e., satisfies  $s_i^k(h^k) \in A_i(h^k)$  for all  $h^k$ ). It should be clear how to find the sequence of actions generated by a profile of such strategies: The stage-0 actions are  $a^0 = s^0(h^0)$ , the stage-1 actions are  $a^1 = s^1(a^0)$ , the stage-2 actions are  $a^2 = s^2(a^0, a^1)$ , and so on. This is called the *path* of the strategy profile. Since the terminal histories represent an entire sequence of play, we can represent each player  $i$ 's payoff as a function  $u_i: H^{K+1} \rightarrow \mathbb{R}$ . In most applications the payoff functions are additively separable over stages (i.e., each player's overall payoff is some weighted average of single-stage payoffs  $g_i(a^k)$ ,  $k = 0, \dots, K$ ), but this restriction is not necessary.

Since we can assign an outcome in  $H^{K+1}$  to each strategy profile, and a payoff vector to each outcome, we can now compute the payoff to any strategy profile; in an abuse of notation, we will represent the payoff vector to profile  $s$  as  $u(s)$ . A (pure-strategy) *Nash equilibrium* in this context is simply a strategy profile  $s$  such that no player  $i$  can do better with a different strategy, which is the familiar condition that  $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$  for all  $s'_i$ .

The Cournot and Bertrand "equilibria" discussed in chapter 1 are trivial examples of Nash equilibria of multi-stage (actually one-stage) games. We

saw two other examples of Nash equilibria when we discussed the Stackelberg game at the beginning of this chapter. We also saw that some of these Nash equilibria may rely on "empty threats" of suboptimal play at histories that are not expected to occur—that is, at histories off the path of the equilibrium.

### 3.2.2 Backward Induction and Subgame Perfection

In the Stackelberg game, it was easy to see how player 2 "ought" to play, because once  $q_1$  was fixed player 2 faced a simple decision problem. This allowed us to solve for player 2's optimal second-stage choice for each  $q_1$  and then work backward to find the optimal choice for player 1. This algorithm can be extended to other games where only one player moves at each stage. We say that a multi-stage game has *perfect information* if, for every stage  $k$  and history  $h^k$ , exactly one player has a nontrivial choice set—a choice set with more than one element—and all the others have the one-element choice set "do nothing." A simple example of such a game has player 1 moving in stages 0, 2, 4, etc. and player 2 moving in stages 1, 3, 5, and so on. More generally, some players could move several times in a row, and which player gets to move in stage  $k$  could depend on the previous history. The key thing is that only one player moves at each stage  $k$ . Since we have assumed that each player knows the past choices of all rivals, this implies that the single player on move at  $k$  is "perfectly informed" of all aspects of the game except those which will occur in the future.

Backward induction can be applied to any finite game of perfect information, where *finite* means that the number of stages is finite and the number of feasible actions at any stage is finite, too.<sup>1</sup> The algorithm begins by determining the optimal choices in the final stage  $K$  for each history  $h^K$ —that is, the action for the player on move, given history  $h^K$ , that maximizes that player's payoff conditional on  $h^K$  being reached. (There may be more than one maximizing choice; in this case backward induction allows the player to choose any of the maximizers.) Then we work back to stage  $K - 1$ , and determine the optimal action for the player on move there, given that the player on move at stage  $K$  with history  $h^K$  will play the action we determined previously. The algorithm proceeds to "roll back," just as in solving decision problems, until the initial stage is reached. At this point we have constructed a strategy profile, and it is easy to verify that this profile is a Nash equilibrium. Moreover, it has the nice property that each player's actions are optimal at every possible history.

The argument for the backward-induction solution in the two-stage Stackelberg game—that player 1 should be able to forecast player 2's second-stage play—strikes us as quite compelling. In a three-stage game,

1. Section 4.6 extends backward induction to infinite games of perfect information, where there is no last period from which to work backward.

the argument is a bit more complex: The player on move at stage 0 must forecast that the player on move at stage 1 will correctly forecast the play of the player on move at stage 2, which clearly is a more demanding hypothesis. And the arguments for backward induction in longer games require correspondingly more involved hypotheses. For this reason, backward-induction arguments may not be compelling in "long" games. For the moment, though, we will pass over the arguments against backward induction; section 3.6 discusses its limitations in more detail.

As defined above, backward induction applies only to games of perfect information. It can be extended to a slightly larger class of games. For instance, in a multi-stage game, if all players have a dominant strategy in the last stage, given the history of the game (or, more generally, if the last stage is solvable by iterated strict dominance), one can replace the last-stage strategies by the dominant strategies, then consider the penultimate stage and apply the same reasoning, and so on. However, this doesn't define backward induction for games that cannot be solved by this backward-induction version of dominance solvability. Yet one would think that the backward-induction idea of predicting what the players are likely to choose in the future ought to carry over to more general games. Suppose that a firm—call it firm 1—has to decide whether or not to invest in a new cost-reducing technology. Its choice will be observed by its only competitor, firm 2. Once the choice is made and observed, the two firms will choose output levels simultaneously, as in Cournot competition. This is a two-stage game, but not one of perfect information. How should firm 1 forecast the second-period output choice of its opponent? In the spirit of equilibrium analysis, a natural conjecture is that the second-period output choices will be those of a Cournot equilibrium for the prevailing cost structure of the industry. That is, each history  $h^1$  generates a simultaneous-move game between the two firms, and firm 1 forecasts that play in this game will correspond to an equilibrium for the payoffs prevailing under  $h^1$ . This is exactly the idea of Selten's (1965) *subgame-perfect equilibrium*.

Defining subgame perfection requires a few preliminary steps. First, since all players know the history  $h^k$  of moves before stage  $k$ , we can view the game from stage  $k$  on with history  $h^k$  as a game in its own right, which we will denote  $G(h^k)$ . To define the payoff functions in this game, note that if the actions in stages  $k$  through  $K$  are  $a^k$  through  $a^K$ , the final history will be  $h^{K+1} = (h^k, a^k, a^{k+1}, \dots, a^K)$ , and so the payoffs will be  $u_i(h^{K+1})$ . Strategies in  $G(h^k)$  are defined in the obvious way: as maps from histories to actions, where the only histories we need consider are those consistent with  $h^k$ . So now we can speak of the Nash equilibria of  $G(h^k)$ .

Next, any strategy profile  $s$  of the whole game induces a strategy profile  $s|h^k$  on any  $G(h^k)$  in the obvious way: For each player  $i$ ,  $s_i|h^k$  is simply the restriction of  $s_i$  to the histories consistent with  $h^k$ .

**Definition 3.1** A strategy profile  $s$  of a multi-stage game with observed actions is a *subgame-perfect equilibrium* if, for every  $h^k$ , the restriction  $s|_{h^k}$  to  $G(h^k)$  is a Nash equilibrium of  $G(h^k)$ .

This definition reduces to backward induction in finite games of perfect information, for the only Nash equilibrium in game  $G(h^k)$  at the final stage is for the player on move to choose (one of) his preferred action(s) as in backward induction, the only Nash-equilibrium choice in the next-to-last stage given Nash play at the last stage is as in backward induction, and so on.

#### Example 3.1

To illustrate the ideas of this section, consider the following model of strategic investment in a duopoly: Firm 1 and firm 2 currently both have a constant average cost of 2 per unit. Firm 1 can install a new technology with an average cost of 0 per unit; installing the technology costs  $f$ . Firm 2 will observe whether or not firm 1 invests in the new technology. Once firm 1's investment decision is observed, the two firms will simultaneously choose output levels  $q_1$  and  $q_2$  as in Cournot competition. Thus, this is a two-stage game.

To define the payoffs, we suppose that the demand is  $p(q) = 14 - q$  and that each firm's goal is to maximize its net revenue minus costs. Firm 1's payoff is then  $[12 - (q_1 + q_2)]q_1$  if it does not invest, and  $[14 - (q_1 + q_2)]q_1 - f$  if it does; firm 2's payoff is  $[12 - (q_1 + q_2)]q_2$ .

To find the subgame-perfect equilibria, we work backward. If firm 1 does not invest, both firms have unit cost 2, and hence their reaction functions are  $r_i(q_j) = 6 - q_j/2$ . These reaction functions intersect at the point (4, 4), with payoffs of 16 each. If firm 1 does invest, its reaction becomes  $\bar{r}_1(q_2) = 7 - q_2/2$ , the second-stage equilibrium is  $(\frac{16}{3}, \frac{10}{3})$ , and firm 1's total payoff is  $256/9 - f$ . Thus, firm 1 should make the investment if  $256/9 - f > 16$ , or  $f < 112/9$ .

Note that making the investment increases firm 1's second-stage profit in two ways. First, firm 1's profit is higher at any fixed pair of outputs, because its cost of production has gone down. Second, firm 1 gains because firm 2's second-stage output is decreased. The reason firm 2's output is lower is because by lowering its cost firm 1 altered its own second-period incentives, and in particular made itself "more aggressive" in the sense that  $\bar{r}_1(q_2) > r_1(q_2)$  for all  $q_2$ . We say more about this kind of "self-commitment" in the next subsection. Note that firm 2's output would not decrease if it continued to believe that firm 1's cost equaled 2.

### 3.2.3 The Value of Commitment and "Time Consistency"

One of the recurring themes in the analysis of dynamic games has been that in many situations players can benefit from the opportunity to make a binding commitment to play in a certain way. In a one-player game—i.e.,

a decision problem—such commitments cannot be of value, as any payoff that a player could attain while playing according to the commitment could be attained by playing in exactly the same way without being committed to do so. With more than one player, though, commitments can be of value, since by committing himself to a given sequence of actions a player may be able to alter the play of his opponents. This "paradoxical" value of commitment is closely related to our observation in chapter 1 that a player can gain by reducing his action set or decreasing his payoff to some outcomes, provided that his opponents are aware of the change. Indeed, some forms of commitment can be represented in exactly this way.

The way to model the possibility of commitments (and related moves like "promises") is to explicitly include them as actions the players can take. (Schelling (1960) was an early proponent of this view.) We have already seen one example of the value of commitment in our study of the Stackelberg game, which describes a situation where one firm (the "leader") can commit itself to an output level that the follower is forced to take as given when making its own output decision. Under the typical assumption that each firm's optimal reaction  $r_i(q_j)$  is a decreasing function of its opponent's output, the Stackelberg leader's payoff is higher than in the "Cournot equilibrium" outcome where the two firms choose their output levels simultaneously.

In the Stackelberg example, commitment is achieved simply by moving earlier than the opponent. Although this corresponds to a different extensive form than the simultaneous moves of Cournot competition, the set of "physical actions" is in some sense the same. The search for a way to commit oneself can also lead to the use of actions that would not otherwise have been considered. Classic examples include a general burning his bridges behind him as a commitment not to retreat and Odysseus having himself lashed to the mast and ordering his sailors to plug their ears with wax as a commitment not to go to the Sirens' island. (Note that the natural way to model the Odysseus story is with two "players," corresponding to Odysseus before and Odysseus after he is exposed to the Sirens.) Both of these cases correspond to a "total commitment": Once the bridge is burned, or Odysseus is lashed to the mast and the sailors' ears are filled with wax, the cost of turning back or escaping from the mast is taken to be infinite. One can also consider partial commitments, which increase the cost of, e.g., turning back without making it infinite.

As a final example of the value of commitment, we consider what is known as the "time-consistency problem" in macroeconomics. This problem was first noted by Kydland and Prescott (1977); our discussion draws on the survey by Mankiw (1988). Suppose that the government sets the inflation rate  $\pi$ , and has preferences over inflation and output  $y$  represented by  $u_\pi(\pi, y) = y - \pi^2$ , so that it is prepared to tolerate inflation if doing so increases the output level. The working of the macroeconomy is such

that only unexpected inflation changes output:

$$y = y^* + (\pi - \hat{\pi}), \quad (3.1)$$

where  $y^*$  is the "natural level" of output and  $\hat{\pi}$  is the expected inflation.<sup>2</sup>

Regardless of the timing of moves, the agents' expectations of inflation are correct in any pure-strategy equilibrium, and so output is at its natural level. (In a mixed-strategy equilibrium the expectations need only be correct on average.) The variable of interest is thus the level of inflation. Suppose first that the government can commit itself to an inflation rate, i.e., the government moves first and chooses a level of  $\pi$  that is observed by the agents. Then output will equal  $y^*$  regardless of the chosen level of  $\pi$ , so the government should choose  $\pi = 0$ .

As Kydland and Prescott point out, this solution to the commitment game is not "time consistent," meaning that if the agents mistakenly believe that  $\pi$  is set equal to 0 when in fact the government is free to choose any level of  $\pi$  it wishes, then the government would prefer to choose a different level of  $\pi$ . That is, the commitment solution is not an equilibrium of the game without commitment.

If the government cannot commit itself, it will choose the level of inflation that equates the marginal benefit from increased output to the marginal cost of increased inflation. The government's utility function is such that this tradeoff is independent of the level of output or the level of expected inflation, and the government will choose  $\pi = \frac{1}{2}$ . Since output is the same in the two cases, the government does strictly worse without commitment. In the context of monetary policy, the "commitment path" can be interpreted as a "money growth rule," and noncommitment corresponds to a "discretionary policy"; hence the conclusion that "rules can be better than discretion."<sup>3</sup>

As a gloss on the time-consistency problem, let us consider the analogous questions in relation to Stackelberg and Cournot equilibria. If we think of the government and the agents as both choosing output levels, the commitment solution corresponds to the Stackelberg outcome  $(q_1^*, q_2^*)$ . This outcome is not an equilibrium of the game where the government cannot commit itself, because in general  $q_1^*$  is not a best response to  $q_2^*$  when  $q_2^*$

2. Equation 3.1 is a reduced form that incorporates the way that the agents' expectations influence their production decisions and in turn influence output. Since the actions of the agents have been suppressed, the model does not directly correspond to an extensive-form game, but the same intuitions apply. Here is an artificial extensive-form game with the same qualitative properties. The government chooses the money supply  $m$ , and a single agent chooses a nominal price  $p$ . Aggregate demand is  $y = \max(0, m - p)$ , and the agent is constrained to supply all demanders. The agent's utility is  $p - p^2/2m$ , and the government's utility is  $y - (m - 1)^2$ . This does not quite give equation 3.1, but the resulting model has very similar properties.

3. In the extensive-game model where the agent chooses prices (see note 2), the agent chooses  $p = m$  and the commitment solution is to set  $m = 1$ . Without commitment this is not an equilibrium, since for fixed  $p$  the government could gain by choosing a larger value of  $m$ .

is held fixed. The no-commitment solution  $\pi = \frac{1}{2}$  derived above corresponds to a situation of simultaneous moves—that is, to the Cournot outcome.

Whether and when a commitment to a monetary rule is credible have been important topics of theoretical and applied research in macroeconomics. This research has started from the observation that decisions about the money supply are not made once and for all, but rather are made repeatedly. Chapter 5, on repeated games, and chapter 9, on reputation effects, discuss game-theoretic analyses of the question of when repeated play makes commitments credible.

Finally, note that a player does not always do better when he moves first (and his choice of action is observed) than when players move simultaneously: In "matching pennies" (example 1.6) each player's equilibrium payoff is 0, whereas if one player moves first his equilibrium payoff is  $-1$ .

### 3.3 The Extensive Form<sup>11</sup>

This section gives a formal development of the idea of an extensive-form game. The extensive form is a fundamental concept in game theory and one to which we will refer frequently, particularly in chapters 8 and 11, but the details of the definitions are not essential for much of the material in the rest of the book. Thus, readers who are primarily interested in applications of the theory should not be discouraged if they do not master all the fine points of the extensive-form methodology. Instead of dwelling on this section, they should proceed along, remembering to review this material before beginning section 8.3.

#### 3.3.1 Definition

The extensive form of a game contains the following information:

- (1) the set of players
- (2) the order of moves—i.e., who moves when
- (3) the players' payoffs as a function of the moves that were made
- (4) what the players' choices are when they move
- (5) what each player knows when he makes his choices
- (6) the probability distributions over any exogenous events.

The set of players is denoted by  $i \in \mathcal{I}$ ; the probability distributions over exogenous events (point 6) are represented as moves by "Nature," which is denoted by  $N$ . The order of play (point 2) is represented by a *game tree*,  $T$ , such as the one shown in figure 3.1.<sup>4</sup> A tree is a finite collection of ordered

4. Our development of the extensive form follows that of Kreps and Wilson 1982 with a simplification suggested by Jim Rattliff. Their assumptions (and ours) are equivalent to those of Kuhn 1953.




Figure 3.1

nodes  $x \in X$  endowed with a precedence relation denoted by  $>$ ;  $x > x'$  means "x is before x'." We assume that the precedence relation is transitive (if x is before x' and x' is before x'', then x is before x'') and asymmetric (if x is before x', then x' is not before x). These assumptions imply that the precedence relation is a *partial order*. (It is not a complete order, because two nodes may not be comparable: In figure 3.1,  $z_3$  is not before  $x''$ , and  $x''$  is not before  $z_3$ .) We include a single initial node  $o \in X$  that is before all other nodes in  $X$ ; this node will correspond to a move by nature if any. Figure 3.1 describes a situation where "nature's move" is trivial, as nature simply gives the move to player 1. As in this figure, we will suppress nature's move whenever it is trivial, and begin the tree with the first "real" choice. The initial node will be depicted with  $\circ$  to distinguish it from the others. In figure 3.1, the precedence order is from the top of the diagram down. Given the assumptions we will impose, the precedence ordering will be clear in most diagrams; when the intended precedence is not clear we will use arrows ( $\rightarrow$ ) to connect a node to its immediate successors.

The assumption that precedence is a partial order rules out cycles of the kind shown in figure 3.2a: If  $x > x' > x'' > x$ , then by transitivity  $x'' > x'$ . Since we already have  $x' > x''$ , this would violate the asymmetry condition. However, the partial ordering does not rule out the situation shown in figure 3.2b, where both x and x' are immediate predecessors of node x''.

We wish to rule out the situation in figure 3.2b, because each node of the tree is meant to be a complete description of all events that preceded it, and not just of the "physical situation" at a given point in time. For example, in figure 3.2c, a firm in each of two markets, A and B, might have entered A and then B (node x and then x'') or B and then A (node x' and then x''), but we want our formalism to distinguish between these two sequences of events instead of describing them by a single node x''. (Of course, we are free to specify that both sequences lead to the same payoff for the firm.) In order to ensure that there is only one path through the tree to a given node, so that each node is a complete description of the path




Figure 3.2




Figure 3.3

preceding it, we require that each node x (except the initial node o) have exactly one immediate predecessor—that is, one node  $x' > x$  such that  $x'' > x$  and  $x'' \neq x'$  implies  $x'' > x'$ . Thus, if x' and x'' are both predecessors of x, then either x' is before x'' or x'' is before x'. (This makes the pair  $(X, >)$  an *arborescence*.)


The nodes that are not predecessors of any other node are called "terminal nodes" and denoted by  $z \in Z$ . Because each z completely determines a path through the tree, we can assign payoffs to sequences of moves using functions  $u_i: Z \rightarrow \mathbb{R}$ , with  $u_i(z)$  being player i's payoff if terminal node z is reached. In drawing extensive forms, the payoff vectors (point 3 in the list above) are displayed next to the corresponding terminal nodes, as in figures 3.3 and 3.4. To complete the specification of point 2 (who moves when), we introduce a map  $i: X \rightarrow \mathcal{P}$  with the interpretation that player  $i(x)$  moves at node x. Next we must describe what player  $i(x)$ 's choices are, which was point 4 of our list. To do so, we introduce a finite set A of actions and a function  $\ell$  that labels each noninitial node x with the last action taken to reach it. We require that  $\ell$  be one-to-one on the set of immediate successors of each node x, so that different successors correspond to different actions, and let  $A(x)$  denote the set of feasible actions at x. (Thus  $A(x)$  is the range of  $\ell$  on the set of immediate successors of x.)

Point 5, the information players have when choosing their actions, is the most subtle of the six points. This information is represented using




information sets  $h \in H$ , which partition the nodes of the tree—that is, every node is in exactly one information set.<sup>5</sup> The interpretation of the information set  $h(x)$  containing node  $x$  is that the player who is choosing an action at  $x$  is uncertain if he is at  $x$  or at some other  $x' \in h(x)$ . We require that if  $x' \in h(x)$  the same player move at  $x$  and  $x'$ . Without this requirement, the players might disagree about who was supposed to move. Also, we require that if  $x' \in h(x)$  then  $A(x') = A(x)$ , so the player on move has the same set of choices at each node of this information set. (Otherwise he might “play” an infeasible action.) Thus, we can let  $A(h)$  denote the action set at information set  $h$ .

A special case of interest is that of games of perfect information, in which all the information sets are singletons. In a game of perfect information, players move one at a time, and each player knows all previous moves when making his decision. The Stackelberg game we discussed at the start of this chapter is a game of perfect information. Figure 3.3 displays a tree for this



a



b

Figure 3.4

5. Note that we use the same notation,  $h$ , for information sets and for histories in multi-stage games. This should not cause too much confusion, especially as information sets can be viewed as a generalization of the idea of a history.

game on the assumption that each player has only three possible output levels: 3, 4, and 6. The vectors at the end of each branch of the tree are the payoffs of players 1 and 2, respectively.

Figure 3.4a displays an extensive form for the Cournot game, where players 1 and 2 choose their output levels simultaneously. Here player 2 does not know player 1's output level when choosing his own output. We model this by placing the nodes corresponding to player 1's three possible actions in the same information set for player 2. This is indicated in the figure by the broken line connecting the three nodes. (Some authors use “loops” around the nodes instead.) Note well the way simultaneous moves are represented: As in figure 3.3, player 1's decision comes “before” player 2's in terms of the precedence ordering of the tree; the difference is in player 2's information set. As this shows, the precedence ordering in the tree need not correspond to calendar time. To emphasize this point, consider the extensive form in figure 3.4b, which begins with a move by player 2. Figures 3.4a and 3.4b describe exactly the same strategic situation: Each player chooses his action not knowing the choice of his opponent. However, the situation represented in figure 3.3, where player 2 observed player 1's move before choosing his own, can only be described by an extensive form in which player 1 moves first.

Almost all games in the economics literature are games of perfect recall: No player ever forgets any information he once knew, and all players know the actions they have chosen previously. To impose this formally, we first require that if  $x$  and  $x'$  are in the same information set then neither is a predecessor of the other. This is not enough to ensure that a player never forgets, as figure 3.5 shows. To rule out this situation, we require that if  $x'' \in h(x')$ , if  $x$  is a predecessor of  $x'$ , and if the same player  $i$  moves at  $x$  and at  $x'$  (and thus at  $x''$ ), then there is a node  $\hat{x}$  (possibly  $x$  itself) that is in the same information set as  $x$ , that  $\hat{x}$  is a predecessor of  $x''$ , and that the action taken at  $x$  along the path to  $x'$  is the same as the action taken at  $\hat{x}$  along the path to  $x''$ . Intuitively, the nodes  $x'$  and  $x''$  are distinguished by information the player doesn't have, so he can't have had it when he was at information set  $h(x)$ ;  $x'$  and  $x''$  must be consistent with the same action at  $h(x)$ , since the player remembers his action there.

When a game involves moves by Nature, the exogenous probabilities are displayed in brackets, as in the two-player extensive form of figure 3.6. In




Figure 3.5




Figure 3.6

figure 3.6, Nature moves first and chooses a "type" or "private information" for player 1. With probability 0.6 player 1 learns that his type is "tough" (T), and with probability 0.4 he learns that his type is "weak" (W). Player 1 then plays left (L) or right (R). Player 2 observes player 1's action but not his type, and chooses between up (U) and down (D). Note that we have allowed both players' payoffs to depend on the choice by Nature even though this choice is initially observed only by player 1. (Player 2 will be able to infer Nature's move from his payoffs.) Figure 3.6 is an example of a "signaling game," as player 1's action may reveal information about his type to player 2. Signaling games, the simplest games of incomplete information, will be studied in detail in chapters 8 and 11.

### 3.3.2 Multi-Stage Games with Observed Actions

Many of the applications of game theory to economics, political science, and biology have used the special class of extensive forms that we discussed in section 3.2: the class of "multi-stage games with observed actions."<sup>6</sup> These games have "stages" such that (1) in each stage  $k$  every player knows all the actions, including those by Nature, that were taken at any previous stage; (2) each player moves at most once within a given stage; and (3) no information set contained in stage  $k$  provides any knowledge of play in that stage. (Exercise 3.4 asks you to give a formal definition of these conditions in terms of a game tree and information sets.)

In a multi-stage game, all past actions are common knowledge at the beginning of stage  $k$ , so there is a well-defined "history"  $h^k$  at the start of each stage  $k$ . Here a pure strategy for player  $i$  is a function  $s_i$  that specifies an action  $a_i \in A_i(h^k)$  for each  $k$  and each history  $h^k$ ; mixed strategies specify probability mixtures over the actions in each stage.

**Caution** Although the idea of a multi-stage game seems natural and intrinsic, it suffers from the following drawback: There may be two exten-

6. Such games are also often called "games of almost-perfect information."




Figure 3.7

sive forms that seem to represent the same real game, with one of them a multi-stage game and the other one not. Consider for example figure 3.7. The extensive form on the left is not a multi-stage game: Player 2's information set is not a singleton, and so it must belong to the first stage and not to a second one. However, player 2 does have some information about player 1's first move (if player 2's information set is reached, then player 1 did not play C), so player 2's information set cannot belong to the first stage either. However, the extensive form on the right is a two-stage game, and the two extensive forms seem to depict the same situation: When player 2 moves, he knows that player 1 is choosing A or B but not C; player 1 chooses A or B without knowing player 2's choice of L or R. The question as to which extensive forms are "equivalent" is still a topic of research—see Elmes and Reny 1988. We will have more to say about this topic when we discuss recent work on equilibrium refinements in chapter 11.

Before proceeding to the next section, we should point out that in applications the extensive form is usually described without using the apparatus of the formal definition, and game trees are virtually never drawn except for very simple "toy" examples. The test of a good informal description is whether it provides enough information to construct the associated extensive form; if the extensive form is not clear, the model has not been well specified.

## 3.4 Strategies and Equilibria in Extensive-Form Games<sup>11</sup>

### 3.4.1 Behavior Strategies

This section defines strategies and equilibria in extensive-form games and relates them to strategies and equilibria of the strategic-form model. Let  $H_i$  be the set of player  $i$ 's information sets, and let  $A_i \equiv \bigcup_{h_i \in H_i} A_i(h_i)$  be the set of all actions for player  $i$ . A pure strategy for player  $i$  is a map  $s_i: H_i \rightarrow A_i$ , with  $s_i(h_i) \in A_i(h_i)$  for all  $h_i \in H_i$ . Player  $i$ 's space of pure strategies,  $S_i$ , is simply the space of all such  $s_i$ . Since each pure strategy is a map from information sets to actions, we can write  $S_i$  as the Cartesian

product of the action spaces at each  $h_i$ :

$$S_i := \prod_{h_i \in H_i} A(h_i).$$

In the Stackelberg example of figure 3.3, player 1 has a single information set and three actions, so that he has three pure strategies. Player 2 has three information sets, corresponding to the three possible choices of player 1, and player 2 has three possible actions at each information set, so player 2 has 27 pure strategies in all. More generally, the number of player  $i$ 's pure strategies,  $\#S_i$ , equals

$$\prod_{h_i \in H_i} \#(A(h_i)).$$

Given a pure strategy for each player  $i$  and the probability distribution over Nature's moves, we can compute a probability distribution over outcomes and thus assign expected payoffs  $u_i(s)$  to each strategy profile  $s$ . The information sets that are reached with positive probability under profile  $s$  are called the *path* of  $s$ .

Now that we have defined the payoffs to each pure strategy, we can proceed to define a pure-strategy Nash equilibrium for an extensive-form game as a strategy profile  $s^*$  such that each player  $i$ 's strategy  $s_i^*$  maximizes his expected payoff given the strategies  $s_{-i}^*$  of his opponents. Note that since the definition of Nash equilibrium holds the strategies of player  $i$ 's opponents fixed in testing whether player  $i$  wishes to deviate, it is as if the players choose their strategies simultaneously. This does *not* mean that in Nash equilibrium players necessarily choose their actions simultaneously. For example, if player 2's fixed strategy in the Stackelberg game of figure 3.3 is the Cournot reaction function  $s_2 = (4,4,3)$ , then when player 1 treats player 2's strategy as fixed he does not presume that player 2's action is unaffected by his own, but rather that player 2 will respond to player 1's action in the way specified by  $s_2$ .

To fill in the details missing from our discussion of the Stackelberg game in the introduction: The "Stackelberg equilibrium" of this game is the outcome  $q_1 = 6, q_2 = 3$ . This outcome corresponds to the Nash-equilibrium strategy profile  $s_1 = 6, s_2 = s_2$ . The Cournot outcome is (4,4); this is the outcome of the Nash equilibrium  $s_1 = 4, s_2 = (4,4,4)$ .

The next order of business is to define mixed strategies and mixed-strategy equilibria for extensive-form games. Such strategies are called *behavior strategies* to distinguish them from the strategic-form mixed strategies we introduced in chapter 1. Let  $\Delta(A(h_i))$  be the probability distributions on  $A(h_i)$ . A *behavior strategy* for player  $i$ , denoted  $b_i$ , is an element of the Cartesian product  $\times_{h_i \in H_i} \Delta(A(h_i))$ . That is, a behavior strategy specifies a probability distribution over actions at each  $h_i$ , and the probability distributions at different information sets are independent. (Note that a

pure strategy is a special kind of behavior strategy in which the distribution at each information set is degenerate.) A profile  $b = (b_1, \dots, b_I)$  of behavior strategies generates a probability distribution over outcomes in the obvious way, and hence gives rise to an expected payoff for each player. A *Nash equilibrium in behavior strategies* is a profile such that no player can increase his expected payoff by using a different behavior strategy.

### 3.4.2 The Strategic-Form Representation of Extensive-Form Games

Our next step is to relate extensive-form games and equilibria to the strategic-form model. To define a strategic form from an extensive form, we simply let the pure strategies  $s \in S$  and the payoffs  $u_i(s)$  be exactly those we defined in the extensive form. A different way of saying this is that the same pure strategies can be interpreted as either extensive-form or strategic-form objects. With the extensive-form interpretation, player  $i$  "waits" until  $h_i$  is reached before deciding how to play there; with the strategic-form interpretation, he makes a complete contingent plan in advance.

Figure 3.8 illustrates this passage from the extensive form to the strategic form in a simple example. We order player 2's information sets from left to right, so that, for example, the strategy  $s_2 = (L, R)$  means that he plays L after U and R after D.

As another example, consider the Stackelberg game illustrated in figure 3.3. We will again order player 2's information sets from left to right, so that player 2's strategy  $s_2 = (4,4,3)$  means that he plays 4 in response to  $q_1 = 3$ , plays 4 in response to 4, and plays 3 in response to 6. (This strategy happens to be player 2's Cournot reaction function.) Since player 2 has three information sets and three possible actions at each of these sets, he has 27 pure strategies. We trust that the reader will forgive our not displaying the strategic form in a matrix diagram!




Figure 3.8

There can be several extensive forms with the same strategic form, as the example of simultaneous moves shows: Figures 3.4a and 3.4b both correspond to the same strategic form for the Cournot game.

At this point we should note that the strategy space as we have defined it may be unnecessarily large, as it may contain pairs of strategies that are "equivalent" in the sense of having the same consequences regardless of how the opponents play.

**Definition 3.2** Two pure strategies  $s_i$  and  $s'_i$  are *equivalent* if they lead to the same probability distribution over outcomes for all pure strategies of the opponents.

Consider the example in figure 3.9. Here player 1 has four pure strategies: (a, c), (a, d), (b, c), and (b, d). However, if player 1 plays b, his second information set is never reached, and the strategies (b, c) and (b, d) are equivalent.

**Definition 3.3** The *reduced strategic form* (or reduced normal form) of an extensive-form game is obtained by identifying equivalent pure strategies (i.e., eliminating all but one member of each equivalence class).

Once we have derived the strategic form from the extensive form, we can (as in chapter 1) define mixed strategies to be probability distributions over pure strategies in the reduced strategic form. Although the extensive form and the strategic form have exactly the same pure strategies, the sets of mixed and behavior strategies are different. With behavior strategies, player  $i$  performs a different randomization at each information set. Luce and Raiffa (1957) use the following analogy to explain the relationship between mixed and behavior strategies: A pure strategy is a book of instructions, where each page tells how to play at a particular information set. The strategy space  $S_i$  is like a library of these books, and a mixed strategy is a probability measure over books—i.e., a random way of making a selection from the library. A given behavior strategy, in contrast, is a single book, but it prescribes a random choice of action on each page.

The reader should suspect that these two kinds of strategies are closely related. Indeed, they are equivalent in games of perfect recall, as was proved




Figure 3.9

by Kuhn (1953). (Here we use "equivalence" as in our earlier definition: Two strategies are equivalent if they give rise to the same distributions over outcomes for all strategies of the opponents.)

**3.4.3 The Equivalence between Mixed and Behavior Strategies in Games of Perfect Recall**

The equivalence between mixed and behavior strategies under perfect recall is worth explaining in some detail, as it also helps to clarify the workings of the extensive-form model. Any mixed strategy  $\sigma_i$  of the strategic form (not of the reduced strategic form) generates a unique behavior strategy  $b_i$  as follows: Let  $R_i(h_i)$  be the set of player  $i$ 's pure strategies that do not preclude  $h_i$ , so that for all  $s_i \in R_i(h_i)$  there is a profile  $s_{-i}$  for player  $i$ 's opponents that reaches  $h_i$ . If  $\sigma_i$  assigns positive probability to some  $s_i \in R_i(h_i)$ , define the probability that  $b_i$  assigns to  $a_i \in A(h_i)$  as

$$b_i(a_i|h_i) = \frac{\sum_{\{s_i \in R_i(h_i) \text{ and } s_i(h_i)=a_i\}} \sigma_i(s_i)}{\sum_{\{s_i \in R_i(h_i)\}} \sigma_i(s_i)}$$

If  $\sigma_i$  assigns probability 0 to all  $s_i \in R_i(h_i)$ , then set

$$b_i(a_i|h_i) = \sum_{\{s_i(h_i)=a_i\}} \sigma_i(s_i)^7$$

In either case, the  $b_i(\cdot|h_i)$  are nonnegative, and

$$\sum_{a_i \in A(h_i)} b_i(a_i|h_i) = 1,$$

because each  $s_i$  specifies an action for player  $i$  at  $h_i$ .

Note that in the notation  $b_i(a_i|h_i)$ , the variable  $h_i$  is redundant, as  $a_i \in A(h_i)$ , but the conditioning helps emphasize that  $a_i$  is an action that is feasible at information set  $h_i$ .

It is useful to work through some examples to illustrate the construction of behavior strategies from mixed strategies. In figure 3.10, a single player (player 1) moves twice. Consider the mixed strategy  $\sigma_1 = (\frac{1}{2}(L, l), \frac{1}{2}(R, r))$ .




Figure 3.10

7. Since  $h_i$  cannot be reached under  $\sigma_i$ , the behavior strategies at  $h_i$  are arbitrary in the same sense that Bayes' rule does not determine posterior probabilities after probability-0 events. Our formula is one of many possible specifications.

The strategy plays  $a$  with probability 1 at information set  $h'_1$ , as only  $(R, a) \in R_1(h'_1)$ .

Figure 3.11 gives another example. Player 2's strategy  $\sigma_2$  assigns probability  $\frac{1}{2}$  each to  $s_2 = (L, L', R'')$  and  $\bar{s}_2 = (R, R', L'')$ . The equivalent behavior strategy is

$$b_2(L|h_2) = b_2(R|h_2) = \frac{1}{2};$$

$$b_2(L'|h'_2) = 0 \text{ and } b_2(R'|h'_2) = 1,$$

and

$$b_2(L''|h''_2) = b_2(R''|h''_2) = \frac{1}{2}.$$

Many different mixed strategies can generate the same behavior strategy. This can be seen from figure 3.12, where player 2 has four pure strategies:  $s_2 = (A, C)$ ,  $s'_2 = (A, D)$ ,  $s''_2 = (B, C)$ , and  $s'''_2 = (B, D)$ .

Now consider two mixed strategies:  $\sigma_2 = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ , which assigns probability  $\frac{1}{4}$  to each pure strategy, and  $\hat{\sigma}_2 = (\frac{1}{2}, 0, 0, \frac{1}{2})$ , which assigns probability  $\frac{1}{2}$  to  $s_2$  and  $\frac{1}{2}$  to  $s'_2$ . Both of these mixed strategies generate the behavior strategy  $b_2$ , where  $b_2(A|h) = b_2(B|h) = \frac{1}{2}$  and  $b_2(C|h') = b_2(D|h') = \frac{1}{2}$ .




Figure 3.11




Figure 3.12

Moreover, for any strategy  $\sigma_1$  of player 1,  $\sigma_2$ ,  $\hat{\sigma}_2$ , and  $b_2$  all lead to the same probability distribution over terminal nodes; for example, the probability of reaching node  $z_1$  equals the probability of player 1's playing U times  $b_2(A|h)$ .

The relationship between mixed and behavior strategies is different in the game illustrated in figure 3.13, which is not a game of perfect recall. (Exercise 3.2 asks you to verify this using the formal definition.) Here, player 1 has four strategies in the strategic form:

$$s_1 = (A, C), s'_1 = (A, D), s''_1 = (B, C), s'''_1 = (B, D).$$

Now consider the mixed strategy  $\sigma_1 = (\frac{1}{2}, 0, 0, \frac{1}{2})$ . As in the last example, this generates the behavior strategy  $b_1 = \{(\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2})\}$ , which says that player 1 mixes  $\frac{1}{2}$ - $\frac{1}{2}$  at each information set. But  $b_1$  is *not* equivalent to the  $\sigma_1$  that generated it. Consider the strategy  $s_2 = L$  for player 2. Then  $(\sigma_1, L)$  generates a  $\frac{1}{2}$  probability of the terminal node corresponding to  $(A, L, C)$ , and a  $\frac{1}{2}$  probability of  $(B, L, D)$ . However, since behavior strategies describe independent randomizations at each information set,  $(b_1, L)$  assigns probability  $\frac{1}{4}$  to each of the four paths  $(A, L, C)$ ,  $(A, L, D)$ ,  $(B, L, C)$ , and  $(B, L, D)$ . Since both A vs. B and C vs. D are choices made by player 1, the strategic-form strategy  $\sigma_1$  can have the property that both A and B have positive probability but C is played wherever A is. Put differently, the strategic-form strategies, where player 1 makes all his decisions at once, allow the decisions at different information sets to be *correlated*. Behavior strategies can't produce this correlation in the example, because when it comes time to choose between C and D player 1 has forgotten whether he chose A or B. This forgetfulness means that there is not perfect recall in this game. If we change the extensive form so that there is perfect recall (by partitioning player 1's second information set into two, corresponding to his choice of A or B), it is easy to see that every mixed strategy is indeed equivalent to the behavior strategy it generates.

**Theorem 3.1** (Kuhn 1953) In a game of perfect recall, mixed and behavior strategies are equivalent. (More precisely: Every mixed strategy is equiv-




Figure 3.13

alent to the unique behavior strategy it generates, and each behavior strategy is equivalent to every mixed strategy that generates it.)

We will restrict our attention to games of perfect recall throughout this book, and will use the terms "mixed strategy" and "Nash equilibrium" to refer to the mixed and behavior formulations interchangeably. This leads us to the following important *notational convention*: In the rest of part II and in most of part IV (except in sections 8.3 and 8.4), we will be studying behavioral strategies. Thus, when we speak of a mixed strategy of an extensive form, we will mean a behavior strategy unless we state otherwise. Although the distinction between the mixed strategy  $\sigma_i$  and the behavior strategy  $b_i$  was necessary to establish their equivalence, we will follow standard usage by denoting both objects by  $\sigma_i$  (thus, the notation  $b_i$  is not used in the rest of the book). In a multi-stage game with observed actions, we will let  $\sigma_i(a_i^k|h^k)$  denote player  $i$ 's probability of playing action  $a_i^k \in A_i(h^k)$  given the history of play  $h^k$  at stage  $k$ . In general extensive forms (with perfect recall), we let  $\sigma_i(a_i|h_i)$  denote player  $i$ 's probability of playing action  $a_i$  at information set  $h_i$ .

#### 3.4.4 Iterated Strict Dominance and Nash Equilibrium

If the extensive form is finite, so is the corresponding strategic form, and the Nash existence theorem yields the existence of a mixed-strategy equilibrium. The notion of iterated strict dominance extends to extensive-form games as well; however, as we mentioned above, this concept turns out to have little force in most extensive forms. The point is that a player cannot strictly prefer one action over another at an information set that is not reached given his opponents' play.

Consider figure 3.14. Here, player 2's strategy R is not strictly dominated, as it is as good as L when player 1 plays U. Moreover, this fact is not "pathological." It obtains for all strategic forms whose payoffs are derived from an extensive form with the tree on the left-hand side of the figure. That is, for any assignment of payoffs to the terminal nodes of the tree, the payoffs to (U, L) and (U, R) must be the same, as both strategy profiles lead to the same terminal node. This shows that the set of strategic-form payoffs of a fixed game tree is of lower dimension than the set of all payoffs of the corresponding strategic form, so theorems based on generic strategic-form payoffs (see chapter 12) do not apply. In particular, there can be an even number of Nash equilibria for an open set of extensive-form payoffs. The game illustrated in figure 3.14 has two Nash equilibria, (U, R) and (D, L), and this number is not changed if the extensive-form payoffs are slightly perturbed. The one case where the odd-number theorem of chapter 12 applies is to a simultaneous-move game such as that of figure 3.4; in such a game, each terminal node corresponds to a unique strategy profile. Put




Figure 3.14

differently: In simultaneous-move games, every strategy profile reaches every information set, and so no player's strategy can involve a choice that is not implemented given his opponents' play.

Recall that a game of perfect information has all its information sets as singletons, as in the games illustrated in figures 3.3 and 3.14.

**Theorem 3.2** (Zermelo 1913; Kuhn 1953) A finite game of perfect information has a pure-strategy Nash equilibrium.

The proof of this theorem constructs the equilibrium strategies using "Zermelo's algorithm," which is a many-player generalization of backward induction in dynamic programming. Since the game is finite, it has a set of penultimate nodes—i.e., nodes whose immediate successors are terminal nodes. Specify that the player who can move at each such node chooses whichever strategy leads to the successive terminal node with the highest payoff for him. (In case of a tie, make an arbitrary selection.) Now specify that each player at nodes whose immediate successors are penultimate nodes chooses the action that maximizes her payoff over the feasible successors, given that players at the penultimate nodes play as we have just specified. We can now roll back through the tree, specifying actions at each node. When we are done, we will have specified a strategy for each player, and it is easy to check that these strategies form a Nash equilibrium. (In fact, the strategies satisfy the more restrictive concept of subgame perfection, which we will introduce in the next section.)

Zermelo's algorithm is not well defined if the hypotheses of the theorem are weakened. First consider infinite games. An infinite game necessarily has either a single node with an infinite number of successors (as do games with a continuum of actions) or a path consisting of an infinite number of nodes (as do multi-stage games with an infinite number of stages). In the first case, an optimal choice need not exist without further restrictions on

the payoff functions<sup>8</sup>; in the second, there need not be a penultimate node on a given path from which to work backward. Finally, consider a game of imperfect information in which some of the information sets are not singletons, as in figure 3.4a. Here there is no way to define an optimal choice for player 2 at his information set without first specifying player 2's belief about the previous choice of player 1; the algorithm fails because it presumes that such an optimal choice exists at every information set given a specification of play at its successors.

We will have much more to say about this issue when we treat equilibrium refinements in detail. We conclude this section with one caveat about the assertion that the Nash equilibrium is a minimal requirement for a "reasonable" point prediction: Although the Nash concept can be applied to any game, the assumption that each player correctly forecasts his opponents' strategy may be less plausible when the strategies correspond to choices of contingent plans than when the strategies are simply choices of actions. The issue here is that when some information sets may not be reached in the equilibrium, Nash equilibrium requires that players correctly forecast their opponents' play at information sets that have 0 probability according to the equilibrium strategies. This may not be a problem if the forecasts are derived from introspection, but if the forecasts are derived from observations of previous play it is less obvious why forecasts should be correct at the information sets that are not reached. This point is examined in detail in Fudenberg and Kreps 1988 and in Fudenberg and Levine 1990.

### 3.5 Backward Induction and Subgame Perfection<sup>11</sup>

As we have seen, the strategic form can be used to represent arbitrarily complex extensive-form games, with the strategies of the strategic form being complete contingent plans of action in the extensive form. Thus, the concept of Nash equilibrium can be applied to all games, not only to games where players choose their actions simultaneously. However, many game theorists doubt that Nash equilibrium is the right solution concept for

<sup>8</sup> The existence of an optimal choice from a compact set of actions requires that payoffs be upper semi-continuous in the choice made. (A real-valued function  $f(x)$  is upper semi-continuous if  $x^* \rightarrow x$  implies  $\lim_{x \rightarrow x^*} f(x) \leq f(x^*)$ .)

Assuming that payoffs  $u_i$  are continuous in  $x$  does not guarantee that an optimal action exists at each node. Although the last mover's payoff is continuous and therefore an optimum exists if his action set is compact, the last mover's optimal action need not be a continuous function of the action chosen by the previous player. In this case, when we replace that last mover by an arbitrary specification of an optimal action on each path, the next-to-last mover's derived payoff function need not be upper semi-continuous, even though that player's payoff is a continuous function of the actions chosen at each node. Thus, the simple backward-induction algorithm defined above cannot be applied. However, subgame-perfect equilibria do exist in infinite-action games of perfect information, as shown by Harris (1985) and by Hellwig and Leininger (1987).

general games. In this section we will present a first look at "equilibrium refinements," which are designed to separate the "reasonable" Nash equilibria from the "unreasonable" ones. In particular, we will discuss the ideas of backward induction and "subgame perfection." Chapters 4, 5 and 13 apply these ideas to some classes of games of interest to economists.

Selten (1965) was the first to argue that in general extensive games some of the Nash equilibria are "more reasonable" than others. He began with the example illustrated here in figure 3.14. This is a finite game of perfect information, and the backward-induction solution (that is, the one obtained using Kuhn's algorithm) is that player 2 should play L if his information set is reached, and so player 1 should play D. Inspection of the strategic form corresponding to this game shows that there is another Nash equilibrium, where player 1 plays U and player 2 plays R. The profile (U, R) is a Nash equilibrium because, given that player 1 plays U, player 2's information set is not reached, and player 2 loses nothing by playing R. But Selten argued, and we agree, that this equilibrium is suspect. After all, if player 2's information set is reached, then, as long as player 2 is convinced that his payoffs are as specified in the figure, player 2 should play L. And if we were player 2, this is how we would play. Moreover, if we were player 1, we would expect player 2 to play L, and so we would play D.

In the now-familiar language, the equilibrium (U, R) is not "credible," because it relies on an "empty threat" by player 2 to play R. The threat is "empty" because player 2 would never wish to carry it out.

The idea that backward induction gives the right answer in simple games like that of figure 3.14 was implicit in the economics literature before Selten's paper. In particular, it is embodied in the idea of Stackelberg equilibrium: The requirement that player 2's strategy be the Cournot reaction function is exactly the idea of backward induction, and all other Nash equilibria of the game are inconsistent with backward induction. So we see that the expression "Stackelberg equilibrium" does not simply refer to the extensive form of the Stackelberg game, but instead is shorthand for "the backward-induction solution to the sequential quantity-choice game." Just as with "Cournot equilibrium," this shorthand terminology can be convenient when no confusion can arise. However, our experience suggests that the terminology can indeed lead to confusion, so we advise the student to use the more precise language instead.

Consider the game illustrated in figure 3.15. Here neither of player 2's choices is dominated at his last information set, and so backward induction does not apply. However, given that one accepts the logic of backward induction, the following argument seems compelling as well: "The game beginning at player 1's second information set is a zero-sum simultaneous-move game ('matching pennies') whose unique Nash equilibrium has expected payoffs (0, 0). Player 2 should choose R only if he expects that there is




Figure 3.15

probability  $\frac{1}{2}$  or better that he will outguess player 1 in the simultaneous-move subgame and end up with +2 instead of -2. Since player 2 assumes that player 1 is as rational as he is, it would be very rash of player 2 to expect to get the better of player 1, especially to such an extent. Thus, player 2 should go L, and so player 1 should go R." This is the logic of subgame perfection: Replace any "proper subgame" of the tree with one of its Nash-equilibrium payoffs, and perform backward induction on the reduced tree. (If the subgame has multiple Nash equilibria, this requires that all players agree on which of them would occur; we will come back to this point in subsection 3.6.1.) Once the subgame starting at player 1's second information set is replaced by its Nash-equilibrium outcome, the games described in figures 3.14 and 3.15 coincide.

To define subgame perfection formally we must first define the idea of a proper subgame. Informally, a proper subgame is a portion of a game that can be analyzed as a game in its own right, like the simultaneous-move game embedded in figure 3.15. The formal definition is not much more complicated:

**Definition 3.4** A proper subgame  $G$  of an extensive-form game  $T$  consists of a single node and all its successors in  $T$ , with the property that if  $x' \in G$  and  $x'' \in h(x')$  then  $x'' \in G$ . The information sets and payoffs of the subgame are inherited from the original game. That is,  $x'$  and  $x''$  are in the same information set in the subgame if and only if they are in the same information set in the original game, and the payoff function on the subgame is just the restriction of the original payoff function to the terminal nodes of the subgame.

Here the word "proper" here does not mean strict inclusion, as it does in the term "proper subset." Any game is always a proper subgame of itself. Proper subgames are particularly easy to identify in the class of deterministic multi-stage games with observed actions. In these games, all




Figure 3.16




Figure 3.17

previous actions are known to all players at the start of each stage, so each stage begins a new proper subgame. (Checking this is part of exercise 3.4.)

The requirements that all the successors of  $x$  be in the subgame and that the subgame not "chop up" any information set ensure that the subgame corresponds to a situation that could arise in the original game. In figure 3.16, the game on the right isn't a subgame of the game on the left, because on the right player 2 knows that player 1 didn't play L, which he did not know in the original game.

Together, the requirements that the subgame begin with a single node  $x$  and that the subgame respect information sets imply that in the original game  $x$  must be a singleton information set, i.e.,  $h(x) = \{x\}$ . This ensures that the payoffs in the subgame, conditional on the subgame being reached, are well defined. In figure 3.17, the "game" on the right has the problem that player 2's optimal choice depends on the relative probabilities of nodes  $x$  and  $x'$ , but the specification of the game does not provide these probabilities. In other words, the diagram on the right cannot be analyzed as an independent game; it makes sense only as a component of the game on the left, which is needed to provide the missing probabilities.

Since payoffs conditional on reaching a proper subgame are well defined, we can test whether strategies yield a Nash equilibrium when restricted to the subgame in the obvious way. That is, if  $\sigma_i$  is a behavior strategy for player  $i$  in the original game, and  $\hat{H}_i$  is the collection of player  $i$ 's information sets in the proper subgame, then the restriction of  $\sigma_i$  to the subgame is the map  $\hat{\sigma}_i$  such that  $\hat{\sigma}_i(\cdot|h_i) = \sigma_i(\cdot|h_i)$  for every  $h_i \in \hat{H}_i$ .

We have now developed the machinery needed to define subgame perfection.

**Definition 3.5** A behavior-strategy profile  $\sigma$  of an extensive-form game is a subgame-perfect equilibrium if the restriction of  $\sigma$  to  $G$  is a Nash equilibrium of  $G$  for every proper subgame  $G$ .



Because any game is a proper subgame of itself, a subgame-perfect equilibrium profile is necessarily a Nash equilibrium. If the only proper subgame is the whole game, the sets of Nash and subgame-perfect equilibria coincide. If there are other proper subgames, some Nash equilibria may fail to be subgame perfect.

It is easy to see that subgame perfection coincides with backward induction in finite games of perfect information. Consider the penultimate nodes of the tree, where the last choices are made. Each of these nodes begins a trivial one-player proper subgame, and Nash equilibrium in these subgames requires that the player now make a choice that maximizes his payoff; thus, any subgame-perfect equilibrium must coincide with a backward-induction solution at every penultimate node, and we can continue up the tree by induction. But subgame perfection is more general than backward induction; for example, it gives the suggested answer in the game of figure 3.15.

We remarked above that in multi-stage games with observed actions every stage begins a new proper subgame. Thus, in these games, subgame perfection is simply the requirement that the restrictions of the strategy profile yield a Nash equilibrium from the start of each stage  $k$  for each history  $h^k$ . If the game has a fixed finite number of stages ( $K + 1$ ), then we can characterize the subgame-perfect equilibria using backward induction: The strategies in the last stage must be a Nash equilibrium of the corresponding one-shot simultaneous-move game, and for each history  $h^k$  we replace the last stage by one of its Nash-equilibrium payoffs. For each such assignment of Nash equilibria to the last stage, we then consider the set of Nash equilibria beginning from each stage  $h^{k-1}$ . (With the last stage replaced by a payoff vector, the game from  $h^{k-1}$  on is a one-shot simultaneous-move game.) The characterization proceeds to "roll back the tree" in the manner of the Kuhn-Zermelo algorithm. Note that even if two different stage- $K$  histories lead to the "same game" in the last stage (that is, if there is a way of identifying strategies in the two games that preserves payoffs), the two histories still correspond to different subgames, and subgame perfection allows us to specify a different Nash equilibrium for each history. This has important consequences, as we will see in section 4.3 and in chapter 5.

### 3.6 Critiques of Backward Induction and Subgame Perfection<sup>††</sup>

This section discusses some of the limitations of the arguments for backwards induction and subgame perfection as necessary conditions for reasonable play. Although these concepts seem compelling in simple two-stage games of perfect information, such as the Stackelberg game we discussed at the start of the chapter, things are more complicated if there are many

players or if each player moves several times; in these games, equilibrium refinements are less compelling.

#### 3.6.1 Critiques of Backward Induction

Consider the  $I$ -player game illustrated in figure 3.18, where each player  $i < I$  can either end the game by playing "D" or play "A" and give the move to player  $i + 1$ . (To readers who skipped sections 3.3-3.5: Figure 3.18 depicts a "game tree." Though you have not seen a formal definition of such trees, we trust that the particular trees we use in this subsection will be clear.) If player  $i$  plays D, each player gets  $1/i$ ; if all players play A, each gets 2.

Since only one player moves at a time, this is a game of perfect information, and we can apply the backward-induction algorithm, which predicts that all players should play A. If  $I$  is small, this seems like a reasonable prediction. If  $I$  is very large, then, as player 1, we ourselves would play D and not A on the basis of a "robustness" argument similar to the one that suggested the inefficient equilibrium in the stag-hunt game of subsection 1.2.4.

First, the payoff 2 requires that all  $I - 1$  other players play A. If the probability that a given player plays A is  $p < 1$ , independent of the others, the probability that all  $I - 1$  other players play A is  $p^{I-1}$ , which can be quite small even if  $p$  is very large. Second, we would worry that player 2 might have these same concerns; that is, player 2 might play D to safeguard against either "mistakes" by future players or the possibility that player 3 might intentionally play D.

A related observation is that longer chains of backward induction presume longer chains of the hypothesis that "player 1 knows that player 2 knows that player 3 knows... the payoffs." If  $I = 2$  in figure 3.18, backward induction supposes that player 1 knows player 2's payoff, or at least that player 1 is fairly sure that player 2's optimal choice is A. If  $I = 3$ , not only must players 1 and 2 know player 3's payoff, in addition, player 1 must know that player 2 knows player 3's payoff, so that player 1 can forecast player 2's forecast of player 3's play. If player 1 thinks that player 2 will forecast player 3's play incorrectly, then player 1 may choose to play D. Traditionally, equilibrium analysis is motivated by the assumption that




Figure 3.18

payoffs are "common knowledge," so that arbitrarily long chains of "i knows that j knows that k knows" are valid, but conclusions that require very long chains of this form are less compelling than conclusions that require less of the power of the common-knowledge assumption. (In part this is because longer chains of backward induction are more sensitive to small changes in the information structure of the game, as we will see in chapter 9.)

The example in figure 3.18 is most troubling if  $I$  is very large. A second complication with backward induction arises whenever the same player can move several times in succession. Consider the game illustrated in figure 3.19. Here the backward-induction solution is that at every information set the player who has the move plays D. Is this solution compelling? Imagine that it is, that you are player 2, and that, contrary to expectation, player 1 plays  $A_1$  at his first move. How should you play? Backward induction says to play  $D_2$  because player 1 will choose  $D_3$  if given a chance, but backward induction also says that player 1 should have played  $D_1$ . In this game, unlike the simple examples we started with, player 2's best action if player 1 deviates from the predicted play  $A_1$  depends on how player 2 expects player 1 to play in the future: If player 2 thinks there is at least a 25 percent chance that player 1 will play  $A_3$ , then player 2 should play  $A_2$ . How should player 2 form these beliefs, and what beliefs are reasonable? In particular, how should player 2 predict how player 1 will play if, contrary to backward induction, player 1 decides to play  $A_1$ ? In some contexts, playing  $A_2$  may seem like a good gamble.

Most analyses of dynamic games in the economics literature continue to use backward induction and its refinements without reservations, but recently the skeptics have become more numerous. The game depicted in figure 3.19 is based on an example provided by Rosenthal (1981), who was one of the first to question the logic of backward induction. Basu (1988, 1990), Bonanno (1988), Binmore (1987, 1988), and Reny (1986) have argued that reasonable theories of play should not try to rule out any behavior once an event to which the theory assigns probability 0 has occurred, because the theory provides no way for players to form their predictions conditional on these events. Chapter 11 discusses the work of Fudenberg, Kreps, and Levine (1988), who propose that players interpret unexpected deviations as being due to the payoffs' differing from those that were



Figure 3.19

originally thought to be most likely. Since any observation of play can be explained by some specification of the opponents' payoffs, this approach sidesteps the difficulty of forming beliefs conditional on probability-0 events, and it recasts the question of how to predict play after a "deviation" as a question of which alternative payoffs are most likely given the observed play. Fudenberg and Kreps (1988) extend this to a methodological principle: They argue that any theory of play should be "complete" in the sense of assigning positive probability to any possible sequence of play, so that, using the theory, the players' conditional forecasts of subsequent play are always well defined.

Payoff uncertainty is not the only way to obtain a complete theory. A second family of complete theories is obtained by interpreting any extensive-form game as implicitly including the fact that players sometimes make small "mistakes" or "trembles" in the sense of Selten 1975. If, as Selten assumes, the probabilities of "trembling" at different information sets are independent, then no matter how often past play has failed to conform to the predictions of backward induction, a player is justified in continuing to use backward induction to predict play in the current subgame. Thus, interpreting "trembles" as deviations is a way to defend backward induction. The relevant question is how likely players view this "trembles" explanation of deviations as opposed to others. In figure 3.19, if player 2 observes  $A_1$ , should she (or will she) interpret this as a "tremble," or as a signal that player 1 is likely to play  $A_3$ ?

### 3.6.2 Critiques of Subgame Perfection

Since subgame perfection is an extension of backward induction, it is vulnerable to the critiques just discussed. Moreover, subgame perfection requires that players all agree on the play in a subgame even if that play cannot be predicted from backward-induction arguments. This point is emphasized by Rabin (1988), who proposes alternative, weaker equilibrium refinements that allow players to disagree about which Nash equilibrium will occur in a subgame off the equilibrium path.

To see the difference this makes, consider the following three-player game. In the first stage, player 1 can either play L, ending the game with payoffs (6, 0, 6), or play R, which gives the move to player 2. Player 2 can then either play R, ending the game with payoffs (8, 6, 8), or play L, in which case players 1 and 3 (but not player 2) play a simultaneous-move "coordination game" in which they each choose F or G. If their choices differ, they each receive 7 and player 2 gets 10; if the choices match, all three players receive 0. This game is depicted in figure 3.20.

The coordination game between players 1 and 3 at the third stage has three Nash equilibria: two in pure strategies with payoffs (7, 10, 7) and a mixed-strategy equilibrium with payoffs  $(3\frac{1}{2}, 5, 3\frac{1}{2})$ . If we specify an equilibrium in which players 1 and 3 successfully coordinate, then player 2 plays




Figure 3.20

L and so player 1 plays R, expecting a payoff of 7. If we specify the inefficient mixed equilibrium in the third stage, then player 2 will play R and again player 1 plays R, this time expecting a payoff of 8. Thus, in all subgame-perfect equilibria of this game, player 1 plays R.

As Rabin argues, it may nevertheless be reasonable for player 1 to play L. He would do so if he saw no way to coordinate in the third stage, and hence expected a payoff of  $3\frac{1}{2}$  conditional on that stage being reached, but feared that player 2 would believe that play in the third stage would result in coordination on an efficient equilibrium.

The point is that subgame perfection supposes not only that the players expect Nash equilibria in all subgames but also that all players expect the same equilibria. Whether this is plausible depends on the reason one thinks an equilibrium might arise in the first place.

Exercises

**Exercise 3.1\*** Players 1 and 2 must decide whether or not to carry an umbrella when leaving home. They know that there is a 50-50 chance of rain. Each player's payoff is  $-5$  if he doesn't carry an umbrella and it rains,  $-2$  if he carries an umbrella and it rains,  $-1$  if he carries an umbrella and it is sunny, and  $1$  if he doesn't carry an umbrella and it is sunny. Player 1 learns the weather before leaving home; player 2 does not, but he can observe player 1's action before choosing his own. Give the extensive and strategic forms of the game. Is it dominance solvable?

**Exercise 3.2.\*** Verify that the game in figure 3.13 does not meet the formal definition of a game of perfect recall.

**Exercise 3.3\*** Player 1, the "government," wishes to influence the choice of player 2. Player 2 chooses an action  $a_2 \in A_2 = \{0, 1\}$  and receives a

transfer  $t \in T = \{0, 1\}$  from the government, which observes  $a_2$ . Player 2's objective is to maximize the expected value of his transfer, minus the cost of his action, which is 0 for  $a_2 = 0$  and  $\frac{1}{2}$  for  $a_2 = 1$ . Player 1's objective is to minimize the sum  $2(a_2 - 1)^2 + t$ . Before player 2 chooses his action, the government can announce a transfer rule  $t(a_2)$ .

- (a) Draw the extensive form for the case where the government's announcement is not binding and has no effect on payoffs.
- (b) Draw the extensive form for the case where the government is constrained to implement the transfer rule it announced.
- (c) Give the strategic forms for both games.
- (d) Characterize the subgame-perfect equilibria of the two games.

**Exercise 3.4\*\*** Define a deterministic multi-stage game with observed actions using conditions on the information sets of an extensive form. Show that in these games the start of each stage begins a proper subgame.

**Exercise 3.5\*\*** Show that subgame-perfect equilibria exist in finite multi-stage games.

**Exercise 3.6\*** There are two players, a seller and a buyer, and two dates. At date 1, the seller chooses his investment level  $I \geq 0$  at cost  $I$ . At date 2, the seller may sell one unit of a good and the seller has cost  $c(I)$  of supplying it, where  $c'(0) = -\infty$ ,  $c' < 0$ ,  $c'' > 0$ , and  $c(0)$  is less than the buyer's valuation. There is no discounting, so the socially optimal level of investment,  $I^*$ , is given by  $1 + c'(I^*) = 0$ .

(a) Suppose that at date 2 the buyer observes the investment  $I$  and makes a take-it-or-leave-it offer to the seller. What is this offer? What is the perfect equilibrium of the game?

(b) Can you think of a contractual way of avoiding the inefficient outcome of (a)? (Assume that contracts cannot be written on the level of  $I$ .)

**Exercise 3.7\*** Consider a voting game in which three players, 1, 2, and 3, are deciding among three alternatives, A, B, and C. Alternative B is the "status quo" and alternatives A and C are "challengers." At the first stage, players choose which of the two challengers should be considered by casting votes for either A or C, with the majority choice being the winner and abstentions not allowed. At the second stage, players vote between the status quo B and whichever alternative was victorious in the first round, with majority rule again determining the winner. Players vote simultaneously in each round. The players care only about the alternative that is finally selected, and are indifferent as to the sequence of votes that leads to a given selection. The payoff functions are  $u_1(A) = 2, u_1(B) = 0, u_1(C) = 1; u_2(A) = 1, u_2(B) = 2, u_2(C) = 0; u_3(A) = 0, u_3(B) = 1, u_3(C) = 2$ .

(a) What would happen if at each stage the players voted for the alternative they would most prefer as the final outcome?

(b) Find the subgame-perfect equilibrium outcome that satisfies the additional condition that no strategy can be eliminated by iterated weak dominance. Indicate what happens if dominated strategies are allowed.

(c) Discuss whether different "agendas" for arriving at a final decision by voting between two alternatives at a time would lead to a different equilibrium outcome.

(This exercise is based on Eckel and Holt 1989, in which the play of this game in experiments is reported.)

**Exercise 3.8\*** Subsection 3.2.3 discussed a player's "strategic incentive" to alter his first-period actions in order to change his own second-period incentives and thus alter the second-period equilibrium. A player may also have a strategic incentive to alter the second-period incentives of others. One application of this idea is the literature on strategic trade policies (e.g. Brander and Spencer 1985; Eaton and Grossman 1986—see Helpman and Krugman 1989, chapters 5 and 6, for a clear review of the arguments). Consider two countries, A and B, and a single good which is consumed only in country B. The inverse demand function is  $p = P(Q)$ , where  $Q$  is the total output produced by firms in countries A and B. Let  $c$  denote the constant marginal cost of production and  $Q_m$  the monopoly output ( $Q_m$  maximizes  $Q(P(Q) - c)$ ).

(a) Suppose that country B does not produce the good. The  $I (\geq 1)$  firms in country A are Cournot competitors. Find conditions under which an optimal policy for the government of country A is to levy a unit export tax equal to  $-P'(Q_m)(I - 1)Q_m/I$ . (The objective of country A's government is to maximize the sum of its own receipts and the profit of its firm.) Give an externality interpretation.

(b) Suppose now that there are two producers, one in each country. The game has two periods. In period 1, the government of country A chooses an export tax or subsidy (per unit of exports); in period 2, the two firms, which have observed the government's choice, simultaneously choose quantities. Suppose that the Cournot reaction curves are downward sloping and intersect only once, at a point at which country A's firm's reaction curve is steeper than country B's firm's reaction curve in the  $(q_A, q_B)$  space. Show that an export *subsidy* is optimal.

(c) What would happen in question (b) if there were more than one firm in country A? If the strategic variables of period 2 gave rise to upward-sloping reaction curves? Caution: The answer to the latter depends on a "stability condition" of the kind discussed in subsection 1.2.5.

**Exercise 3.9\*\*** Consider the three-player extensive-form game depicted in figure 3.21.

(a) Show that (A, A) is not the outcome of a Nash equilibrium.

(b) Consider the nonequilibrium situation where player 1 expects player 3 to play R, player 2 expects player 3 to play L, and consequently players




Figure 3.21

1 and 2 both play A. When might this be a fixed point of a learning process like those discussed in chapter 1? When might learning be expected to lead players 1 and 2 to have the same beliefs about player 3's action, as required for Nash equilibrium? (Give an informal answer.) For more on this question see Fudenberg and Kreps 1988 and Fudenberg and Levine 1990.

**Exercise 3.10\*\*\*** In the class of zero-sum games, the sets of outcomes of Nash and subgame-perfect equilibria are the same. That is, for every outcome (probability distribution over terminal nodes) of a Nash-equilibrium strategy profile, there is a perfect equilibrium profile with the same outcome. This result has limited interest, because most games in the social sciences are not zero-sum; however, its proof, which we give in the context of a multi-stage game with observed actions, is a nice way to get acquainted with the logic of perfect equilibrium. Consider a two-person game and let  $u_1(\sigma_1, \sigma_2)$  denote player 1's expected payoff (by definition of a zero-sum game,  $u_2 = -u_1$ ). Let  $u_1(\sigma_1, \sigma_2 | h')$  denote player 1's expected payoff conditional on history  $h'$  having been reached at date  $t$  (for simplicity, we identify "stages" with "dates"). Last, let  $\sigma_i / \delta_i^{h'}$  denote player  $i$ 's strategy  $\sigma_i$ , except that if  $h'$  is reached at date  $t$ , player  $i$  adopts strategy  $\delta_i^{h'}$  in the subgame associated with history  $h'$  (henceforth called "the subgame").

(a) Let  $(\sigma_1, \sigma_2)$  be a Nash equilibrium. If  $(\sigma_1, \sigma_2)$  is not perfect, there is a date  $t$ , a history  $h'$ , and a player (say player 1) such that this player does not maximize his payoff conditional on history  $h'$  being reached. (Of course, this history  $h'$  must have probability 0 of being reached according to strategies  $(\sigma_1, \sigma_2)$ ; otherwise player 1 will not be maximizing his unconditional payoff  $u_1(\sigma_1, \sigma_2)$  given  $\sigma_2$ .)

Let  $\delta_1^{h'}$  denote the strategy that maximizes  $u_1(\sigma_1 / \delta_1^{h'}, \sigma_2 | h')$ . Last, let  $(\sigma_1^{*h'}, \sigma_2^{*h'})$  denote a Nash equilibrium of the subgame. Show that for any  $\bar{\sigma}_1$

$$u_1(\bar{\sigma}_1 / \delta_1^{h'}, \sigma_2 | h') \geq u_1(\bar{\sigma}_1 / \sigma_1^{*h'}, \sigma_2 / \sigma_2^{*h'} | h').$$

(Hint: Use the facts that  $\sigma_2^{*h'}$  is a best response to  $\sigma_1^{*h'}$  in subgame  $h'$ , that the game is a zero-sum game, and that  $\delta_1^{h'}$  is an optimal response to  $\sigma_2$  in the subgame.)




Figure 3.22

(b) Show that the strategy profile  $(\sigma_1/\sigma_1^{*M}, \sigma_2/\sigma_2^{*M})$  is also a Nash equilibrium. (Hint: Use the fact that subgame  $h'$  is not reached under  $(\sigma_1, \sigma_2)$  and the definition of Nash behavior in the subgame.)

(c) Conclude that the Nash-equilibrium outcome (the probability distribution on terminal nodes generated by  $\sigma_1$  and  $\sigma_2$ ) is also a perfect-equilibrium outcome.

Note that although outcomes coincide, the Nash-equilibrium strategies need not be perfect-equilibrium strategies—as is demonstrated in figure 3.22, where  $(R_1, R_2)$  is a Nash, but not a perfect, equilibrium.

**Exercise 3.11\*** Consider the agenda-setter model of Romer and Rosenthal (1978) (see also Shepsle 1981). The object of the game is to make a one-dimensional decision. There are two players. The “agenda-setter” (player 1, who may stand for a committee in a closed-rule voting system) offers a point  $s_1 \in \mathbb{R}$ . The “voter” (player 2, who may stand for the median voter in the legislature) can then accept  $s_1$  or refuse it; in the latter case, the decision is the status quo or reversion point  $s_0$ . Thus,  $s_2 \in \{s_0, s_1\}$ . The adopted policy is thus  $s_2$ . The voter has quadratic preferences:  $-(s_2 - \delta_2)^2$ , where  $\delta_2$  is his bliss point.

(a) Suppose that the agenda setter’s objective is  $s_2$  (she prefers higher policy levels). Show that, in perfect equilibrium, the setter offers  $s_1 = s_0$  if  $s_0 \geq \delta_2$  and  $s_1 = 2\delta_2 - s_0$  if  $s_0 < \delta_2$ .

(b) Suppose that the agenda setter’s objective function is quadratic as well:  $-(s_2 - \delta_1)^2$ . Fixing  $\delta_1$  and  $\delta_2$  ( $\delta_1 \geq \delta_2$ ), depict how the perfect-equilibrium policy varies with the reversion  $s_0$ .

**Exercise 3.12\*\*** Consider the twice-repeated version of the agenda-setter model developed in the previous exercise. The new status quo in period 2 is whatever policy (agenda setter’s proposal or initial status quo) was adopted in period 1. Suppose that the objective function of the agenda setter is the sum of the two periods’ policies, and that the voter’s preferences are  $-(s_1^2 - 4)^2 - (s_2^2 - 12)^2$  (that is, his bliss point is 4 for the first-period policy and 12 for the second-period one). The initial status quo is 2.

(a) Suppose first that the voter is myopic (acts as if his discount factor were 0 instead of 1), but that the agenda setter is not. Show that the agenda

setter offers 6 in period 1, and that the payoffs are  $-24$  for the agenda setter and  $-40$  for the voter. Assume in this exercise that the voter chooses the higher acceptable policy when indifferent. If you are courageous, show that this policy is uniquely optimal when the agenda setter’s discount factor is slightly less than 1 instead of 1.

(b) Suppose now that both players are rational. Show that the agenda setter’s utility is higher and the voter’s utility is lower than in question (a). What point does this comparison illustrate? (See Ingberman 1985 and Rosenthal 1990.)

## References

- Basu, K. 1988. Strategic irrationality in extensive games. *Mathematical Social Sciences* 15: 247–260.
- Basu, K. 1990. On the non-existence of a rationality definition for extensive games. *International Journal of Game Theory*.
- Binmore, K. 1987. Modeling rational players: Part I. *Economics and Philosophy* 3: 179–214.
- Binmore, K. 1988. Modeling rational players: Part II. *Economics and Philosophy* 4: 9–55.
- Bonanno, G. 1988. The logic of rational play in extensive games of perfect information. Mimeo, University of California, Davis.
- Brander, J., and B. Spencer. 1985. Export subsidies and market share rivalry. *Journal of International Economics* 18: 83–100.
- Dixit, A. 1979. A model of duopoly suggesting a theory of entry barriers. *Bell Journal of Economics* 10: 20–32.
- Eaton, J., and G. Grossman. 1986. Optimal trade and industrial policy under oligopoly. *Quarterly Journal of Economics* 101: 383–406.
- Eckel, C., and C. Holt. 1989. Strategic voting in agenda-controlled committee experiments. *American Economic Review* 79: 763–773.
- Elmes, S., and P. Reny. 1988. The equivalence of games with perfect recall. Mimeo.
- Fudenberg, D., and D. Kreps. 1988. A theory of learning, experimentation, and equilibrium in games. Mimeo, MIT.
- Fudenberg, D., D. Kreps, and D. Levine. 1988. On the robustness of equilibrium refinements. *Journal of Economic Theory* 44: 354–380.
- Fudenberg, D., and D. Levine. 1990. Steady-state learning and self-confirming equilibrium. Mimeo.
- Harris, C. 1985. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica* 53: 613–627.
- Hellwig, M., and W. Leininger. 1987. On the existence of subgame-perfect equilibrium in infinite-action games of perfect information. *Journal of Economic Theory*.
- Helpman, E., and P. Krugman. 1989. *Trade Policy and Market Structure*. MIT Press.
- Ingberman, D. 1985. Running against the status-quo. *Public Choice* 146: 14–44.
- Kreps, D., and R. Wilson. 1982. Sequential equilibria. *Econometrica* 50: 863–894.
- Kuhn, H. 1953. Extensive games and the problem of information. *Annals of Mathematics Studies*, no. 28. Princeton University Press.
- Kydland, F., and E. Prescott. 1977. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy* 85: 473–491.
- Luce, R., and H. Raiffa. 1957. *Games and Decisions*. Wiley.

- Mankiw, G. 1988. Recent developments in macroeconomics: A very quick refresher course. *Journal of Money, Credit, and Banking* 20: 436-459.
- Rabin, M. 1988. Consistency and robustness criteria for game theory. Mimeo, MIT.
- Reny, P. 1986. Rationality, common knowledge, and the theory of games. Ph.D. Dissertation, Princeton University.
- Romer, T., and H. Rosenthal. 1978. Political resource allocation, controlled agendas, and the status-quo. *Public Choice* 33: 27-44.
- Rosenthal, H. 1990. The setter model. In *Readings in the Spatial Theory of Elections*, ed. Enelow and Hinich. Cambridge University Press.
- Rosenthal, R. 1981. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory* 25: 92-100.
- Schelling, T. 1960. *The Strategy of Conflict*. Harvard University Press.
- Selten, R. 1965. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträglichkeit. *Zeitschrift für die gesamte Staatswissenschaft* 12: 301-324.
- Selten, R. 1975. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4: 25-55.
- Shepsle, K. 1981. Structure-induced equilibrium and legislative choice. *Public Choice* 37: 503-520.
- Spence, A. M. 1977. Entry, capacity, investment and oligopolistic pricing. *Bell Journal of Economics* 8: 534-544.
- Zermelo, E. 1913. Über eine Anwendung der Mengenlehre auf der Theorie des Schachspiels. In *Proceedings of the Fifth International Congress on Mathematics*.

4.1 Introduction<sup>†</sup>

In chapter 3 we introduced a class of extensive-form games that we called "multi-stage games with observed actions," where the players move simultaneously within each stage and know the actions that were chosen in all past stages. Although these games are very special, they have been used in many applications in economics, political science, and biology. The repeated games we study in chapter 5 belong to this class, as do the games of resource extraction, preemptive investment, and strategic bequests discussed in chapter 13. This chapter develops a basic fact about dynamic optimization and presents a few interesting examples of multi-stage games. The chapter concludes with discussions of what is meant by "open-loop" and "closed-loop" equilibria, of the notion of iterated conditional dominance, and of the relationship between equilibria of finite-horizon and infinite-horizon games.

Recall that in a multi-stage game with observed actions the history  $h^t$  at the beginning of stage  $t$  is simply the sequence of actions  $(a^0, a^1, \dots, a^{t-1})$  chosen in previous periods, and that a pure strategy  $s_i$  for player  $i$  is a sequence of maps  $s_i^t$  from histories  $h^t$  to actions  $a_i^t$  in the feasible sets  $A_i(h^t)$ . Player  $i$ 's payoff  $u_i$  is a function of the terminal history  $h^{T+1}$ , i.e., of the entire sequence of actions from the initial stage 0 through the terminal stage  $T$ , where  $T$  is sometimes taken to be infinite. In some of the examples of this section, payoffs take the special form of the discounted sum  $\sum_{t=0}^T \delta^t g_i^t(a^t)$  of per-period payoffs  $g_i^t(a^t)$ .

Section 4.3 presents a first look at the subclass of repeated games, where the payoffs are given by averages as above and where the sets of feasible actions at each stage and the per-period payoffs are independent of previous play and time, so that the "physical environment" of the game is memoryless. Nevertheless, the fact that the game is repeated means that the players can condition their current play on the past play of their opponents, and indeed there can be equilibria in strategies of this kind. Section 4.3 considers only a few examples of repeated games, and does not try to characterize all the equilibria of the examples it examines; chapter 5 gives a more thorough treatment.

In this chapter we consider mostly games with an infinite horizon as opposed to a horizon that is long but finite. Games with a long but finite horizon represent a situation where the horizon is long but well foreseen; infinite-horizon games describe a situation where players are fairly uncertain as to which period will be the last. This latter assumption seems to be a better model of many situations with a large number of stages; we will say more about this point when discussing some of the examples.

When the horizon is infinite, the set of subgame-perfect equilibria cannot be determined by backward induction from the terminal date, as it can