

Basic Elements of Noncooperative Games

A Introduction

In this chapter, we begin our study of noncooperative game theory by introducing some of its basic building blocks. This material serves as a prelude to our analysis of games in Chapters 8 and 9.

Section 7.B begins with an informal introduction to the concept of a *game*. It describes the four basic elements of any setting of strategic interaction that we must know to specify a game.

In Section 7.C, we show how a game can be described by means of what is called its *extensive form representation*. The extensive form representation provides a very rich description of a game, capturing who moves when, what they can do, what they know when it is their turn to move, and the outcomes associated with any collection of actions taken by the individuals playing the game.

In Section 7.D, we introduce a central concept of game theory, a player's *strategy*. A player's strategy is a complete contingent plan describing the actions she will take in each conceivable evolution of the game. We then show how the notion of a strategy can be used to derive a much more compact representation of a game, known as its *normal (or strategic) form representation*.

In Section 7.E, we consider the possibility that a player might randomize her choices. This gives rise to the notion of a *mixed strategy*.

What Is a Game?

A *game* is a formal representation of a situation in which a number of individuals interact in a setting of *strategic interdependence*. By that, we mean that each individual's welfare depends not only on her own actions but also on the actions of the other individuals. Moreover, the actions that are best for her to take may depend on what she expects the other players to do.

To describe a situation of strategic interaction, we need to know four things:

- (i) *The players:* Who is involved?
- (ii) *The rules:* Who moves when? What do they know when they move?
What can they do?

large set of pure strategies in S_i , she could randomize separately over the possible actions at each of her information sets $H \in \mathcal{H}_i$. This way of randomizing is called a *behavior strategy*.

Definition 7.E.2: Given an extensive form game Γ_E , a *behavior strategy* for player i specifies, for every information set $H \in \mathcal{H}_i$ and action $a \in C(H)$, a probability $\lambda_i(a, H) \geq 0$, with $\sum_{a \in C(H)} \lambda_i(a, H) = 1$ for all $H \in \mathcal{H}_i$.

As might seem intuitive, for games of perfect recall (and we deal only with these), the two types of randomization are equivalent. For any behavior strategy of player i , there is a mixed strategy for that player that yields exactly the same distribution over outcomes for any strategies, mixed or behavior, that might be played by i 's rivals, and vice versa [this result is due to Kuhn (1953); see Exercise 7.E.1]. Which form of randomized strategy we consider is therefore a matter of analytical convenience; we typically use behavior strategies when analyzing the extensive form representation of a game and mixed strategies when analyzing the normal form.

Because the way we introduce randomization is solely a matter of analytical convenience, we shall be a bit loose in our terminology and refer to all randomized strategies as *mixed strategies*.

REFERENCES

- Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics* 4: 1236–39.
 Kreps, D. M., and R. Wilson. (1982). Sequential equilibrium. *Econometrica* 50: 863–94.
 Kuhn, H. W. (1953). Extensive games and the problem of information. In *Contributions to the Theory of Games*, vol. 2, edited by H. W. Kuhn and A. W. Tucker. Princeton, N.J.: Princeton University Press, 193–216.
 Milgrom, P. (1981). An axiomatic characterization of common knowledge. *Econometrica* 49: 219–22.

EXERCISES

7.C.1^A Suppose that in the Meeting in New York game (Example 7.B.3), there are two possible places where the two players can meet: Grand Central Station and the Empire State Building. Draw an extensive form representation (game tree) for this game.

7.D.1^B In a game where player i has N information sets indexed $n = 1, \dots, N$ and M_n possible actions at information set n , how many strategies does player i have?

7.D.2^A In text.

7.E.1^B Consider the two-player game whose extensive form representation (excluding payoffs) is depicted in Figure 7.Ex.1.

- (a) What are player 1's possible strategies? Player 2's?
 (b) Show that for any behavior strategy that player 1 might play, there is a realization equivalent mixed strategy; that is, a mixed strategy that generates the same probability distribution over the terminal nodes for any mixed strategy choice by player 2.




Figure 7E

(c) Show that the converse is also true: For any mixed strategy that player 1 might play, there is a realization equivalent behavior strategy.

(d) Suppose that we change the game by merging the information sets at player 1's second round of moves (so that all four nodes are now in a single information set). Argue that the game is no longer one of perfect recall. Which of the two results in (b) and (c) still holds?

- (iii) *The outcomes:* For each possible set of actions by the players, what is the outcome of the game?
- (iv) *The payoffs:* What are the players' preferences (i.e., utility functions) over the possible outcomes?

We begin by considering items (i) to (iii). A simple example is provided by the school-yard game of *Matching Pennies*.

Example 7.B.1: *Matching Pennies.* Items (i) to (iii) are as follows:

- Players:* There are two players, denoted 1 and 2.
- Rules:* Each player simultaneously puts a penny down, either heads up or tails up.
- Outcomes:* If the two pennies match (either both heads up or both tails up), player 1 pays 1 dollar to player 2; otherwise, player 2 pays 1 dollar to player 1. ■

Consider another example, the game of *Tick-Tack-Toe*.

Example 7.B.2: *Tick-Tack-Toe.* Items (i) to (iii) are as follows:

- Players:* There are two players, X and O.
- Rules:* The players are faced with a board that consists of nine squares arrayed with three rows of three squares each stacked on one another (see Figure 7.B.1). The players take turns putting their marks (an X or an O) into an as-yet-unmarked square. Player X moves first. Both players observe all choices previously made.
- Outcomes:* The first player to have three of her marks in a row (horizontally, vertically, or diagonally) wins and receives 1 dollar from the other player. If no one succeeds in doing so after all nine boxes are marked, the game is a tie and no payments are made or received by either player. ■

To complete our description of these two games, we need to say what the players' preferences are over the possible outcomes [item (iv) in our list]. As a general matter, we describe a player's preferences by a utility function that assigns a utility level for each possible outcome. It is common to refer to the player's utility function as her *payoff function* and the utility level as her *payoff*. Throughout, we assume that these utility functions take an expected utility form (see Chapter 6) so that when we consider situations in which outcomes are random, we can evaluate the random prospect by means of the player's expected utility.




FIGURE 7.B.1
A 3x3 board

In later references to Matching Pennies and Tick-Tack-Toe, we assume that each player's payoff is simply equal to the amount of money she gains or loses. Note that in both examples, the actions that maximize a player's payoff depend on what she expects her opponent to do.

Examples 7.B.1 and 7.B.2 involve situations of pure conflict: What one player wins, the other player loses. Such games are called *zero-sum games*. But strategic interaction and game theory are not limited to situations of pure or even partial conflict. Consider the situation in Example 7.B.3.

Example 7.B.3: *Meeting in New York*. Items (i) to (iv) are as follows:

Players: Two players, Mr. Thomas and Mr. Schelling.

Rules: The two players are separated and cannot communicate. They are supposed to meet in New York City at noon for lunch but have forgotten to specify where. Each must decide where to go (each can make only one choice).

Outcomes: If they meet each other, they get to enjoy each other's company at lunch. Otherwise, they must eat alone.

Payoffs: They each attach a monetary value of 100 dollars to the other's company (their payoffs are each 100 dollars if they meet, 0 dollars if they do not).

In this example, the two players' interests are completely aligned. Their problem is simply one of coordination. Nevertheless, each player's payoff depends on what the other player does; and more importantly, *each player's optimal action depends on what he thinks the other will do*. Thus, even the task of coordination can have a strategic nature. ■

Although the information given in items (i) to (iv) fully describe a game, it is useful for purposes of analysis to represent this information in particular ways. We examine one of these ways in Section 7.C.

C The Extensive Form Representation of a Game

If we know the items (i) to (iv) described in Section 7.B (the players, the rules, the outcomes, and the payoffs), then we can formally represent the game in what is called its *extensive form*. The extensive form captures who moves when, what actions each player can take, what players know when they move, what the outcome is as a function of the actions taken by the players, and the players' payoffs from each possible outcome.

We begin by informally introducing the elements of the extensive form representation through a series of examples. After doing so, we then provide a formal specification of the extensive form (some readers may want to begin with this and then return to the examples).

The extensive form relies on the conceptual apparatus known as a *game tree*. As our starting point, it is useful to begin with a very simple variation of Matching Pennies, which we call *Matching Pennies Version B*.

Example 7.C.1: *Matching Pennies Version B and Its Extensive Form*. Matching Pennies Version B is identical to Matching Pennies (see Example 7.B.1) except

Figure 7.B.1
A Tick-Tack-Toe
board.




Figure
Extensive
Match
Vers

that the two players move sequentially, rather than simultaneously. In particular, player 1 puts her penny down (heads up or tails up) first. Then, after seeing player 1's choice, player 2 puts her penny down. (This is a very nice game for player 2!)

The extensive form representation of this game is depicted in Figure 7.C.1. The game starts at an *initial decision node* (represented by an open circle), where player 1 makes her move, deciding whether to place her penny heads up or tails up. Each of the two possible choices for player 1 is represented by a *branch* from this initial decision node. At the end of each branch is another decision node (represented by a solid dot), at which player 2 can choose between two actions, heads up or tails up, after seeing player 1's choice. The initial decision node is referred to as *player 1's decision node*; the latter two as *player 2's decision nodes*. After player 2's move, we reach the end of the game, represented by *terminal nodes*. At each terminal node, we list the players' payoffs arising from the sequence of moves leading to that terminal node.

Note the treelike structure of Figure 7.C.1: Like an actual tree, it has a unique connected path of branches from the initial node (sometimes also called the *root*) to each point in the tree. This type of figure is known as a *game tree*. ■

Example 7.C.2: The Extensive Form of Tick-Tack-Toe. The more elaborate game tree shown in Figure 7.C.2 depicts the extensive form for Tick-Tack-Toe (to conserve space, many parts are omitted). Note that every path through the tree represents a unique sequence of moves by the players. In particular, when a given board position (such as the two left corners filled by X and the two right corners filled by O) can be reached through several different sequences of moves, each of these sequences is depicted separately in the game tree. Nodes represent not only the current position but also *how it was reached*. ■

In both Matching Pennies Version B and Tick-Tack-Toe, when it is a player's turn to move, she is able to observe all her rival's previous moves. They are games of *perfect information* (we give a precise definition of this term in Definition 7.C.1). The concept of an *information set* allows us to accommodate the possibility that this is not so. Formally, the elements of an information set are a subset of a particular player's decision nodes. The interpretation is that when play has reached one of the decision nodes in the information set and it is that player's turn to move, she does

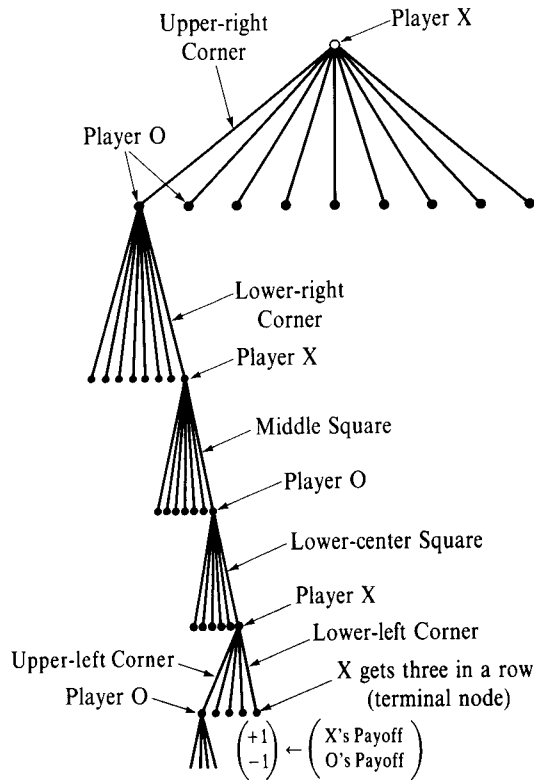


Figure 7.C.2
Part of the extensive form for Tick-Tack-Toe.

not know which of these nodes she is actually at. The reason for this ignorance is that the player does not observe something about what has previously transpired in the game. A further variation of Matching Pennies, which we call *Matching Pennies Version C*, helps make this concept clearer.

Example 7.C.3: Matching Pennies Version C and Its Extensive Form. This version of Matching Pennies is just like Matching Pennies Version B (in Example 7.C.1) except that when player 1 puts her penny down, she keeps it covered with her hand. Hence, player 2 cannot see player 1's choice until after player 2 has moved.

The extensive form for this game is represented in Figure 7.C.3. It is identical to Figure 7.C.1 except that we have drawn a circle around player 2's two decision nodes to indicate that these two nodes are in a single information set. The meaning of this information set is that when it is player 2's turn to move, she cannot tell which of these two nodes she is at because she has not observed player 1's previous move. Note that player 2 has the same two possible actions at each of the two nodes in her information set. This must be the case if player 2 is unable to distinguish the two nodes; otherwise, she could figure out which move player 1 had taken simply by what her own possible actions are.

In principle, we could also associate player 1's decision node with an information set. Because player 1 knows that nothing has happened before it is her turn to move, this information set has only one member (player 1 knows exactly which node she is at when she moves). To be fully rigorous, we should therefore also draw an information set circle around player 1's decision node in Figure 7.C.3. It is common, however, to




Figure 7.C.3
Extensive form
Matching Pennies
Version C.

simplify the diagrammatic depiction of a game in extensive form by not drawing the information sets that contain a single node. Thus, any uncircled decision nodes are understood to be elements of *singleton* information sets. In Figures 7.C.1 and 7.C.2, for example, every decision node belongs to a singleton information set. ■

A listing of all of a player's information sets gives a listing, from the player's perspective, of all of the possible distinguishable "events" or "circumstances" in which she might be called upon to move. For example, in Example 7.C.1, from player 2's perspective there are two distinguishable events that might arise in which she would be called upon to move, each one corresponding to play having reached one of her two (singleton) information sets. By way of contrast, player 2 foresees only one possible circumstance in which she would need to move in Example 7.C.3 (this circumstance is, however, certain to arise).

In Example 7.C.3, we noted a natural restriction on information sets: At every node within a given information set, a player must have the same set of possible actions. Another restriction we impose is that players possess what is known as *perfect recall*. Loosely speaking, perfect recall means that a player does not forget what she once knew, including her own actions. Figure 7.C.4 depicts two games in which this condition is not met. In Figure 7.C.4(a), as the game progresses, player 2 forgets a move by player 1 that she once knew (namely, whether player 1 chose ℓ or r). In Figure 7.C.4(b), player 1 forgets her own previous move.¹ All the games we consider in this book satisfy the property of perfect recall.

The use of information sets also allows us to capture play that is simultaneous rather than sequential. This is illustrated in Example 7.C.4 for the game of (standard) Matching Pennies introduced in Example 7.B.1.

1. In terms of the formal specification of the extensive form given later in this section, if we denote the information set containing decision node x by $H(x)$, a game is formally characterized as one of perfect recall if the following two conditions hold: (i) If $H(x) = H(x')$, x is neither a predecessor nor a successor of x' ; and (ii) if x and x' are two decision nodes for player i with $H(x) = H(x')$, and if x'' is a predecessor of x (not necessarily an immediate one) that is also in one of player i 's information sets, with a'' being the action at $H(x'')$ on the path to x , then there must be a predecessor node to x' that is an element of $H(x'')$ and the action at this predecessor node that is on the path to x' must also be a'' .

Figure 7.C.3
Extensive form for
Matching Pennies
Version C.

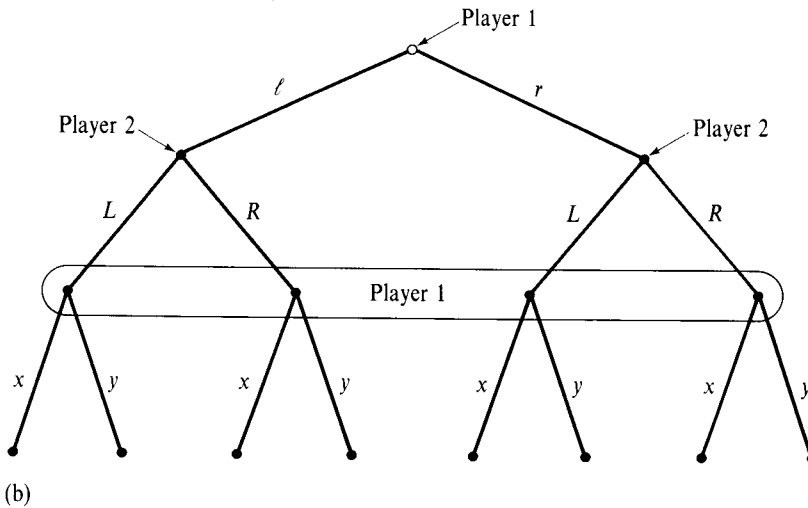



Figure 7.C.4
Two games not
satisfying perfect recall.

Example 7.C.4: *The Extensive Form for Matching Pennies.* Suppose now that the players put their pennies down simultaneously. For each player, this game is strategically equivalent to the Version C game. In Version C, player 1 was unable to observe player 2's choice because player 1 moved first, and player 2 was unable to observe player 1's choice because player 1 kept it covered; here each player is unable to observe the other's choice because they move simultaneously. As long as they cannot observe each other's choices, the timing of moves is irrelevant. Thus, we can use the game tree in Figure 7.C.3 to describe the game of (standard) Matching Pennies. Note that by this logic we can also describe this game with a game tree that reverses the decision nodes of players 1 and 2 in Figure 7.C.3. ■

We can now return to the notion of a game of perfect information and offer a formal definition.




Figure 7.C.5 Extensive form for Matching Pennies Version D.

Definition 7.C.1: A game is one of *perfect information* if each information set contains a single decision node. Otherwise, it is a game of *imperfect information*.

Up to this point, the outcome of a game has been a deterministic function of the players' choices. In many games, however, there is an element of chance. This, too, can be captured in the extensive form representation by including *random moves of nature*. We illustrate this point with still another variation, *Matching Pennies Version D*.

Example 7.C.5: Matching Pennies Version D and Its Extensive Form. Suppose that prior to playing Matching Pennies Version B, the two players flip a coin to see who will move first. Thus, with equal probability either player 1 will put her penny down first, or player 2 will. In Figure 7.C.5, this game is depicted as beginning with a *move of nature* at the initial node that has two branches, each with probability $\frac{1}{2}$. Note that this is drawn as if nature were an additional player who must play its two actions with fixed probabilities. (In the figure, H stands for "heads up" and T stands for "tails up".) ■

It is a basic postulate of game theory that all players know the structure of the game, know that their rivals know it, know that their rivals know that they know it, and so on. In theoretical parlance, we say that the structure of the game is *common knowledge* [see Aumann (1976) and Milgrom (1981) for discussions of this concept].

In addition to being depicted graphically, the extensive form can be described mathematically. The basic components are fairly easily explained and can help you keep in mind the fundamental building blocks of a game. Formally, a game represented in extensive form consists of the following items:²

2. To be a bit more precise about terminology: A collection of items (i) to (vi) is formally known as an extensive *game form*; adding item (vii), the players' preferences over the outcomes, leads to a *game* represented in extensive form. We will not make anything of this distinction here. See Kuhn (1953) or Section 2 of Kreps and Wilson (1982) for additional discussion of this and other points regarding the extensive form.

- (i) A finite set of nodes \mathcal{X} , a finite set of possible actions \mathcal{A} , and a finite set of players $\{1, \dots, I\}$.
- (ii) A function $p: \mathcal{X} \rightarrow \{\mathcal{X} \cup \emptyset\}$ specifying a single immediate predecessor of each node x ; $p(x)$ is nonempty for all $x \in \mathcal{X}$ but one, designated as the *initial node* x_0 . The immediate successor nodes of x are then $s(x) = p^{-1}(x)$, and the set of *all* predecessors and *all* successors of node x can be found by iterating $p(x)$ and $s(x)$. To have a tree structure, we require that these sets be disjoint (a predecessor of node x cannot also be a successor to it). The set of *terminal nodes* is $T = \{x \in \mathcal{X}: s(x) = \emptyset\}$. All other nodes $\mathcal{X} \setminus T$ are known as *decision nodes*.
- (iii) A function $\alpha: \mathcal{X} \setminus \{x_0\} \rightarrow \mathcal{A}$ giving the action that leads to any noninitial node x from its immediate predecessor $p(x)$ and satisfying the property that if $x', x'' \in s(x)$ and $x' \neq x''$, then $\alpha(x') \neq \alpha(x'')$. The set of choices available at decision node x is $c(x) = \{a \in \mathcal{A}: a = \alpha(x') \text{ for some } x' \in s(x)\}$.
- (iv) A collection of information sets \mathcal{H} , and a function $H: \mathcal{X} \rightarrow \mathcal{H}$ assigning each decision node x to an information set $H(x) \in \mathcal{H}$. Thus, the information sets in \mathcal{H} form a partition of \mathcal{X} . We require that all decision nodes assigned to a single information set have the same choices available; formally, $c(x) = c(x')$ if $H(x) = H(x')$. We can therefore write the choices available at information set H as $C(H) = \{a \in \mathcal{A}: a \in c(x) \text{ for } x \in H\}$.
- (v) A function $\iota: \mathcal{H} \rightarrow \{0, 1, \dots, I\}$ assigning each information set in \mathcal{H} to the player (or to nature: formally, player 0) who moves at the decision nodes in that set. We can denote the collection of player i 's information sets by $\mathcal{H}_i = \{H \in \mathcal{H}: i = \iota(H)\}$.
- (vi) A function $\rho: \mathcal{H}_0 \times \mathcal{A} \rightarrow [0, 1]$ assigning probabilities to actions at information sets where nature moves and satisfying $\rho(H, a) = 0$ if $a \notin C(H)$ and $\sum_{a \in C(H)} \rho(H, a) = 1$ for all $H \in \mathcal{H}_0$.
- (vii) A collection of payoff functions $u = \{u_1(\cdot), \dots, u_I(\cdot)\}$ assigning utilities to the players for each terminal node that can be reached, $u_i: T \rightarrow \mathbb{R}$. As we noted in Section 7.B, because we want to allow for a random realization of outcomes we take each $u_i(\cdot)$ to be a Bernoulli utility function.

Thus, formally, a game in extensive form is specified by the collection $\Gamma_E = \{\mathcal{X}, \mathcal{A}, I, p(\cdot), \alpha(\cdot), \mathcal{H}, H(\cdot), \iota(\cdot), \rho(\cdot), u\}$.

We should note that there are three implicit types of finiteness hidden in the formulation just presented. Because we will often encounter games not sharing these features in the economic applications discussed in later chapters, we briefly identify them here, although without any formal treatment. The formal definition of an extensive form representation of a game can be extended to these infinite cases without much difficulty, although there can be important differences in the predicted outcomes of finite and infinite economic models, as we shall see later (e.g., in Chapters 12 and 20).

First, we have assumed that players have a finite number of actions available at each decision node. This would rule out a game in which, say, a player can choose any number from some interval $[a, b] \subset \mathbb{R}$. In fact, allowing for an infinite set of actions requires that we allow for an infinite set of nodes as well. But with this change, items (i) to (vii) remain the basic elements of an extensive form representation (e.g., decision nodes and terminal nodes are still associated with a unique path through the tree).

Second, we have described the extensive form of a game that must end after a finite number of moves (because the set of decision nodes is finite). Indeed, all the examples we have considered so far fall into this category. There are, however, other types of games. For example, suppose that two players with infinite life spans (perhaps two firms) play Matching Pennies repeatedly every January 1. The players discount the money gained or lost at future dates with interest rate r and seek to maximize their discounted net gains. In this game, there are no terminal nodes. Even so, we can still associate discounted payoffs for the two players with every (infinite) sequence of moves the players make. Of course, actually drawing a complete game tree would be impossible, but the basic elements of the extensive form can nonetheless be captured as before (with payoffs being associated with paths through the tree rather than with terminal nodes).

Third, we may at times also imagine that there are an infinite number of players who take actions in a game. For example, models involving overlapping generations of players (as in various macroeconomic models) have this feature, as do models of entry in which we want to allow for an infinite number of potential firms. In the games of this type that we consider, this issue can be handled in a simple and natural manner.

Note that all three of these extensions require that we relax the assumption that there is a finite set of nodes. Games with a finite number of nodes, such as those we have been considering, are known as *finite games*.

For pedagogical purposes, we restrict our attention in Part II to finite games except where specifically indicated otherwise. The extension of the formal concepts we discuss here to the economic games studied later in the book that do not share these finiteness properties is straightforward.

7.D Strategies and the Normal Form Representation of a Game

A central concept of game theory is the notion of a player's *strategy*. A strategy is a *complete contingent plan*, or *decision rule*, that specifies how the player will act in *every possible distinguishable circumstance* in which she might be called upon to move. Recall that, from a player's perspective, the set of such circumstances is represented by her collection of information sets, with each information set representing a different distinguishable circumstance in which she may need to move (see Section 7.C). Thus, a player's strategy amounts to a specification of how she plans to move at each one of her information sets, should it be reached during play of the game. This is stated formally in Definition 7.D.1.

Definition 7.D.1: Let \mathcal{H}_i denote the collection of player i 's information sets, \mathcal{A} the set of possible actions in the game, and $C(H) \subset \mathcal{A}$ the set of actions possible at information set H . A *strategy* for player i is a function $s_i: \mathcal{H}_i \rightarrow \mathcal{A}$ such that $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$.

The fact that a strategy is a complete contingent plan cannot be overemphasized, and it is often a source of confusion to those new to game theory. When a player specifies her strategy, it is as if she had to write down an instruction book prior to play so that a representative could act on her behalf merely by consulting that book.

As a complete contingent plan, a strategy often specifies actions for a player at information sets that may not be reached during the actual play of the game.

For example, in Tick-Tack-Toe, player O's strategy describes what she will do on her first move if player X starts the game by marking the center square. But in the actual play of the game, player X might not begin in the center; she may instead mark the lower-right corner first, making this part of player O's plan no longer relevant.

In fact, there is an even subtler point: A player's strategy may include plans for actions that her own strategy makes irrelevant. For example, a complete contingent plan for player X in Tick-Tack-Toe includes a description of what she will do after she plays "center" and player O then plays "lower-right corner," even though her own strategy may call for her first move to be "upper-left corner." This probably seems strange: its importance will become apparent only when we talk about dynamic games in Chapter 9. Nevertheless, remember: *A strategy is a complete contingent plan that says what a player will do at each of her information sets if she is called on to play there.*

It is worthwhile to consider what the players' possible strategies are for some of the simple Matching Pennies games.

Example 7.D.1: *Strategies in Matching Pennies Version B.* In Matching Pennies Version B, a strategy for player 1 simply specifies her move at the game's initial node. She has two possible strategies: She can play heads (H) or tails (T). A strategy for player 2, on the other hand, specifies how she will play (H or T) at each of her two information sets, that is, how she will play if player 1 picks H and how she will play if player 1 picks T. Thus, player 2 has four possible strategies.

Strategy 1 (s_1): Play H if player 1 plays H; play H if player 1 plays T.

Strategy 2 (s_2): Play H if player 1 plays H; play T if player 1 plays T.

Strategy 3 (s_3): Play T if player 1 plays H; play H if player 1 plays T.

Strategy 4 (s_4): Play T if player 1 plays H; play T if player 1 plays T. ■

Example 7.D.2: *Strategies in Matching Pennies Version C.* In Matching Pennies Version C, player 1's strategies are exactly the same as in Version B; but player 2 now only has two possible strategies, "play H" and "play T", because she now has only one information set. She can no longer condition her action on player 1's previous action. ■

We will often find it convenient to represent a profile of players' strategy choices in an I -player game by a vector $s = (s_1, \dots, s_I)$, where s_i is the strategy chosen by player i . We will also sometimes write the strategy profile s as (s_i, s_{-i}) , where s_{-i} is the $(I - 1)$ vector of strategies for players other than i .

The Normal Form Representation of a Game

Every profile of strategies for the players $s = (s_1, \dots, s_I)$ induces an outcome of the game: a sequence of moves actually taken and a probability distribution over the terminal nodes of the game. Thus, for any profile of strategies (s_1, \dots, s_I) , we can deduce the payoffs received by each player. We might think, therefore, of specifying the game directly in terms of strategies and their associated payoffs. This second way to represent a game is known as the *normal* (or *strategic*) *form*. It is, in essence, a condensed version of the extensive form.

		Player 2			
		s_1	s_2	s_3	s_4
Player 1	H	-1, +1	-1, +1	+1, -1	+1, -1
	T	+1, -1	-1, +1	+1, -1	-1, +1

Figure 7.D.1
The normal form representation of Matching Pennies Version B.

Definition 7.D.2: For a game with I players, the *normal form representation* Γ_N specifies for each player i a set of strategies S_i (with $s_i \in S_i$) and a payoff function $u_i(s_1, \dots, s_I)$ giving the von Neumann–Morgenstern utility levels associated with the (possibly random) outcome arising from strategies (s_1, \dots, s_I) . Formally, we write $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$.

In fact, when describing a game in its normal form, there is no need to keep track of the specific moves associated with each strategy. Instead, we can simply number the various possible strategies of a player, writing player i 's strategy set as $S_i = \{s_{1i}, s_{2i}, \dots\}$ and then referring to each strategy by its number.

A concrete example of a game in normal form is presented in Example 7.D.3 for Matching Pennies Version B.

Example 7.D.3: *The Normal Form of Matching Pennies Version B.* We have already described the strategy sets of the two players in Example 7.D.1. The payoff functions are

$$u_1(s_1, s_2) = \begin{cases} +1 & \text{if } (s_1, s_2) = (\text{H, strategies 3 or 4}) \text{ or } (\text{T, strategies 1 or 3}), \\ -1 & \text{if } (s_1, s_2) = (\text{H, strategies 1 or 2}) \text{ or } (\text{T, strategies 2 or 4}), \end{cases}$$

and $u_2(s_1, s_2) = -u_1(s_1, s_2)$. A convenient way to summarize this information is in the “game box” depicted in Figure 7.D.1. The different rows correspond to the strategies of player 1, and the columns to those of player 2. Within each cell, the payoffs of the two players are depicted as $(u_1(s_1, s_2), u_2(s_1, s_2))$. ■

Exercise 7.D.2: Depict the normal forms for Matching Pennies Version C and the standard version of Matching Pennies.

The idea behind using the normal form representation to study behavior in a game is that a player's decision problem can be thought of as one of choosing her strategy (her contingent plan of action) given the strategies that she thinks her rivals will be adopting. Because each player is faced with this problem, we can think of the players as simultaneously choosing their strategies from the sets $\{S_i\}$. It is as if the players each simultaneously write down their strategies on slips of paper and hand them to a referee, who then computes the outcome of the game from the players' submitted strategies.

From the previous discussion, it is clear that for any extensive form representation of a game, there is a unique normal form representation (more precisely, it is unique up to any renaming or renumbering of the strategies). The converse is not true, however. Many different extensive forms may be represented by the same normal form. For example, the normal form shown in Figure 7.D.1 represents not only the extensive form in Figure 7.C.1 but also the




Figure 7.D.2 An extensive form whose normal form is that depicted in Figure 7.D.1.

extensive form in Figure 7.D.2. In the latter game, players move simultaneously, player 1 choosing between two strategies, L and R , and player 2 choosing among four strategies: a , b , c , and d . In terms of their representations in a game box, the only difference between the normal forms for these games lies in the “labels” given to the rows and columns.

Because the condensed representation of the game in the normal form generally omits some of the details present in the extensive form, we may wonder whether this omission is important or whether the normal form summarizes all of the strategically relevant information (as the last paragraph in regular type seems to suggest). The question can be put a little differently: Is the scenario in which players simultaneously write down their strategies and submit them to a referee really equivalent to their playing the game over time as described in the extensive form? This question is currently a subject of some controversy among game theorists. The debate centers on issues arising in dynamic games such as those studied in Chapter 9.

For the simultaneous-move games that we study in Chapter 8, in which all players choose their actions at the same time, the normal form captures *all* the strategically relevant information. In simultaneous-move games, a player's strategy is a simple non-contingent choice of an action. In this case, players' simultaneous choice of strategies in the normal form is clearly equivalent to their simultaneous choice of actions in the extensive form (captured there by having players not observing each other's choices).

E Randomized Choices

Up to this point, we have assumed that players make their choices with certainty. However, there is no a priori reason to exclude the possibility that a player could randomize when faced with a choice. Indeed, we will see in Chapters 8 and 9 that in certain circumstances the possibility of randomization can play an important role in the analysis of games.

As stated in Definition 7.D.1, a deterministic strategy for player i , which we now call a *pure strategy*, specifies a deterministic choice $s_i(H)$ at each of her information sets $H \in \mathcal{H}_i$. Suppose that player i 's (finite) set of pure strategies is S_i . One way for

the player to randomize is to choose randomly one element of this set. This kind of randomization gives rise to what is called a *mixed strategy*.

Definition 7.E.1: Given player i 's (finite) pure strategy set S_i , a *mixed strategy* for player i , $\sigma_i: S_i \rightarrow [0, 1]$, assigns to each pure strategy $s_i \in S_i$ a probability $\sigma_i(s_i) \geq 0$ that it will be played, where $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

Suppose that player i has M pure strategies in set $S_i = \{s_{1i}, \dots, s_{Mi}\}$. Player i 's set of possible mixed strategies can therefore be associated with the points of the following simplex (recall our use of a simplex to represent lotteries in Chapter 6):

$$\Delta(S_i) = \{(\sigma_{1i}, \dots, \sigma_{Mi}) \in \mathbb{R}^M: \sigma_{mi} \geq 0 \text{ for all } m = 1, \dots, M \text{ and } \sum_{m=1}^M \sigma_{mi} = 1\}.$$

This simplex is called the *mixed extension* of S_i . Note that a pure strategy can be viewed as a special case of a mixed strategy in which the probability distribution over the elements of S_i is degenerate.

When players randomize over their pure strategies, the induced outcome is itself random, leading to a probability distribution over the terminal nodes of the game. Since each player i 's normal form payoff function $u_i(s)$ is of the von Neumann–Morgenstern type, player i 's payoff given a profile of mixed strategies $\sigma = (\sigma_1, \dots, \sigma_I)$ for the I players is her expected utility $E_\sigma[u_i(s)]$, the expectation being taken with respect to the probabilities induced by σ on pure strategy profiles $s = (s_1, \dots, s_I)$. That is, letting $S = S_1 \times \dots \times S_I$, player i 's von Neumann–Morgenstern utility from mixed strategy profile σ is

$$\sum_{s \in S} [\sigma_1(s_1) \sigma_2(s_2) \dots \sigma_I(s_I)] u_i(s),$$

which, with a slight abuse of notation, we denote by $u_i(\sigma)$. Note that because we assume that each player randomizes on her own, we take the realizations of players' randomizations to be independent of one another.³

The basic definition of the normal form representation need not be changed to accommodate the possibility that players might choose to play mixed strategies. We can simply consider the normal form game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ in which players' strategy sets are extended to include both pure and mixed strategies.

Note that we can equivalently think of a player forming her mixed strategy as follows: Player i has access to a private signal θ_i that is uniformly distributed on the interval $[0, 1]$ and is independent of other players' signals, and she forms her mixed strategy by making her plan of action contingent on the realization of the signal. That is, she specifies a pure strategy $s_i(\theta_i) \in S_i$ for each realization of θ_i . We shall return to this alternative interpretation of mixed strategies in Chapter 8.

If we use the extensive form description of a game, there is another way that player i could randomize. Rather than randomizing over the potentially very

3. In Chapter 8, however, we discuss the possibility that players' randomizations could be correlated.