

Simultaneous-Move Games

8.A Introduction

We now turn to the central question of game theory: What should we expect to observe in a game played by rational players who are fully knowledgeable about the structure of the game and each others' rationality? In this chapter, we study *simultaneous-move* games, in which all players move only once and at the same time. Our motivation for beginning with these games is primarily pedagogic; they allow us to concentrate on the study of strategic interaction in the simplest possible setting and to defer until Chapter 9 some difficult issues that arise in more general, dynamic games.

In Section 8.B, we introduce the concepts of *dominant* and *dominated* strategies. These notions and their extension in the concept of *iterated dominance* provide a first and compelling restriction on the strategies rational players should choose to play.

In Section 8.C, we extend these ideas by defining the notion of a *rationalizable strategy*. We argue that the implication of players' common knowledge of each others' rationality and of the structure of the game is precisely that they will play rationalizable strategies.

Unfortunately, in many games, the set of rationalizable strategies does not yield a very precise prediction of the play that will occur. In the remaining sections of the chapter, we therefore study solution concepts that yield more precise predictions by adding "equilibrium" requirements regarding players' behavior.

Section 8.D begins our study of equilibrium-based solution concepts by introducing the important and widely applied concept of *Nash equilibrium*. This concept adds to the assumption of common knowledge of players' rationality a requirement of *mutually correct expectations*. By doing so, it often greatly narrows the set of predicted outcomes of a game. We discuss in some detail the reasonableness of this requirement, as well as the conditions under which we can be assured that a Nash equilibrium exists.

In Sections 8.E and 8.F, we examine two extensions of the Nash equilibrium concept. In Section 8.E, we broaden the notion of a Nash equilibrium to cover situations with *incomplete information*, where each player's payoffs may, to some extent, be known only by the player. This yields the concept of *Bayesian Nash*

equilibrium. In Section 8.F, we explore the implications of players entertaining the possibility that, with some small but positive probability, their opponents might make a mistake in choosing their strategies. We define the notion of a (*normal form*) *trembling-hand perfect Nash equilibrium*, an extension of the Nash equilibrium concept that requires that equilibria be robust to the possibility of small mistakes.

Throughout the chapter, we study simultaneous-move games using their normal form representations (see Section 7.D). Thus, we use $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ when we consider only pure (nonrandom) strategy choices and $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ when we allow for the possibility of randomized choices by the players (see Section 7.E for a discussion of randomized choices). We often denote a profile of pure strategies for player i 's opponents by $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I)$, with a similar meaning applying to the profile of mixed strategies σ_{-i} . We then write $s = (s_i, s_{-i})$ and $\sigma = (\sigma_i, \sigma_{-i})$. We also let $S = S_1 \times \dots \times S_I$ and $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_I$.

8.B Dominant and Dominated Strategies

We begin our study of simultaneous-move games by considering the predictions that can be made based on a relatively simple means of comparing a player's possible strategies: that of *dominance*.

To keep matters as simple as possible, we initially ignore the possibility that players might randomize in their strategy choices. Hence, our focus is on games $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ whose strategy sets allow for only pure strategies.

Consider the game depicted in Figure 8.B.1, the famous *Prisoner's Dilemma*. The story behind this game is as follows: Two individuals are arrested for allegedly engaging in a serious crime and are held in separate cells. The district attorney (the DA) tries to extract a confession from each prisoner. Each is privately told that if he is the only one to confess, then he will be rewarded with a light sentence of 1 year while the recalcitrant prisoner will go to jail for 10 years. However, if he is the only one not to confess, then it is he who will serve the 10-year sentence. If both confess, they will both be shown some mercy: they will each get 5 years. Finally, if neither confesses, it will still be possible to convict both of a lesser crime that carries a sentence of 2 years. Each player wishes to minimize the time he spends in jail (or maximize the negative of this, the payoffs that are depicted in Figure 8.B.1).

What will the outcome of this game be? There is only one plausible answer: (confess, confess). To see why, note that playing "confess" is each player's best strategy *regardless of what the other player does*. This type of strategy is known as a *strictly dominant strategy*.

		Prisoner 2	
		Don't Confess	Confess
Prisoner 1	Don't Confess	-2, -2	-10, -1
	Confess	-1, -10	-5, -5

Figure 8.B.1
The Prisoner's Dilemma.

Definition 8.B.1: A strategy $s_i \in S_i$ is a *strictly dominant strategy* for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for all $s'_i \neq s_i$, we have

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$.

In words, a strategy s_i is a strictly dominant strategy for player i if it maximizes uniquely player i 's payoff for any strategy that player i 's rivals might play. (The reason for the modifier *strictly* in Definition 8.B.1 will be made clear in Definition 8.B.3.) If a player has a strictly dominant strategy, as in the Prisoner's Dilemma, we should expect him to play it.

The striking aspect of the (confess, confess) outcome in the Prisoner's Dilemma is that although it is the one we expect to arise, it is not the best outcome for the players *jointly*; both players would prefer that neither of them confess. For this reason, the Prisoner's Dilemma is the paradigmatic example of self-interested, rational behavior *not* leading to a socially optimal result.

One way of viewing the outcome of the Prisoner's Dilemma is that, in seeking to maximize his own payoff, each prisoner has a negative effect on his partner; by moving away from the (don't confess, don't confess) outcome, a player reduces his jail time by 1 year but increases that of his partner by 8 (in Chapter 11, we shall see this as an example of an *externality*).

Dominated Strategies

Although it is compelling that players should play strictly dominant strategies if they have them, it is rare for such strategies to exist. Often, one strategy of player i 's may be best when his rivals play s_{-i} and another when they play some other strategies s'_{-i} (think of the standard Matching Pennies game in Chapter 7). Even so, we might still be able to use the idea of dominance to eliminate some strategies as possible choices. In particular, we should expect that player i will not play *dominated* strategies, those for which there is some alternative strategy that yields him a greater payoff regardless of what the other players do.

Definition 8.B.2: A strategy $s_i \in S_i$ is *strictly dominated* for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there exists another strategy $s'_i \in S_i$ such that for all $s_{-i} \in S_{-i}$,

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i}).$$

In this case, we say that strategy s'_i *strictly dominates* strategy s_i .

With this definition, we can restate our definition of a strictly dominant strategy (Definition 8.B.1) as follows: Strategy s_i is a strictly dominant strategy for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if it strictly dominates every other strategy in S_i .

Example 8.B.1: Consider the game shown in Figure 8.B.2. There is no strictly dominant strategy, but strategy D for player 1 is strictly dominated by strategy M (and also by strategy U). ■

Definition 8.D.3 presents a related, weaker notion of a dominated strategy that is of some importance.

		Player 2	
		L	R
Player 1	U	1, -1	-1, 1
	M	-1, 1	1, -1
	D	-2, 5	-3, 2

		Player 2	
		L	R
Player 1	U	5, 1	4, 0
	M	6, 0	3, 1
	D	6, 4	4, 4

Figure 8.B.2 (left)
Strategy D is strictly dominated.

Figure 8.B.3 (right)
Strategies U and M are weakly dominated.

Definition 8.B.3: A strategy $s_i \in S_i$ is *weakly dominated* in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there exists another strategy $s'_i \in S_i$ such that for all $s_{-i} \in S_{-i}$,

$$u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i}),$$

with strict inequality for *some* s_{-i} . In this case, we say that strategy s'_i *weakly dominates* strategy s_i . A strategy is a *weakly dominant strategy* for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if it weakly dominates every other strategy in S_i .

Thus, a strategy is weakly dominated if another strategy does at least as well for all s_{-i} and strictly better for some s_{-i} .

Example 8.B.2: Figure 8.B.3 depicts a game in which player 1 has two weakly dominated strategies, U and M. Both are weakly dominated by strategy D. ■

Unlike a strictly dominated strategy, a strategy that is only weakly dominated cannot be ruled out based solely on principles of rationality. For any alternative strategy that player i might pick, there is at least one profile of strategies for his rivals for which the weakly dominated strategy does as well. In Figure 8.B.3, for example, player 1 could rationally pick M if he was *absolutely sure* that player 2 would play L. Yet, if the probability of player 2 choosing strategy R was perceived by player 1 as positive (no matter how small), then M would not be a rational choice for player 1. *Caution* might therefore rule out M. More generally, weakly dominated strategies could be dismissed if players always believed that there was at least some positive probability that any strategies of their rivals could be chosen. We do not pursue this idea here, although we return to it in Section 8.F. For now, we continue to allow a player to entertain any conjecture about what an opponent might play, even a perfectly certain one.

Iterated Deletion of Strictly Dominated Strategies

As we have noted, it is unusual for elimination of strictly dominated strategies to lead to a unique prediction for a game (e.g., recall the game in Figure 8.B.2). However, the logic of eliminating strictly dominated strategies can be pushed further, as demonstrated in Example 8.B.3.

Example 8.B.3: In Figure 8.B.4, we depict a modification of the Prisoner's Dilemma, which we call the *DA's Brother*.

The story (a somewhat far-fetched one!) is now as follows: One of the prisoners, prisoner 1, is the DA's brother. The DA has some discretion in the fervor with which

		Prisoner 2	
		Don't Confess	Confess
Prisoner 1	Don't Confess	0, -2	-10, -1
	Confess	-1, -10	-5, -5

Figure 8.B.4
The DA's Brother.

he prosecutes and, in particular, can allow prisoner 1 to go free if neither of the prisoners confesses. With this change, if prisoner 2 confesses, then prisoner 1 should also confess; but “don’t confess” has become prisoner 1’s best strategy if prisoner 2 plays “don’t confess.” Thus, we are unable to rule out either of prisoner 1’s strategies as being dominated, and so elimination of strictly dominated (or, for that matter, weakly dominated) strategies does not lead to a unique prediction.

However, we can still derive a unique prediction in this game if we push the logic of eliminating strictly dominated strategies further. Note that “don’t confess” is still strictly dominated for prisoner 2. Furthermore, once prisoner 1 eliminates “don’t confess” as a possible action by prisoner 2, “confess” is prisoner 1’s unambiguously optimal action; that is, it is his strictly dominant strategy once the strictly dominated strategy of prisoner 2 has been deleted. Thus, the unique predicted outcome in the DA’s Brother game should still be (confess, confess). ■

Note the way players’ common knowledge of each other’s payoffs and rationality is used to solve the game in Example 8.B.3. Elimination of strictly dominated strategies requires only that each player be rational. What we have just done, however, requires not only that prisoner 2 be rational but also that prisoner 1 *know* that prisoner 2 is rational. Put somewhat differently, a player need not know anything about his opponents’ payoffs or be sure of their rationality to eliminate a strictly dominated strategy from consideration as his own strategy choice; but for the player to eliminate one of his strategies from consideration because it is dominated if his opponents never play *their* dominated strategies *does* require this knowledge.

As a general matter, if we are willing to assume that all players are rational *and* that this fact and the players’ payoffs are common knowledge (so everybody knows that everybody knows that . . . everybody is rational), then we do not need to stop after only two iterations. We can eliminate not only strictly dominated strategies and strategies that are strictly dominated after the first deletion of strategies but also strategies that are strictly dominated after this *next* deletion of strategies, and so on. Note that with each elimination of strategies, it becomes possible for additional strategies to become dominated because the fewer strategies that a player’s opponents might play, the more likely that a particular strategy of his is dominated. However, each additional iteration requires that players’ knowledge of each others’ rationality be one level deeper. A player must now know not only that his rivals are rational but also that they know that he is, and so on.

One feature of the process of iteratively eliminating strictly dominated strategies is that the order of deletion does not affect the set of strategies that remain in the end (see Exercise 8.B.4). That is, if at any given point several strategies (of one or

several players) are strictly dominated, then we can eliminate them all at once or in any sequence without changing the set of strategies that we ultimately end up with. This is fortunate, since we would worry if our prediction depended on the arbitrarily chosen order of deletion.

Exercise 8.B.5 presents an interesting example of a game for which the iterated removal of strictly dominated strategies yields a unique prediction: the *Cournot duopoly game* (which we will discuss in detail in Chapter 12).

The iterated deletion of *weakly* dominated strategies is harder to justify. As we have already indicated, the argument for deletion of a weakly dominated strategy for player i is that he contemplates the possibility that every strategy combination of his rivals occurs with positive probability. However, this hypothesis clashes with the logic of iterated deletion, which assumes, precisely, that eliminated strategies are not expected to occur. This inconsistency leads the iterative elimination of weakly dominated strategies to have the undesirable feature that it *can* depend on the order of deletion. The game in Figure 8.B.3 provides an example. If we first eliminate strategy U , we next eliminate strategy L , and we can then eliminate strategy M ; (D, R) is therefore our prediction. If, instead, we eliminate strategy M first, we next eliminate strategy R , and we can then eliminate strategy U ; now (D, L) is our prediction.

Allowing for Mixed Strategies

When we recognize that players may randomize over their pure strategies, the basic definitions of strictly dominated and dominant strategies can be generalized in a straightforward way.

Definition 8.B.4: A strategy $\sigma_i \in \Delta(S_i)$ is *strictly dominated* for player i in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if there exists another strategy $\sigma'_i \in \Delta(S_i)$ such that for all $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$,

$$u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}).$$

In this case, we say that strategy σ'_i *strictly dominates* strategy σ_i . A strategy σ_i is a *strictly dominant strategy* for player i in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if it strictly dominates every other strategy in $\Delta(S_i)$.

Using this definition and the structure of mixed strategies, we can say a bit more about the set of strictly dominated strategies in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$.

Note first that when we test whether a strategy σ_i is strictly dominated by strategy σ'_i for player i , we need only consider these two strategies' payoffs against the *pure* strategies of i 's opponents. That is,

$$u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}) \quad \text{for all } \sigma_{-i}$$

if and only if

$$u_i(\sigma'_i, s_{-i}) > u_i(\sigma_i, s_{-i}) \quad \text{for all } s_{-i}.$$

This follows because we can write

$$u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \left[\prod_{k \neq i} \sigma_k(s_k) \right] [u_i(\sigma'_i, s_{-i}) - u_i(\sigma_i, s_{-i})].$$




Figure 8.B.5

Domination of a pure strategy by a randomized strategy.

This expression is positive for all σ_{-i} if and only if $[u_i(\sigma'_i, s_{-i}) - u_i(\sigma_i, s_{-i})]$ is positive for all s_{-i} . One implication of this point is presented in Proposition 8.B.1.

Proposition 8.B.1: Player i 's pure strategy $s_i \in S_i$ is strictly dominated in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if and only if there exists another strategy $\sigma'_i \in \Delta(S_i)$ such that

$$u_i(\sigma'_i, s_{-i}) > u_i(s_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$.

Proposition 8.B.1 tells us that to test whether a pure strategy s_i is dominated when randomized play is possible, the test given in Definition 8.B.2 need only be augmented by checking whether any of player i 's mixed strategies does better than s_i against every possible profile of pure strategies by i 's rivals.

In fact, this extra requirement can eliminate additional pure strategies because a pure strategy s_i may be dominated only by a randomized combination of other pure strategies; that is, to dominate a strategy, even a pure one, it may be necessary to consider alternative strategies that involve randomization. To see this, consider the two-player game depicted in Figure 8.B.5(a). Player 1 has three strategies: U , M , and D . We can see that U is an excellent strategy when player 2 plays L but a poor one against R and that D is excellent against R and poor against L . Strategy M , on the other hand, is a good but not great strategy against both L and R . None of these three pure strategies is strictly dominated by any of the others. But if we allow player 1 to randomize, then playing U and D each with probability $\frac{1}{2}$ yields player 1 an expected payoff of 5 regardless of player 2's strategy, strictly dominating M (remember, payoffs are levels of von Neumann–Morgenstern utilities). This is shown in Figure 8.B.5(b), where player 1's expected payoffs from playing U , D , M , and the randomized strategy $\frac{1}{2}U + \frac{1}{2}D$ are plotted as points in \mathbb{R}^2 (the two dimensions correspond to a strategy's expected payoff for player 1 when player 2 plays R , denoted by u_R , and L , denoted by u_L). In the figure, the payoff vectors achievable by randomizing over U and D , and that from the randomized strategy $\frac{1}{2}U + \frac{1}{2}D$ in particular, lie on the line connecting points $(0, 10)$ and $(10, 0)$. As can be seen, the payoffs from $\frac{1}{2}U + \frac{1}{2}D$ strictly dominate those from strategy M .

Once we have determined the set of undominated pure strategies for player i , we need to consider which mixed strategies are undominated. We can immediately eliminate any mixed strategy that uses a dominated pure strategy; if pure strategy s_i is strictly dominated for player i , then so is every mixed strategy that assigns a positive probability to this strategy.

Exercise 8.B.6: Prove that if pure strategy s_i is a strictly dominated strategy in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, then so is any strategy that plays s_i with positive probability.

But these are not the only mixed strategies that may be dominated. A mixed strategy that randomizes over undominated pure strategies may itself be dominated. For example, if strategy M in Figure 8.B.5(a) instead gave player 1 a payoff of 6 for either strategy chosen by player 2, then although neither strategy U nor strategy D would be strictly dominated, the randomized strategy $\frac{1}{2}U + \frac{1}{2}D$ would be strictly dominated by strategy M [look where the point (6, 6) would lie in Figure 8.B.5(b)].

In summary, to find the set of strictly dominated strategies for player i in $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, we can first eliminate those pure strategies that are strictly dominated by applying the test in Proposition 8.B.1. Call player i 's set of undominated pure strategies $S_i^u \subset S_i$. Next, eliminate any mixed strategies in set $\Delta(S_i^u)$ that are dominated. Player i 's set of undominated strategies (pure and mixed) is exactly the remaining strategies in set $\Delta(S_i^u)$.

As when we considered only pure strategies, we can push the logic of removal of strictly dominated strategies in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ further through iterative elimination. The preceding discussion implies that this iterative procedure can be accomplished with the following two-stage procedure: First iteratively eliminate dominated pure strategies using the test in Proposition 8.B.1, applied at each stage using the remaining set of pure strategies. Call the remaining sets of pure strategies $\{\bar{S}_1^u, \dots, \bar{S}_I^u\}$. Then, eliminate any mixed strategies in sets $\{\Delta(\bar{S}_1^u), \dots, \Delta(\bar{S}_I^u)\}$ that are dominated.

8.C Rationalizable Strategies

In Section 8.B, we eliminated strictly dominated strategies based on the argument that a rational player would never choose such a strategy regardless of the strategies that he anticipates his rivals will play. We then used players' common knowledge of each others' rationality and the structure of the game to justify iterative removal of strictly dominated strategies.

In general, however, players' common knowledge of each others' rationality and the game's structure allows us to eliminate more than just those strategies that are iteratively strictly dominated. Here, we develop this point, leading to the concept of a *rationalizable strategy*. The set of rationalizable strategies consists precisely of those strategies that may be played in a game where the structure of the game and the players' rationality are common knowledge among the players. Throughout this section, we focus on games of the form $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ (mixed strategies are permitted).

We begin with Definition 8.C.1.

Definition 8.C.1: In game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, strategy σ_i is a *best response* for player i to his rivals' strategies σ_{-i} if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Delta(S_i)$. Strategy σ_i is *never a best response* if there is no σ_{-i} for which σ_i is a best response.

Strategy σ_i is a best response to σ_{-i} if it is an optimal choice when player i conjectures that his opponents will play σ_{-i} . Player i 's strategy σ_i is never a best response if there is no belief that player i may hold about his opponents' strategy

choices σ_{-i} that justifies choosing strategy σ_i .¹ Clearly, a player should not play a strategy that is never a best response.

Note that a strategy that is strictly dominated is never a best response. However, as a general matter, a strategy might never be a best response even though it is not strictly dominated (we say more about this relation at the end of this section in small type). Thus, eliminating strategies that are never a best response must eliminate at least as many strategies as eliminating just strictly dominated strategies and may eliminate more.

Moreover, as in the case of strictly dominated strategies, common knowledge of rationality and the game's structure implies that we can iterate the deletion of strategies that are never a best response. In particular, a rational player should not play a strategy that is never a best response once he eliminates the possibility that any of his rivals might play a strategy that is never a best response for them, and so on.

Equally important, the strategies that remain after this iterative deletion are the strategies that a rational player can *justify*, or *rationalize*, affirmatively with some reasonable conjecture about the choices of his rivals; that is, with a conjecture that does not assume that any player will play a strategy that is never a best response or one that is only a best response to a conjecture that someone else will play such a strategy, and so on. (Example 8.C.1 provides an illustration of this point.) As a result, the set of strategies surviving this iterative deletion process can be said to be precisely the set of strategies that can be played by rational players in a game in which the players' rationality and the structure of the game are common knowledge. They are known as *rationalizable strategies* [a concept developed independently by Bernheim (1984) and Pearce (1984)].

Definition 8.C.2: In game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, the strategies in $\Delta(S_i)$ that survive the iterated removal of strategies that are never a best response are known as player i 's *rationalizable strategies*.

Note that the set of rationalizable strategies can be no larger than the set of strategies surviving iterative removal of strictly dominated strategies because, at each stage of the iterative process in Definition 8.C.2, all strategies that are strictly dominated at that stage are eliminated. As in the case of iterated deletion of strictly dominated strategies, the order of removal of strategies that are never a best response can be shown not to affect the set of strategies that remain in the end (see Exercise 8.C.2).

1. We speak here as if a player's conjecture is necessarily deterministic in the sense that the player believes it is certain that his rivals will play a particular profile of mixed strategies σ_{-i} . One might wonder about conjectures that are probabilistic, that is, that take the form of a nondegenerate probability distribution over possible profiles of mixed strategy choices by his rivals. In fact, a strategy σ_i is an optimal choice for player i given some probabilistic conjecture (that treats his opponents' choices as independent random variables) only if it is an optimal choice given some deterministic conjecture. The reason is that if σ_i is an optimal choice given some probabilistic conjecture, then it must be a best response to the profile of mixed strategies σ_{-i} that plays each possible pure strategy profile $s_{-i} \in S_{-i}$ with exactly the compound probability implied by the probabilistic conjecture.

		Player 2			
		b_1	b_2	b_3	b_4
Player 1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

Figure 8.C.1

$\{a_1, a_2, a_3\}$
are rationalizable
strategies for player 1;
 $\{b_1, b_2, b_3\}$ are
rationalizable
strategies for player 2.

Example 8.C.1: Consider the game depicted in Figure 8.C.1, which is taken from Bernheim (1984). What is the set of rationalizable pure strategies for the two players? In the first round of deletion, we can eliminate strategy b_4 , which is never a best response because it is strictly dominated by a strategy that plays strategies b_1 and b_3 each with probability $\frac{1}{2}$. Once strategy b_4 is eliminated, strategy a_4 can be eliminated because it is strictly dominated by a_2 once b_4 is deleted. At this point, no further strategies can be ruled out: a_1 is a best response to b_3 , a_2 is a best response to b_2 , and a_3 is a best response to b_1 . Similarly, you can check that b_1 , b_2 , and b_3 are each best responses to one of a_1 , a_2 , and a_3 . Thus, the set of rationalizable pure strategies for player 1 is $\{a_1, a_2, a_3\}$, and the set $\{b_1, b_2, b_3\}$ is rationalizable for player 2.

Note that for each of these rationalizable strategies, a player can construct a *chain of justification* for his choice that never relies on any player believing that another player will play a strategy that is never a best response.² For example, in the game in Figure 8.C.1, player 1 can justify choosing a_2 by the belief that player 2 will play b_2 , which player 1 can justify to himself by believing that player 2 will think that he is going to play a_2 , which is reasonable if player 1 believes that player 2 is thinking that he, player 1, thinks player 2 will play b_2 , and so on. Thus, player 1 can construct an (infinite) chain of justification for playing strategy a_2 , $(a_2, b_2, a_2, b_2, \dots)$, where each element is justified using the next element in the sequence.

Similarly, player 1 can rationalize playing strategy a_1 with the chain of justification $(a_1, b_3, a_3, b_1, a_1, b_3, a_3, b_1, a_1, \dots)$. Here player 1 justifies playing a_1 by believing that player 2 will play b_3 . He justifies the belief that player 2 will play b_3 by thinking that player 2 believes that he, player 1, will play a_3 . He justifies this belief by thinking that player 2 thinks that he, player 1, believes that player 2 will play b_1 . And so on.

Suppose, however, that player 1 tried to justify a_4 . He could do so only by a belief that player 2 would play b_4 , but there is *no* belief that player 2 could have that would justify b_4 . Hence, player 1 cannot justify playing the nonrationalizable strategy a_4 . ■

2. In fact, this chain-of-justification approach to the set of rationalizable strategies is used in the original definition of the concept [for a formal treatment, consult Bernheim (1984) and Pearce (1984)].

It can be shown that under fairly weak conditions a player always has at least one rationalizable strategy.³ Unfortunately, players may have many rationalizable strategies, as in Example 8.C.1. If we want to narrow our predictions further, we need to make additional assumptions beyond common knowledge of rationality. The solution concepts studied in the remainder of this chapter do so by imposing “equilibrium” requirements on players’ strategy choices.

We have said that the set of rationalizable strategies is no larger than the set remaining after iterative deletion of strictly dominated strategies. It turns out, however, that for the case of two-player games ($I = 2$), these two sets are identical because in two-player games a (mixed) strategy σ_i is a best response to some strategy choice of a player’s rival whenever σ_i is not strictly dominated.

To see that this is plausible, reconsider the game in Figure 8.B.5 (Exercise 8.C.3 asks you for a general proof). Suppose that the payoffs from strategy M are altered so that M is not strictly dominated. Then, as depicted in Figure 8.C.2, the payoffs from M lie somewhere above




Figure 8.C.2

In a two-player game, a strategy is a best response if it is not strictly dominated.

the line connecting the points for strategies U and D . Is M a best response here? Yes. To see this, note that if player 2 plays strategy R with probability $\sigma_2(R)$, then player 1’s expected payoff from choosing a strategy with payoffs (u_R, u_L) is $\sigma_2(R)u_R + (1 - \sigma_2(R))u_L$. Points yielding the same expected payoff as strategy M therefore lie on a hyperplane with normal vector $(1 - \sigma_2(R), \sigma_2(R))$. As can be seen, strategy M is a best response to $\sigma_2(R) = \frac{1}{2}$; it yields an expected payoff strictly larger than any expected payoff achievable by playing strategies U and/or D .

With more than two players, however, there can be strategies that are never a best response and yet are not strictly dominated. The reason can be traced to the fact that players’ randomizations are independent. If the randomizations by i ’s rivals can be correlated (we discuss how this might happen at the end of Sections 8.D and 8.E), the equivalence reemerges. Exercise 8.C.4 illustrates these points.

3. This will be true, for example, whenever a Nash equilibrium (introduced in Section 8.D) exists.

8.D Nash Equilibrium

In this section, we present and discuss the most widely used solution concept in applications of game theory to economics, that of *Nash equilibrium* [due to Nash (1951)]. Throughout the rest of the book, we rely on it extensively.

For ease of exposition, we initially ignore the possibility that players might randomize over their pure strategies, restricting our attention to game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$. Mixed strategies are introduced later in the section.

We begin with Definition 8.D.1.

Definition 8.D.1: A strategy profile $s = (s_1, \dots, s_I)$ constitutes a *Nash equilibrium* of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for every $i = 1, \dots, I$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for all $s'_i \in S_i$.

In a Nash equilibrium, each player's strategy choice is a best response (see Definition 8.C.1) to the strategies *actually played* by his rivals. The italicized words distinguish the concept of Nash equilibrium from the concept of rationalizability studied in Section 8.C. Rationalizability, which captures the implications of the players' common knowledge of each others' rationality and the structure of the game, requires only that a player's strategy be a best response to some reasonable conjecture about what his rivals will be playing, where *reasonable* means that the conjectured play of his rivals can also be so justified. Nash equilibrium adds to this the requirement that players be *correct* in their conjectures.

Examples 8.D.1 and 8.D.2 illustrate the use of the concept.

Example 8.D.1: Consider the two-player simultaneous-move game shown in Figure 8.D.1. We can see that (M, m) is a Nash equilibrium. If player 1 chooses M , then the best response of player 2 is to choose m ; the reverse is true for player 2. Moreover, (M, m) is the only combination of (pure) strategies that is a Nash equilibrium. For example, strategy profile (U, r) cannot be a Nash equilibrium because player 1 would prefer to deviate to strategy D given that player 2 is playing r . (Check the other possibilities for yourself.) ■

Example 8.D.2: *Nash Equilibrium in the Game of Figure 8.C.1.* In this game, the unique Nash equilibrium profile of (pure) strategies is (a_2, b_2) . Player 1's best response to b_2 is a_2 , and player 2's best response to a_2 is b_2 , so (a_2, b_2) is a Nash equilibrium.

		Player 2		
		l	m	r
Player 1	U	5, 3	0, 4	3, 5
	M	4, 0	5, 5	4, 0
	D	3, 5	0, 4	5, 3

Figure 8.D.1
A Nash equilibrium.

		Mr. Schelling	
		Empire State	Grand Central
Mr. Thomas	Empire State	100, 100	0, 0
	Grand Central	0, 0	100, 100

Figure 8.D.2

Nash equilibria in the Meeting in New York game.

At any other strategy profile, one of the players has an incentive to deviate. [In fact, (a_2, b_2) is the unique Nash equilibrium even when randomization is permitted; see Exercise 8.D.1.]

This example illustrates a general relationship between the concept of Nash equilibrium and that of rationalizable strategies: *Every strategy that is part of a Nash equilibrium profile is rationalizable* because each player's strategy in a Nash equilibrium can be justified by the Nash equilibrium strategies of the other players. Thus, as a general matter, the Nash equilibrium concept offers at least as sharp a prediction as does the rationalizability concept. In fact, it often offers a *much* sharper prediction. In the game of Figure 8.C.1, for example, the rationalizable strategies a_1 , a_3 , b_1 , and b_3 are eliminated as predictions because they cannot be sustained when players' beliefs about each other's play are required to be correct. ■

In the previous two examples, the Nash equilibrium concept yields a unique prediction. However, this is not always the case. Consider the Meeting in New York game.

Example 8.D.3: *Nash Equilibria in the Meeting in New York Game.* Figure 8.D.2 depicts a simple version of the Meeting in New York game. Mr. Thomas and Mr. Schelling each have two choices: They can meet either at noon at the top of the Empire State Building or at noon at the clock in Grand Central Station. There are two Nash equilibria (ignoring the possibility of randomization): (Empire State, Empire State) and (Grand Central, Grand Central). ■

Example 8.D.3 emphasizes how strongly the Nash equilibrium concept uses the assumption of mutually correct expectations. The theory of Nash equilibrium is silent on *which* equilibrium we should expect to see when there are many. Yet, the players are assumed to correctly forecast which one it will be.

A compact restatement of the definition of a Nash equilibrium can be obtained through the introduction of the concept of a player's *best-response correspondence*. Formally, we say that player i 's best-response correspondence $b_i: S_{-i} \rightarrow S_i$ in the game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$, is the correspondence that assigns to each $s_{-i} \in S_{-i}$ the set

$$b_i(s_{-i}) = \{s_i \in S_i: u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \text{ for all } s'_i \in S_i\}.$$

With this notion, we can restate the definition of a Nash equilibrium as follows: The strategy profile (s_1, \dots, s_I) is a Nash equilibrium of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if and only if $s_i \in b_i(s_{-i})$ for $i = 1, \dots, I$.

Discussion of the Concept of Nash Equilibrium

Why might it be reasonable to expect players' conjectures about each other's play to be correct? Or, in sharper terms, why should we concern ourselves with the concept of Nash equilibrium?

A number of arguments have been put forward for the Nash equilibrium concept and you will undoubtedly react to them with varying degrees of satisfaction. Moreover, one argument might seem compelling in one application but not at all convincing in another. Until very recently, all these arguments have been informal, as will be our discussion. The issue is one of the more important open questions in game theory, particularly given the Nash equilibrium concept's widespread use in applied problems, and it is currently getting some formal attention.

(i) *Nash equilibrium as a consequence of rational inference.* It is sometimes argued that because each player can think through the strategic considerations faced by his opponents, rationality alone implies that players must be able to correctly forecast what their rivals will play. Although this argument may seem appealing, it is faulty. As we saw in Section 8.C, the implication of common knowledge of the players' rationality (and of the game's structure) is precisely that each player must play a rationalizable strategy. Rationality need not lead players' forecasts to be correct.

(ii) *Nash equilibrium as a necessary condition if there is a unique predicted outcome to a game.* A more satisfying version of the previous idea argues that if there is a unique predicted outcome for a game, then rational players will understand this. Therefore, for no player to wish to deviate, this predicted outcome must be a Nash equilibrium. Put somewhat differently [as in Kreps (1990)], if players think and share the belief that there is an *obvious* (in particular, a unique) way to play a game, then it must be a Nash equilibrium.

Of course, this argument is only relevant if there is a unique prediction for how players will play a game. The discussion of rationalizability in Section 8.C, however, shows that common knowledge of rationality alone does not imply this. Therefore, this argument is really useful only in conjunction with some reason why a particular profile of strategies might be the obvious way to play a particular game. The other arguments for Nash equilibrium that we discuss can be viewed as combining this argument with a reason why there might be an "obvious" way to play a game.

(iii) *Focal points.* It sometimes happens that certain outcomes are what Schelling (1960) calls *focal*. For example, take the Meeting in New York game depicted in Figure 8.D.2, and suppose that restaurants in the Grand Central area are so much better than those around the Empire State Building that the payoffs to meeting at Grand Central are (1000, 1000) rather than (100, 100). Suddenly, going to Grand Central seems like the obvious thing to do. Focal outcomes could also be culturally determined. As Schelling pointed out in his original discussion, two people who do not live in New York will tend to find meeting at the top of the Empire State building (a famous tourist site) to be focal, whereas two native New Yorkers will find Grand

Central Station (the central railroad station) a more compelling choice. In both examples, one of the outcomes has a natural appeal. The implication of argument (ii) is that this kind of appeal can lead an outcome to be the clear prediction in a game only if the outcome is a Nash equilibrium.

(iv) *Nash equilibrium as a self-enforcing agreement.* Another argument for Nash equilibrium comes from imagining that the players can engage in nonbinding communication prior to playing the game. If players agree to an outcome to be played, this naturally becomes the obvious candidate for play. However, because players cannot bind themselves to their agreed-upon strategies, any agreement that the players reach must be self-enforcing if it is to be meaningful. Hence, any meaningful agreement must involve the play of a Nash equilibrium strategy profile. Of course, even though players have reached an agreement to play a Nash equilibrium, they could still deviate from it if they expect others to do so. In essence, this justification assumes that once the players have agreed to a choice of strategies, this agreement becomes focal.

(v) *Nash equilibrium as a stable social convention.* A particular way to play a game might arise over time if the game is played repeatedly and some stable social convention emerges. If it does, it may be “obvious” to all players that the convention will be maintained. The convention, so to speak, becomes focal.

A good example is the game played by New Yorkers every day: Walking in Downtown Manhattan. Every day, people who walk to work need to decide which side of the sidewalk they will walk on. Over time, the stable social convention is that everyone walks on the right side, a convention that is enforced by the fact that any individual who unilaterally deviates from it is sure to be severely trampled. Of course, on any given day, it is *possible* that an individual might decide to walk on the left by conjecturing that everyone else suddenly expects the convention to change. Nevertheless, the prediction that we will remain at the Nash equilibrium “everyone walks on the right” seems reasonable in this case. Note that if an outcome is to become a stable social convention, it must be a Nash equilibrium. If it were not, then individuals would deviate from it as soon as it began to emerge.

The notion of an equilibrium as a rest point for some dynamic adjustment process underlies the use and the traditional appeal of equilibrium notions in economics. In this sense, the stable social convention justification of Nash equilibrium is closest to the tradition of economic theory.

To formally model the emergence of stable social conventions is not easy. One difficulty is that the repeated one-day game may itself be viewed as a larger dynamic game. Thus, when we consider rational players choosing their strategies in this overall game, we are merely led back to our original conundrum: Why should we expect a Nash equilibrium in this larger game? One response to this difficulty currently getting some formal attention imagines that players follow simple rules of thumb concerning their opponents’ likely play in situations where play is repeated (note that this implies a certain withdrawal from the assumption of complete rationality). For example, a player could conjecture that whatever his opponents did yesterday will be repeated today. If so, then each day players will play a best response to yesterday’s play. If a combination of strategies arises that is a stationary point of this process (i.e., the

		Player 2	
		Heads	Tails
Player 1	Heads	-1, +1	+1, -1
	Tails	+1, -1	-1, +1

Figure 8.D.3
Matching Pennies.

play today is the same as it was yesterday), it must be a Nash equilibrium. However, it is less clear that from any initial position, the process will converge to a stationary outcome; convergence turns out to depend on the game.⁴

Mixed Strategy Nash Equilibria

It is straightforward to extend the definition of Nash equilibrium to games in which we allow the players to randomize over their pure strategies.

Definition 8.D.2: A mixed strategy profile $\sigma = (\sigma_1, \dots, \sigma_I)$ constitutes a *Nash equilibrium* of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if for every $i = 1, \dots, I$,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Delta(S_i)$.

Example 8.D.4: As a very simple example, consider the standard version of Matching Pennies depicted in Figure 8.D.3. This is a game with no pure strategy equilibrium. On the other hand, it is fairly intuitive that there is a mixed strategy equilibrium in which each player chooses H or T with equal probability. When a player randomizes in this way, it makes his rival indifferent between playing heads or tails, and so his rival is also willing to randomize between heads and tails with equal probability. ■

It is not an accident that a player who is randomizing in a Nash equilibrium of Matching Pennies is indifferent between playing heads and tails. As Proposition 8.D.1 confirms, this indifference among strategies played with positive probability is a general feature of mixed strategy equilibria.

Proposition 8.D.1: Let $S_i^+ \subset S_i$ denote the set of pure strategies that player i plays with positive probability in mixed strategy profile $\sigma = (\sigma_1, \dots, \sigma_I)$. Strategy profile σ is a Nash equilibrium in game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ if and only if for all $i = 1, \dots, I$,

- (i) $u_i(s_i, \sigma_{-i}) = u_i(s'_i, \sigma_{-i})$ for all $s_i, s'_i \in S_i^+$;
- (ii) $u_i(s_i, \sigma_{-i}) \geq u_i(s'_i, \sigma_{-i})$ for all $s_i \in S_i^+$ and all $s'_i \notin S_i^+$.

Proof: For necessity, note that if either of conditions (i) or (ii) does not hold for some player i , then there are strategies $s_i \in S_i^+$ and $s'_i \in S_i$ such that $u_i(s'_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. If so, player i can strictly increase his payoff by playing strategy s'_i whenever he would have played strategy s_i .

4. This approach actually dates to Cournot's (1838) myopic adjustment procedure. A recent example can be found in Milgrom and Roberts (1990). Interestingly, this work explains the "ultrarational" Nash outcome by *relaxing* the assumption of rationality. It also can be used to try to identify the likelihood of various Nash equilibria arising when multiple Nash equilibria exist.

For sufficiency, suppose that conditions (i) and (ii) hold but that σ is not a Nash equilibrium. Then there is some player i who has a strategy σ'_i with $u_i(\sigma'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. But if so, then there must be some pure strategy s'_i that is played with positive probability under σ'_i for which $u_i(s'_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. Since $u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i^+$, this contradicts conditions (i) and (ii) being satisfied. ■

Hence, a necessary and sufficient condition for mixed strategy profile σ to be a Nash equilibrium of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ is that each player, given the distribution of strategies played by his opponents, is indifferent among all the pure strategies that he plays with positive probability and that these pure strategies are at least as good as any pure strategy he plays with zero probability.

An implication of Proposition 8.D.1 is that to test whether a strategy profile σ is a Nash equilibrium it suffices to consider only pure strategy deviations (i.e., changes in a player's strategy σ_i to some pure strategy s'_i). As long as no player can improve his payoff by switching to any pure strategy, σ is a Nash equilibrium. We therefore get the comforting result given in Corollary 8.D.1.

Corollary 8.D.1: Pure strategy profile $s = (s_1, \dots, s_I)$ is a Nash equilibrium of game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if and only if it is a (degenerate) mixed strategy Nash equilibrium of game $\Gamma'_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$.

Corollary 8.D.1 tells us that to identify the pure strategy equilibria of game $\Gamma'_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, it suffices to restrict attention to the game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ in which randomization is not permitted.

Proposition 8.D.1 can also be of great help in the computation of mixed strategy equilibria as Example 8.D.5 illustrates.

Example 8.D.5: Mixed Strategy Equilibria in the Meeting in New York Game. Let us try to find a mixed strategy equilibrium in the variation of the Meeting in New York game where the payoffs of meeting at Grand Central are (1000, 1000). By Proposition 8.D.1, if Mr. Thomas is going to randomize between Empire State and Grand Central, he must be indifferent between them. Suppose that Mr. Schelling plays Grand Central with probability σ_s . Then Mr. Thomas' expected payoff from playing Grand Central is $1000\sigma_s + 0(1 - \sigma_s)$, and his expected payoff from playing Empire State is $100(1 - \sigma_s) + 0\sigma_s$. These two expected payoffs are equal only when $\sigma_s = 1/11$. Now, for Mr. Schelling to set $\sigma_s = 1/11$, he must also be indifferent between his two pure strategies. By a similar argument, we find that Mr. Thomas' probability of playing Grand Central must also be $1/11$. We conclude that each player going to Grand Central with a probability of $1/11$ is a Nash equilibrium. ■

Note that in accordance with Proposition 8.D.1, the players in Example 8.D.5 have no real preference over the probabilities that they assign to the pure strategies they play with positive probability. What determines the probabilities that each player uses is an equilibrium consideration: the need to make the *other* player indifferent over *his* strategies.

This fact has led some economists and game theorists to question the usefulness of mixed strategy Nash equilibria as predictions of play. They raise two concerns: First, if players always have a pure strategy that gives them the same expected payoff as their equilibrium mixed strategy, it is not clear why they will bother to randomize.

One answer to this objection is that players may not actually randomize. Rather, they may make definite choices that are affected by seemingly inconsequential variables ("signals") that only they observe. For example, consider how a pitcher for a major league baseball team "mixes his pitches" to keep batters guessing. He may have a completely deterministic plan for what he will do, but it may depend on which side of the bed he woke up on that day or on the number of red traffic lights he came to on his drive to the stadium. As a result, batters view the behavior of the pitcher as random even though it is not. We touched briefly on this interpretation of mixed strategies as behavior contingent on realizations of a signal in Section 7.E, and we will examine it in more detail in Section 8.E.

The second concern is that the stability of mixed strategy equilibria seems tenuous. Players must randomize with exactly the correct probabilities, but they have no positive incentive to do so. One's reaction to this problem may depend on why one expects a Nash equilibrium to arise in the first place. For example, the use of the correct probabilities may be unlikely to arise as a stable social convention, but may seem more plausible when the equilibrium arises as a self-enforcing agreement.

Up to this point, we have assumed that players' randomizations are independent. In the Meeting in New York game in Example 8.D.5, for instance, we could describe a mixed strategy equilibrium as follows: Nature provides *private and independently distributed* signals $(\theta_1, \theta_2) \in [0, 1] \times [0, 1]$ to the two players, and each player i assigns decisions to the various possible realizations of his signal θ_i .

However, suppose that there are also *public* signals available that both players observe. Let $\theta \in [0, 1]$ be such a signal. Then many new possibilities arise. For example, the two players could both decide to go to Grand Central if $\theta < \frac{1}{2}$ and to Empire State if $\theta \geq \frac{1}{2}$. Each player's strategy choice is still random, but the coordination of their actions is now perfect and they always meet. More importantly, the decisions have an equilibrium character. If one player decides to follow this decision rule, then it is also optimal for the other player to do so. This is an example of a *correlated equilibrium* [due to Aumann (1974)]. More generally, we could allow for correlated equilibria in which nature's signals are partly private and partly public.

Allowing for such correlation may be important because economic agents observe many public signals. Formally, a correlated equilibrium is a special case of a Bayesian Nash equilibrium, a concept that we introduce in Section 8.E; hence, we defer further discussion to the end of that section.

Existence of Nash Equilibria

Does a Nash equilibrium necessarily exist in a game? Fortunately, the answer turns out to be "yes" under fairly broad circumstances. Here we describe two of the more important existence results; their proofs, based on mathematical fixed point theorems, are given in Appendix A of this chapter. (Proposition 9.B.1 of Section 9.B provides another existence result.)

Proposition 8.D.2: Every game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ in which the sets S_1, \dots, S_I have a finite number of elements has a mixed strategy Nash equilibrium.

Thus, for the class of games we have been considering, a Nash equilibrium always exists as long as we are willing to accept equilibria in which players randomize. (If you want to be convinced without going through the proof, try Exercise 8.D.6.) Allowing

for randomization is essential for this result. We have already seen in (standard) Matching Pennies, for example, that a pure strategy equilibrium may not exist in a game with a finite number of pure strategies.

Up to this point, we have focused on games with finite strategy sets. However, in economic applications, we frequently encounter games in which players have strategies naturally modeled as continuous variables. This can be helpful for the existence of a pure strategy equilibrium. In particular, we have the result given in Proposition 8.D.3.

Proposition 8.D.3: A Nash equilibrium exists in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for all $i = 1, \dots, I$,

- (i) S_i is a nonempty, convex, and compact subset of some Euclidean space \mathbb{R}^M .
- (ii) $u_i(s_1, \dots, s_I)$ is continuous in (s_1, \dots, s_I) and quasiconcave in s_i .

Proposition 8.D.3 provides a significant result whose requirements are satisfied in a wide range of economic applications. The convexity of strategy sets and the nature of the payoff functions help to smooth out the structure of the model, allowing us to achieve a pure strategy equilibrium.⁵

Further existence results can also be established. In situations where quasiconcavity of the payoff functions $u_i(\cdot)$ fails but they are still continuous, existence of a mixed strategy equilibrium can still be demonstrated. In fact, even if continuity of the payoff functions fails to hold, a mixed strategy equilibrium can be shown to exist in a variety of cases [see Dasgupta and Maskin (1986)].

Of course, these results do not mean that we *cannot* have an equilibrium if the conditions of these existence results do not hold. Rather, we just cannot be *assured* that there is one.

8.E Games of Incomplete Information: Bayesian Nash Equilibrium

Up to this point, we have assumed that players know all relevant information about each other, including the payoffs that each receives from the various outcomes of the game. Such games are known as games of *complete information*. A moment of thought, however, should convince you that this is a very strong assumption. Do two firms in an industry necessarily know each other's costs? Does a firm bargaining with a union necessarily know the disutility that union members will feel if they go out on strike for a month? Clearly, the answer is "no." Rather, in many circumstances, players have what is known as *incomplete information*.

The presence of incomplete information raises the possibility that we may need to consider a player's beliefs about other players' preferences, his beliefs about their beliefs about his preferences, and so on, much in the spirit of rationalizability.⁶

5. Note that a finite strategy set S_i cannot be convex. In fact, the use of mixed strategies in Proposition 8.D.2 helps us to obtain existence of equilibrium in much the same way that Proposition 8.D.3's assumptions assure existence of a pure strategy Nash equilibrium: It convexifies players' strategy sets and yields well-behaved payoff functions. (See Appendix A for details.)

6. For more on this problem, see Mertens and Zamir (1985).

Fortunately, there is a widely used approach to this problem, originated by Harsanyi (1967–68), that makes this unnecessary. In this approach, one imagines that each player's preferences are determined by the realization of a random variable. Although the random variable's actual realization is observed only by the player, its *ex ante* probability distribution is assumed to be common knowledge among all the players. Through this formulation, the situation of incomplete information is reinterpreted as a game of imperfect information: Nature makes the first move, choosing realizations of the random variables that determine each player's preference *type*, and each player observes the realization of only his own random variable. A game of this sort is known as a *Bayesian game*.

Example 8.E.1: Consider a modification of the DA's Brother game discussed in Example 8.B.3. With probability μ , prisoner 2 has the preferences in Figure 8.B.4 (we call these *type I preferences*), while with probability $(1 - \mu)$, prisoner 2 hates to rat on his accomplice (this is *type II*). In this case, he pays a psychic penalty equal to 6 years in prison for confessing. Prisoner 1, on the other hand, always has the preferences depicted in Figure 8.B.4. The extensive form of this Bayesian game is represented in Figure 8.E.1 (in the figure, "C" and "DC" stand for "confess" and "don't confess" respectively).

In this game, a pure strategy (a complete contingent plan) for player 2 can be viewed as a function that for each possible realization of his preference type

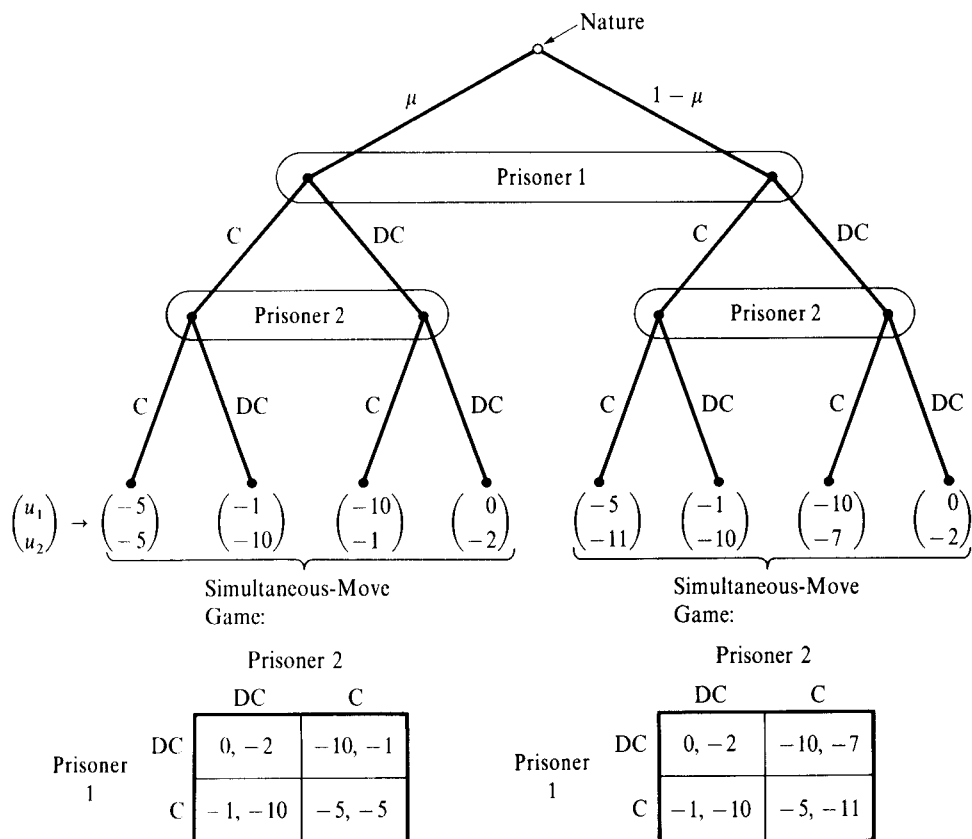


Figure 8.E.1
The DA's Brother game with incomplete information.

indicates what action he will take. Hence, prisoner 2 now has four possible pure strategies:

- (confess if type I, confess if type II);
- (confess if type I, don't confess if type II);
- (don't confess if type I, confess if type II);
- (don't confess if type I, don't confess if type II).

Notice, however, that player 1 does not observe player 2's type, and so a pure strategy for player 1 in this game is simply a (noncontingent) choice of either "confess" or "don't confess." ■

Formally, in a Bayesian game, each player i has a payoff function $u_i(s_i, s_{-i}, \theta_i)$, where $\theta_i \in \Theta_i$ is a random variable chosen by nature that is observed only by player i . The joint probability distribution of the θ_i 's is given by $F(\theta_1, \dots, \theta_I)$, which is assumed to be common knowledge among the players. Letting $\Theta = \Theta_1 \times \dots \times \Theta_I$, a Bayesian game is summarized by the data $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$.

A pure strategy for player i in a Bayesian game is a function $s_i(\theta_i)$, or *decision rule*, that gives the player's strategy choice for each realization of his type θ_i . Player i 's pure strategy set \mathcal{S}_i is therefore the set of all such functions. Player i 's expected payoff given a profile of pure strategies for the I players $(s_1(\cdot), \dots, s_I(\cdot))$ is then given by

$$\tilde{u}_i(s_1(\cdot), \dots, s_I(\cdot)) = E_\theta[u_i(s_1(\theta_1), \dots, s_I(\theta_I), \theta_i)]. \quad (8.E.1)$$

We can now look for an ordinary (pure strategy) Nash equilibrium of this game of imperfect information, which is known in this context as a *Bayesian Nash equilibrium*.⁷

Definition 8.E.1: A (pure strategy) *Bayesian Nash equilibrium* for the Bayesian game $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$ is a profile of decision rules $(s_1(\cdot), \dots, s_I(\cdot))$ that constitutes a Nash equilibrium of game $\Gamma_N = [I, \{\mathcal{S}_i\}, \{\tilde{u}_i(\cdot)\}]$. That is, for every $i = 1, \dots, I$,

$$\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot)) \geq \tilde{u}_i(s'_i(\cdot), s_{-i}(\cdot))$$

for all $s'_i(\cdot) \in \mathcal{S}_i$, where $\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot))$ is defined as in (8.E.1).

A very useful point to note is that in a (pure strategy) Bayesian Nash equilibrium each player must be playing a best response to the conditional distribution of his opponents' strategies *for each type that he might end up having*. Proposition 8.E.1 provides a more formal statement of this point.

Proposition 8.E.1: A profile of decision rules $(s_1(\cdot), \dots, s_I(\cdot))$ is a Bayesian Nash equilibrium in Bayesian game $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$ if and only if, for all i and

7. We shall restrict our attention to pure strategies here; mixed strategies involve randomization over the strategies in \mathcal{S}_i . Note also that we have not been very explicit about whether the Θ_i 's are finite sets. If they are, then the strategy sets \mathcal{S}_i are finite; if they are not, then the sets \mathcal{S}_i include an infinite number of possible functions $s_i(\cdot)$. Either way, however, the basic definition of a Bayesian Nash equilibrium is the same.

all $\bar{\theta}_i \in \Theta_i$ occurring with positive probability⁸

$$E_{\theta_{-i}}[u_i(s_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) | \bar{\theta}_i] \geq E_{\theta_{-i}}[u_i(s'_i, s_{-i}(\theta_{-i}), \bar{\theta}_i) | \bar{\theta}_i] \quad (8.E.2)$$

for all $s'_i \in S_i$, where the expectation is taken over realizations of the other players' random variables conditional on player i 's realization of his signal $\bar{\theta}_i$.

Proof: For necessity, note that if (8.E.2) did not hold for some player i for some $\bar{\theta}_i \in \Theta_i$ that occurs with positive probability, then player i could do better by changing his strategy choice in the event he gets realization $\bar{\theta}_i$, contradicting $(s_1(\cdot), \dots, s_i(\cdot))$ being a Bayesian Nash equilibrium. In the other direction, if condition (8.E.2) holds for all $\bar{\theta}_i \in \Theta_i$ occurring with positive probability, then player i cannot improve on the payoff he receives by playing strategy $s_i(\cdot)$. ■

Proposition 8.E.1 tells us that, in essence, we can think of each type of player i as being a separate player who maximizes his payoff given his conditional probability distribution over the strategy choices of his rivals.

Example 8.E.1 Continued: To solve for the (pure strategy) Bayesian Nash equilibrium of this game, note first that type I of prisoner 2 must play “confess” with probability 1 because this is that type’s dominant strategy. Likewise, type II of prisoner 2 also has a dominant strategy: “don’t confess.” Given this behavior by prisoner 2, prisoner 1’s best response is to play “don’t confess” if $[-10\mu + 0(1 - \mu)] > [-5\mu - 1(1 - \mu)]$, or equivalently, if $\mu < \frac{1}{6}$, and is to play “confess” if $\mu > \frac{1}{6}$. (He is indifferent if $\mu = \frac{1}{6}$.) ■

Example 8.E.2: The Alphabeta research and development consortium has two (noncompeting) members, firms 1 and 2. The rules of the consortium are that any independent invention by one of the firms is shared fully with the other. Suppose that there is a new invention, the “Zigger,” that either of the two firms could potentially develop. To develop this new product costs a firm $c \in (0, 1)$. The benefit of the Zigger to each firm i is known only by that firm. Formally, each firm i has a type θ_i that is independently drawn from a uniform distribution on $[0, 1]$, and its benefit from the Zigger when its type is θ_i is $(\theta_i)^2$. The timing is as follows: The two firms each privately observe their own type. Then they each simultaneously choose either to develop the Zigger or not.

Let us now solve for the Bayesian Nash equilibrium of this game. We shall write $s_i(\theta_i) = 1$ if type θ_i of firm i develops the Zigger and $s_i(\theta_i) = 0$ if it does not. If firm i develops the Zigger when its type is θ_i , its payoff is $(\theta_i)^2 - c$ regardless of whether firm j does so. If firm i decides not to develop the Zigger when its type is θ_i , it will have an expected payoff equal to $(\theta_i)^2 \text{Prob}(s_j(\theta_j) = 1)$. Hence, firm i 's best response is to develop the Zigger if and only if its type θ_i is such that (we assume firm i develops the Zigger if it is indifferent):

$$\theta_i \geq \left[\frac{c}{1 - \text{Prob}(s_j(\theta_j) = 1)} \right]^{1/2}. \quad (8.E.3)$$

8. The formulation given here (and the proof) is for the case in which the sets Θ_i are finite. When a player i has an infinite number of possible types, condition (8.E.2) must hold on a subset of Θ_i that is of full measure (i.e., that occurs with probability equal to one). It is then said that (8.E.2) holds for *almost every* $\bar{\theta}_i \in \Theta_i$.

Note that for any given strategy of firm j , firm i 's best response takes the form of a *cutoff rule*: It optimally develops the Zigger for all θ_i above the value on the right-hand side of (8.E.3) and does not for all θ_i below it. [Note that if firm i existed in isolation, it would be indifferent about developing the Zigger when $\theta_i = \sqrt{c}$. But (8.E.3) tells us that when firm i is part of the consortium, its cutoff is always (weakly) above this. This is true because each firm hopes to *free-ride* on the other firm's development effort; see Chapter 11 for more on this.]

Suppose then that $\hat{\theta}_1, \hat{\theta}_2 \in (0, 1)$ are the cutoff values for firms 1 and 2 respectively in a Bayesian Nash equilibrium (it can be shown that $0 < \hat{\theta}_i < 1$ for $i = 1, 2$ in any Bayesian Nash equilibrium of this game). If so, then using the fact that $\text{Prob}(s_j(\theta_j) = 1) = 1 - \hat{\theta}_j$, condition (8.E.3) applied first for $i = 1$ and then for $i = 2$ tells us that we must have

$$(\hat{\theta}_1)^2 \hat{\theta}_2 = c$$

and

$$(\hat{\theta}_2)^2 \hat{\theta}_1 = c.$$

Because $(\hat{\theta}_1)^2 \hat{\theta}_2 = (\hat{\theta}_2)^2 \hat{\theta}_1$ implies that $\hat{\theta}_1 = \hat{\theta}_2$, we see that any Bayesian Nash equilibrium of this game involves an identical cutoff value for the two firms, $\theta^* = (c)^{1/3}$. In this equilibrium, the probability that neither firm develops the Zigger is $(\theta^*)^2$, the probability that exactly one firm develops it is $2\theta^*(1 - \theta^*)$, and the probability that both do is $(1 - \theta^*)^2$. ■

The exercises at the end of this chapter consider several other examples of Bayesian Nash equilibria. Another important application arises in the theory of implementation with incomplete information, studied in Chapter 23.

In Section 8.D, we argued that mixed strategies could be interpreted as situations where players play deterministic strategies conditional on seemingly irrelevant signals (recall the baseball pitcher). We can now say a bit more about this. Suppose we start with a game of complete information that has a unique mixed strategy equilibrium in which players actually randomize. Now consider changing the game by introducing many different types (formally, a continuum) of each player, with the realizations of the various players' types being statistically independent of one another. Suppose, in addition, that all types of a player have *identical* preferences. A (pure strategy) Bayesian Nash equilibrium of this Bayesian game is then precisely equivalent to a mixed strategy Nash equilibrium of the original complete information game. Moreover, in many circumstances, one can show that there are also “nearby” Bayesian games in which preferences of the different types of a player differ only slightly from one another, the Bayesian Nash equilibria are close to the mixed strategy distribution, and each type has a strict preference for his strategy choice. Such results are known as *purification theorems* [see Harsanyi (1973)].

We can also return to the issue of *correlated equilibria* raised in Section 8.D. In particular, if we allow the realizations of the various players' types in the previous paragraph to be statistically correlated, then a (pure strategy) Bayesian Nash equilibrium of this Bayesian game is a correlated equilibrium of the original complete information game. The set of all correlated equilibria in game $[I, \{S_i\}, \{u_i(\cdot)\}]$ is identified by considering all possible Bayesian games of this sort (i.e., we allow for all possible signals that the players might observe).

8.F The Possibility of Mistakes: Trembling-Hand Perfection

In Section 8.B, we noted that although rationality per se does not rule out the choice of a weakly dominated strategy, such strategies are unappealing because they are dominated unless a player is absolutely sure of what his rivals will play. In fact, as the game depicted in Figure 8.F.1 illustrates, the Nash equilibrium concept also does not preclude the use of such strategies. In this game, (D, R) is a Nash equilibrium in which both players play a weakly dominated strategy with certainty.

Here, we elaborate on the idea, raised in Section 8.B, that *caution* might preclude the use of such strategies. The discussion leads us to define a refinement of the concept of Nash equilibrium, known as a (*normal form*) *trembling-hand perfect Nash equilibrium*, which identifies Nash equilibria that are robust to the possibility that, with some very small probability, players make mistakes.

Following Selten (1975), for any normal form game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, we can define a *perturbed* game $\Gamma_\varepsilon = [I, \{\Delta_\varepsilon(S_i)\}, \{u_i(\cdot)\}]$ by choosing for each player i and strategy $s_i \in S_i$ a number $\varepsilon_i(s_i) \in (0, 1)$, with $\sum_{s_i \in S_i} \varepsilon_i(s_i) < 1$, and then defining player i 's perturbed strategy set to be

$$\Delta_\varepsilon(S_i) = \{\sigma_i: \sigma_i(s_i) \geq \varepsilon_i(s_i) \text{ for all } s_i \in S_i \text{ and } \sum_{s_i \in S_i} \sigma_i(s_i) = 1\}.$$

That is, perturbed game Γ_ε is derived from the original game Γ_N by requiring that each player i play every one of his strategies, say s_i , with at least some minimal positive probability $\varepsilon_i(s_i)$; $\varepsilon_i(s_i)$ is interpreted as the unavoidable probability that strategy s_i gets played by mistake.

Having defined this perturbed game, we want to consider as predictions in game Γ_N only those Nash equilibria σ that are robust to the possibility that players make mistakes. The robustness test we employ can be stated roughly as: To consider σ as a robust equilibrium, we want there to be at least some slight perturbations of Γ_N whose equilibria are close to σ . The formal definition of a (*normal form*) *trembling-hand perfect Nash equilibrium* (the name comes from the idea of players making mistakes because of their trembling hands) is presented in Definition 8.F.1.

Definition 8.F.1: A Nash equilibrium σ of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ is (*normal form*) *trembling-hand perfect* if there is *some* sequence of perturbed games $\{\Gamma_{\varepsilon^k}\}_{k=1}^\infty$ that converges to Γ_N [in the sense that $\lim_{k \rightarrow \infty} \varepsilon_i^k(s_i) = 0$ for all i and $s_i \in S_i$], for which there is *some* associated sequence of Nash equilibria $\{\sigma^k\}_{k=1}^\infty$ that converges to σ (i.e., such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$).

We use the modifier *normal form* because Selten (1975) also proposes a slightly different form of trembling-hand perfection for dynamic games; we discuss this version of the concept in Chapter 9.⁹

Note that the concept of a (*normal form*) trembling-hand perfect Nash equilibrium provides a relatively mild test of robustness: We require only that *some* perturbed games exist that have equilibria arbitrarily close to σ . A stronger test would

9. In fact, Selten (1975) is primarily concerned with the problem of identifying desirable equilibria in dynamic games. See Chapter 9, Appendix B for more on this.

	<i>L</i>	<i>R</i>
<i>U</i>	1, 1	0, -3
<i>D</i>	-3, 0	0, 0

Figure 8.F.1

(D, R) is a Nash equilibrium involving play of weakly dominated strategies.

require that the equilibrium σ be robust to *all* perturbations close to the original game.

In general, the criterion in Definition 8.F.1 can be difficult to work with because it requires that we compute the equilibria of many possible perturbed games. The result presented in Proposition 8.F.1 provides a formulation that makes checking whether a Nash equilibrium is trembling-hand perfect much easier (in its statement, a *totally mixed* strategy is a mixed strategy in which every pure strategy receives positive probability).

Proposition 8.F.1: A Nash equilibrium σ of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ is (normal form) trembling-hand perfect if and only if there is some sequence of totally mixed strategies $\{\sigma^k\}_{k=1}^\infty$ such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$ and σ_i is a best response to every element of sequence $\{\sigma_{-i}^k\}_{k=1}^\infty$ for all $i = 1, \dots, I$.

You are asked to prove this result in Exercise 8.F.1 [or consult Selten (1975)]. The result presented in Proposition 8.F.2 is an immediate consequence of Definition 8.F.1 and Proposition 8.F.1.

Proposition 8.F.2: If $\sigma = (\sigma_1, \dots, \sigma_I)$ is a (normal form) trembling-hand perfect Nash equilibrium, then σ_i is not a weakly dominated strategy for any $i = 1, \dots, I$. Hence, in any (normal form) trembling-hand perfect Nash equilibrium, no weakly dominated pure strategy can be played with positive probability.

The converse, that any Nash equilibrium not involving play of a weakly dominated strategy is necessarily trembling-hand perfect, turns out to be true for two-player games but not for games with more than two players. Thus, trembling-hand perfection can rule out more than just Nash equilibria involving weakly dominated strategies. The reason is tied to the fact that when a player's rivals make mistakes with small probability, this can give rise to only a limited set of probability distributions over their nonequilibrium strategies. For example, if a player's two rivals each have a small probability of making a mistake, there is a much greater probability that one will make a mistake than that both will. If the player's equilibrium strategy is a unique best response only when both of his rivals make a mistake, his strategy may not be a best response to any local perturbation of his rivals' strategies even though his strategy is not weakly dominated. (Exercise 8.F.2 provides an example.) However, if players' trembles are allowed to be correlated (e.g., as in the correlated equilibrium concept), then the converse of Proposition 8.F.2 would hold regardless of the number of players.

Selten (1975) also proves an existence result that parallels Proposition 8.D.2: Every game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ with finite strategy sets S_1, \dots, S_I has a trembling-hand perfect Nash equilibrium. An implication of this result is that every such game has at least one Nash equilibrium in which no player plays any weakly dominated strategy with positive probability. Hence, if we decide to accept only Nash

equilibria that do not involve the play of weakly dominated strategies, with great generality there is at least one such equilibrium.¹⁰

Myerson (1978) proposes a refinement of Selten's idea in which players are less likely to make more costly mistakes (the idea is that they will try harder to avoid these mistakes). He establishes that the resulting solution concept, called a *proper Nash equilibrium*, exists under the conditions described in the previous paragraph for trembling-hand perfect Nash equilibria. van Damme (1983) presents a good discussion of this and other refinements of trembling-hand perfection.

APPENDIX A: EXISTENCE OF NASH EQUILIBRIUM

In this appendix, we prove Propositions 8.D.2 and 8.D.3. We begin with Lemma 8.AA.1, which provides a key technical result.

Lemma 8.AA.1: If the sets S_1, \dots, S_I are nonempty, S_i is compact and convex, and $u_i(\cdot)$ is continuous in (s_1, \dots, s_I) and quasiconcave in s_i , then player i 's best-response correspondence $b_i(\cdot)$ is nonempty, convex-valued, and upper hemicontinuous.¹¹

Proof: Note first that $b_i(s_{-i})$ is the set of maximizers of the continuous function $u_i(\cdot, s_{-i})$ on the compact set S_i . Hence, it is nonempty (see Theorem M.F.2 of the Mathematical Appendix). The convexity of $b_i(s_{-i})$ follows because the set of maximizers of a quasiconcave function [here, the function $u_i(\cdot, s_{-i})$] on a convex set (here, S_i) is convex. Finally, for upper hemicontinuity, we need to show that for any sequence $(s_i^n, s_{-i}^n) \rightarrow (s_i, s_{-i})$ such that $s_i^n \in b_i(s_{-i}^n)$ for all n , we have $s_i \in b_i(s_{-i})$. To see this, note that for all n , $u_i(s_i^n, s_{-i}^n) \geq u_i(s_i', s_{-i}^n)$ for all $s_i' \in S_i$. Therefore, by the continuity of $u_i(\cdot)$, we have $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$. ■

It is convenient to prove Proposition 8.D.3 first.

Proposition 8.D.3: A Nash equilibrium exists in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if for all $i = 1, \dots, I$,

- (i) S_i is a nonempty, convex, and compact subset of some Euclidean space \mathbb{R}^M .
- (ii) $u_i(s_1, \dots, s_I)$ is continuous in (s_1, \dots, s_I) and quasiconcave in s_i .

10. The Bertrand duopoly game discussed in Chapter 12 provides one example of a game in which this is not the case; its unique Nash equilibrium involves the play of weakly dominated strategies. The problem arises because the strategies in that game are continuous variables (and so the sets S_i are not finite). Fortunately, this equilibrium can be viewed as the limit of undominated equilibria in "nearby" discrete versions of the game. (See Exercise 12.C.3 for more on this point.)

11. See Section M.H of the Mathematical Appendix for a discussion of upper hemicontinuous correspondences.

Proof: Define the correspondence $b: S \rightarrow S$ by

$$b(s_1, \dots, s_I) = b_1(s_{-1}) \times \dots \times b_I(s_{-I}).$$

Note that $b(\cdot)$ is a correspondence from the nonempty, convex, and compact set $S = S_1 \times \dots \times S_I$ to itself. In addition, by Lemma 8.AA.1, $b(\cdot)$ is a nonempty, convex-valued, and upper hemicontinuous correspondence. Thus, all the conditions of the Kakutani fixed point theorem are satisfied (see Section M.I of the Mathematical Appendix). Hence, there exists a fixed point for this correspondence, a strategy profile $s \in S$ such that $s \in b(s)$. The strategies at this fixed point constitute a Nash equilibrium because by construction $s_i \in b_i(s_{-i})$ for all $i = 1, \dots, I$. ■

Now we move to the proof of Proposition 8.D.2.

Proposition 8.D.2: Every game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ in which the sets S_1, \dots, S_I have a finite number of elements has a mixed strategy Nash equilibrium.

Proof: The game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, viewed as a game with strategy sets $\{\Delta(S_i)\}$ and payoff functions $u_i(\sigma_1, \dots, \sigma_I) = \sum_{s \in S} [\prod_{k=1}^I \sigma_k(s_k)] u_i(s)$ for all $i = 1, \dots, I$, satisfies all the assumptions of Proposition 8.D.3. Hence, Proposition 8.D.2 is a direct corollary of that result. ■

REFERENCES

- Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1: 67–96.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica* 52: 1007–28.
- Bernheim, B. D. (1986). Axiomatic characterizations of rational choice in strategic environments. *Scandinavian Journal of Economics* 88: 473–88.
- Brandenberger, A., and E. Dekel. (1987). Rationalizability and correlated equilibria. *Econometrica* 55: 1391–1402.
- Cournot, A. (1838). *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. [English edition: *Researches into the Mathematical Principles of the Theory of Wealth*. New York: Macmillan, 1897.]
- Dasgupta, P., and E. Maskin. (1986). The existence of equilibrium in discontinuous economic games. *Review of Economic Studies* 53: 1–41.
- Harsanyi, J. (1967–68). Games with incomplete information played by Bayesian players. *Management Science* 14: 159–82, 320–34, 486–502.
- Harsanyi, J. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2: 1–23.
- Kreps, D. M. (1990). *Game Theory and Economic Modelling*. Oxford: Oxford University Press.
- Mertens, J. F., and S. Zamir. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* 10: 619–32.
- Milgrom, P., and J. Roberts. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58: 1255–78.
- Myerson, R. B. (1978). Refinements of the Nash equilibrium concept. *International Journal of Game Theory* 7: 73–80.
- Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics* 54: 289–95.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52: 1029–50.
- Schelling, T. (1960). *The Strategy of Conflict*. Cambridge, Mass.: Harvard University Press.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4: 25–55.
- van Damme, E. (1983). *Refinements of the Nash Equilibrium Concept*. Berlin: Springer-Verlag.

EXERCISES

8.B.1^A There are I firms in an industry. Each can try to convince Congress to give the industry a subsidy. Let h_i denote the number of hours of effort put in by firm i , and let $c_i(h_i) = w_i(h_i)^2$, where w_i is a positive constant, be the cost of this effort to firm i . When the effort levels of the firms are (h_1, \dots, h_I) , the value of the subsidy that gets approved is $\alpha \sum_i h_i + \beta (\prod_i h_i)$, where α and β are constants.

Consider a game in which the firms decide simultaneously and independently how many hours they will each devote to this effort. Show that each firm has a strictly dominant strategy if and only if $\beta = 0$. What is firm i 's strictly dominant strategy when this is so?

8.B.2^B (a) Argue that if a player has two weakly dominant strategies, then for every strategy choice by his opponents, the two strategies yield him equal payoffs.

(b) Provide an example of a two-player game in which a player has two weakly dominant pure strategies but his opponent prefers that he play one of them rather than the other.

8.B.3^B Consider the following auction (known as a *second-price*, or *Vickrey*, auction). An object is auctioned off to I bidders. Bidder i 's valuation of the object (in monetary terms) is v_i . The auction rules are that each player submit a bid (a nonnegative number) in a sealed envelope. The envelopes are then opened, and the bidder who has submitted the highest bid gets the object but pays the auctioneer the amount of the *second-highest* bid. If more than one bidder submits the highest bid, each gets the object with equal probability. Show that submitting a bid of v_i with certainty is a weakly dominant strategy for bidder i . Also argue that this is bidder i 's unique weakly dominant strategy.

8.B.4^C Show that the order of deletion does not matter for the set of strategies surviving a process of iterated deletion of strictly dominated strategies.

8.B.5^C Consider the Cournot duopoly model (discussed extensively in Chapter 12) in which two firms, 1 and 2, simultaneously choose the quantities they will sell on the market, q_1 and q_2 . The price each receives for each unit given these quantities is $P(q_1, q_2) = a - b(q_1 + q_2)$. Their costs are c per unit sold.

(a) Argue that successive elimination of strictly dominated strategies yields a unique prediction in this game.

(b) Would this be true if there were three firms instead of two?

8.B.6^B In text.

8.B.7^B Show that any strictly dominant strategy in game $[I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ must be a pure strategy.

8.C.1^A Argue that if elimination of strictly dominated strategies yields a unique prediction in a game, this prediction also results from eliminating strategies that are never a best response.

8.C.2^C Prove that the order of removal does not matter for the set of strategies that survives a process of iterated deletion of strategies that are never a best response.

8.C.3^C Prove that in a two-player game (with finite strategy sets), if a pure strategy s_i for player i is never a best response for any mixed strategy by i 's opponent, then s_i is strictly dominated by some mixed strategy $\sigma_i \in \Delta(S_i)$. [Hint: Try using the supporting hyperplane theorem presented in Section M.G of the Mathematical Appendix.]

8.C.4^B Consider a game Γ_N with players 1, 2, and 3 in which $S_1 = \{L, M, R\}$, $S_2 = \{U, D\}$, and $S_3 = \{\ell, r\}$. Player 1's payoffs from each of his three strategies conditional on the strategy choices of players 2 and 3 are depicted as (u_L, u_M, u_R) in each of the four boxes shown below, where $(\pi, \varepsilon, \eta) \gg 0$. Assume that $\eta < 4\varepsilon$.

		Player 3's Strategy	
		ℓ	r
Player 2's Strategy	U	$\pi - \pi + 4\varepsilon, \pi - \pi + \eta, \pi - \pi + 4\varepsilon$	$\pi - \pi + 4\varepsilon, \pi - \frac{\eta}{2}, \pi - \pi + 4\varepsilon$
	D	$\eta + 4\varepsilon, \pi + \frac{\eta}{2}, \pi - 4\varepsilon$	$\pi - 4\varepsilon, \pi - \eta, \pi + 4\varepsilon$

(a) Argue that (pure) strategy M is never a best response for player 1 to any independent randomizations by players 2 and 3.

(b) Show that (pure) strategy M is not strictly dominated.

(c) Show that (pure) strategy M can be a best response if player 2's and player 3's randomizations are allowed to be correlated.

8.D.1^B Show that (a_2, b_2) being played with certainty is the unique mixed strategy Nash equilibrium in the game depicted in Figure 8.C.1.

8.D.2^B Show that if there is a unique profile of strategies that survives iterated removal of strictly dominated strategies, this profile is a Nash equilibrium.

8.D.3^B Consider a first-price sealed-bid auction of an object with two bidders. Each bidder i 's valuation of the object is v_i , which is known to both bidders. The auction rules are that each player submits a bid in a sealed envelope. The envelopes are then opened, and the bidder who has submitted the highest bid gets the object and pays the auctioneer the amount of his bid. If the bidders submit the same bid, each gets the object with probability $\frac{1}{2}$. Bids must be in dollar multiples (assume that valuations are also).

(a) Are any strategies strictly dominated?

(b) Are any strategies weakly dominated?

(c) Is there a Nash equilibrium? What is it? Is it unique?

8.D.4^B Consider a bargaining situation in which two individuals are considering undertaking a business venture that will earn them 100 dollars in profit, but they must agree on how to split the 100 dollars. Bargaining works as follows: The two individuals each make a demand simultaneously. If their demands sum to more than 100 dollars, then they fail to agree, and each gets nothing. If their demands sum to less than 100 dollars, they do the project, each gets his demand, and the rest goes to charity.

(a) What are each player's strictly dominated strategies?

(b) What are each player's weakly dominated strategies?

(c) What are the pure strategy Nash equilibria of this game?

8.D.5^B Consumers are uniformly distributed along a boardwalk that is 1 mile long. Ice-cream prices are regulated, so consumers go to the nearest vendor because they dislike walking (assume that at the regulated prices all consumers will purchase an ice cream even if they

have to walk a full mile). If more than one vendor is at the same location, they split the business evenly.

(a) Consider a game in which two ice-cream vendors pick their locations simultaneously. Show that there exists a unique pure strategy Nash equilibrium and that it involves both vendors locating at the midpoint of the boardwalk.

(b) Show that with three vendors, no pure strategy Nash equilibrium exists.

8.D.6^B Consider any two-player game of the following form (where letters indicate arbitrary payoffs):

		Player 2	
		b_1	b_2
Player 1	a_1	u, v	ℓ, m
	a_2	w, x	y, z

Show that a mixed strategy Nash equilibrium always exists in this game. [Hint: Define player 1's strategy to be his probability of choosing action a_1 and player 2's to be his probability of choosing b_1 ; then examine the best-response correspondences of the two players.]

8.D.7^C (The Minimax Theorem) A two-player game with finite strategy sets $\Gamma_N = [I, \{S_1, S_2\}, \{u_1(\cdot), u_2(\cdot)\}]$ is a *zero-sum* game if $u_2(s_1, s_2) = -u_1(s_1, s_2)$ for all $(s_1, s_2) \in S_1 \times S_2$.

Define i 's *maximin* expected utility level \underline{w}_i to be the level he can guarantee himself in game $[I, \{\Delta(S_1), \Delta(S_2)\}, \{u_1(\cdot), u_2(\cdot)\}]$:

$$\underline{w}_i = \max_{\sigma_i} \left[\min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \right].$$

Define player i 's *minimax* utility level \underline{v}_i to be the worst expected utility level he can be forced to receive if he gets to respond to his rival's actions:

$$\underline{v}_i = \min_{\sigma_{-i}} \left[\max_{\sigma_i} u_i(\sigma_i, \sigma_{-i}) \right].$$

(a) Show that $\underline{v}_i \geq \underline{w}_i$ in any game.

(b) Prove that in any mixed strategy Nash equilibrium of the zero-sum game $\Gamma_N = [I, \{\Delta(S_1), \Delta(S_2)\}, \{u_1(\cdot), u_2(\cdot)\}]$, player i 's expected utility u_i° satisfies $u_i^\circ = \underline{v}_i = \underline{w}_i$. [Hint: Such an equilibrium must exist by Proposition 8.D.2.]

(c) Show that if (σ'_1, σ'_2) and (σ''_1, σ''_2) are both Nash equilibria of the zero-sum game $\Gamma_N = [I, \{\Delta(S_1), \Delta(S_2)\}, \{u_1(\cdot), u_2(\cdot)\}]$, then so are (σ'_1, σ''_2) and (σ''_1, σ'_2) .

8.D.8^C Consider a simultaneous-move game with normal form $[I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$. Suppose that, for all i , S_i is a convex set and $u_i(\cdot)$ is strictly quasiconvex. Argue that any mixed strategy Nash equilibrium of this game must be degenerate, with each player playing a single pure strategy with probability 1.

8.D.9^B Consider the following game [based on an example from Kreps (1990)]:

		Player 2			
		LL	L	M	R
Player 1	U	100, 2	-100, 1	0, 0	-100, -100
	D	-100, -100	100, -49	1, 0	100, 2

- (a) If you were player 2 in this game and you were playing it once without the ability to engage in preplay communication with player 1, what strategy would you choose?
- (b) What are all the Nash equilibria (pure and mixed) of this game?
- (c) Is your strategy choice in (a) a component of any Nash equilibrium strategy profile? Is it a rationalizable strategy?
- (d) Suppose now that preplay communication were possible. Would you expect to play something different from your choice in (a)?

8.E.1^B Consider the following strategic situation. Two opposed armies are poised to seize an island. Each army's general can choose either "attack" or "not attack." In addition, each army is either "strong" or "weak" with equal probability (the draws for each army are independent), and an army's type is known only to its general. Payoffs are as follows: The island is worth M if captured. An army can capture the island either by attacking when its opponent does not or by attacking when its rival does if it is strong and its rival is weak. If two armies of equal strength both attack, neither captures the island. An army also has a "cost" of fighting, which is s if it is strong and w if it is weak, where $s < w$. There is no cost of attacking if its rival does not.

Identify all pure strategy Bayesian Nash equilibria of this game.

8.E.2^C Consider the first-price sealed-bid auction of Exercise 8.D.3, but now suppose that each bidder i observes only his own valuation v_i . This valuation is distributed uniformly and independently on $[0, \bar{v}]$ for each bidder.

(a) Derive a symmetric (pure strategy) Bayesian Nash equilibrium of this auction. (You should now suppose that bids can be any real number.) [Hint: Look for an equilibrium in which bidder i 's bid is a linear function of his valuation.]

(b) What if there are I bidders? What happens to each bidder's equilibrium bid function $s(v_i)$ as I increases?

8.E.3^B Consider the linear Cournot model described in Exercise 8.B.5. Now, however, suppose that each firm has probability μ of having unit costs of c_L and $(1 - \mu)$ of having unit costs of c_H , where $c_H > c_L$. Solve for the Bayesian Nash equilibrium.

8.F.1^C Prove Proposition 8.F.1.

8.F.2^B Consider the following three-player game [taken from van Damme (1983)], in which player 1 chooses rows ($S_1 = \{U, D\}$), player 2 chooses columns ($S_2 = \{L, R\}$), and player 3 chooses boxes ($S_3 = \{B_1, B_2\}$):

		B_1				B_2	
		L	R			L	R
U	D	(1, 1, 1)	(1, 0, 1)	U	D	(1, 1, 0)	(0, 0, 0)
		(1, 1, 1)	(0, 0, 1)			(0, 1, 0)	(1, 0, 0)

Each cell describes the payoffs to the three players (u_1, u_2, u_3) from that strategy combination. Both (D, L, B_1) and (U, L, B_1) are pure strategy Nash equilibria. Show that (D, L, B_1) is not (normal form) trembling-hand perfect even though none of these three strategies is weakly dominated.

8.F.3^C Prove that every game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ in which the S_i are finite sets has a (normal form) trembling-hand perfect Nash equilibrium. [*Hint*: Show that every perturbed game has an equilibrium and that for any sequence of perturbed games converging to the original game Γ_N and corresponding sequence of equilibria, there is a subsequence that converges to an equilibrium of Γ_N .]