

Draft chapter from *An introduction to game theory* by Martin J. Osborne. Version: 2002/7/23.
 Martin.Osborne@utoronto.ca
<http://www.economics.utoronto.ca/osborne>
 Copyright © 1995–2002 by Martin J. Osborne. All rights reserved. No part of this book may be reproduced by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from Oxford University Press, except that one copy of up to six chapters may be made by any individual for private study.

2

Nash Equilibrium: Theory

2.1	Strategic games	11
2.2	Example: the <i>Prisoner's Dilemma</i>	12
2.3	Example: <i>Bach or Stravinsky?</i>	16
2.4	Example: <i>Matching Pennies</i>	17
2.5	Example: the <i>Stag Hunt</i>	18
2.6	Nash equilibrium	19
2.7	Examples of Nash equilibrium	24
2.8	Best response functions	33
2.9	Dominated actions	43
2.10	Equilibrium in a single population: symmetric games and symmetric equilibria	49
<i>Prerequisite:</i> Chapter 1.		

2.1 Strategic games

A STRATEGIC GAME is a model of interacting decision-makers. In recognition of the interaction, we refer to the decision-makers as *players*. Each player has a set of possible *actions*. The model captures interaction between the players by allowing each player to be affected by the actions of *all* players, not only her own action. Specifically, each player has *preferences* about the action *profile*—the list of all the players' actions. (See Section 17.4, in the mathematical appendix, for a discussion of profiles.)

More precisely, a strategic game is defined as follows. (The qualification “with ordinal preferences” distinguishes this notion of a strategic game from a more general notion studied in Chapter 4.)

► DEFINITION 11.1 (*Strategic game with ordinal preferences*) A **strategic game** (with ordinal preferences) consists of

- a set of **players**
- for each player, a set of **actions**
- for each player, **preferences** over the set of action profiles.

A very wide range of situations may be modeled as strategic games. For example, the players may be firms, the actions prices, and the preferences a reflection of the firms' profits. Or the players may be candidates for political office, the actions campaign expenditures, and the preferences a reflection of the candidates' probabilities of winning. Or the players may be animals fighting over some prey, the actions concession times, and the preferences a reflection of whether an animal wins or loses. In this chapter I describe some simple games designed to capture fundamental conflicts present in a variety of situations. The next chapter is devoted to more detailed applications to specific phenomena.

As in the model of rational choice by a single decision-maker (Section 1.2), it is frequently convenient to specify the players' preferences by giving *payoff functions* that represent them. Bear in mind that these payoffs have only *ordinal* significance. If a player's payoffs to the action profiles a , b , and c are 1, 2, and 10, for example, the only conclusion we can draw is that the player prefers c to b and b to a ; the numbers do *not* imply that the player's preference between c and b is stronger than her preference between a and b .

Time is absent from the model. The idea is that each player chooses her action once and for all, and the players choose their actions "simultaneously" in the sense that no player is informed, when she chooses her action, of the action chosen by any other player. (For this reason, a strategic game is sometimes referred to as a "simultaneous move game".) Nevertheless, an action may involve activities that extend over time, and may take into account an unlimited number of contingencies. An action might specify, for example, "if company X 's stock falls below \$10, buy 100 shares; otherwise, do not buy any shares". (For this reason, an action is sometimes called a "strategy".) However, the fact that time is absent from the model means that when analyzing a situation as a strategic game, we abstract from the complications that may arise if a player is allowed to change her plan as events unfold: we assume that actions are chosen once and for all.

2.2 Example: the *Prisoner's Dilemma*

One of the most well-known strategic games is the *Prisoner's Dilemma*. Its name comes from a story involving suspects in a crime; its importance comes from the huge variety of situations in which the participants face incentives similar to those faced by the suspects in the story.

- ◆ EXAMPLE 12.1 (*Prisoner's Dilemma*) Two suspects in a major crime are held in separate cells. There is enough evidence to convict each of them of a minor offense, but not enough evidence to convict either of them of the major crime unless one of them acts as an informer against the other (finks). If they both stay quiet, each will be convicted of the minor offense and spend one year in prison. If one and only one of them finks, she will be freed and used as a witness against the other, who will spend four years in prison. If they both fink, each will spend three years in prison.

This situation may be modeled as a strategic game:

Players The two suspects.

Actions Each player's set of actions is $\{Quiet, Fink\}$.

Preferences Suspect 1's ordering of the action profiles, from best to worst, is $(Fink, Quiet)$ (she finks and suspect 2 remains quiet, so she is freed), $(Quiet, Quiet)$ (she gets one year in prison), $(Fink, Fink)$ (she gets three years in prison), $(Quiet, Fink)$ (she gets four years in prison). Suspect 2's ordering is $(Quiet, Fink)$, $(Quiet, Quiet)$, $(Fink, Fink)$, $(Fink, Quiet)$.

We can represent the game compactly in a table. First choose payoff functions that represent the suspects' preference orderings. For suspect 1 we need a function u_1 for which

$$u_1(Fink, Quiet) > u_1(Quiet, Quiet) > u_1(Fink, Fink) > u_1(Quiet, Fink).$$

A simple specification is $u_1(Fink, Quiet) = 3$, $u_1(Quiet, Quiet) = 2$, $u_1(Fink, Fink) = 1$, and $u_1(Quiet, Fink) = 0$. For suspect 2 we can similarly choose the function u_2 for which $u_2(Quiet, Fink) = 3$, $u_2(Quiet, Quiet) = 2$, $u_2(Fink, Fink) = 1$, and $u_2(Fink, Quiet) = 0$. Using these representations, the game is illustrated in Figure 13.1. In this figure the two rows correspond to the two possible actions of player 1, the two columns correspond to the two possible actions of player 2, and the numbers in each box are the players' payoffs to the action profile to which the box corresponds, with player 1's payoff listed first.

		Suspect 2	
		Quiet	Fink
Suspect 1	Quiet	2, 2	0, 3
	Fink	3, 0	1, 1

Figure 13.1 The Prisoner's Dilemma (Example 12.1).

The *Prisoner's Dilemma* models a situation in which there are gains from cooperation (each player prefers that both players choose *Quiet* than they both choose *Fink*) but each player has an incentive to "free ride" (choose *Fink*) whatever the other player does. The game is important not because we are interested in understanding the incentives for prisoners to confess, but because many other situations have similar structures. Whenever each of two players has two actions, say C (corresponding to *Quiet*) and D (corresponding to *Fink*), player 1 prefers (D, C) to (C, C) to (D, D) to (C, D) , and player 2 prefers (C, D) to (C, C) to (D, D) to (D, C) , the *Prisoner's Dilemma* models the situation that the players face. Some examples follow.

2.2.1 Working on a joint project

You are working with a friend on a joint project. Each of you can either work hard or goof off. If your friend works hard then you prefer to goof off (the outcome of

the project would be better if you worked hard too, but the increment in its value to you is not worth the extra effort). You prefer the outcome of your both working hard to the outcome of your both goofing off (in which case nothing gets accomplished), and the worst outcome for you is that you work hard and your friend goofs off (you hate to be “exploited”). If your friend has the same preferences then the game that models the situation you face is given in Figure 14.1, which, as you can see, differs from the *Prisoner’s Dilemma* only in the names of the actions.

	<i>Work hard</i>	<i>Goof off</i>
<i>Work hard</i>	2, 2	0, 3
<i>Goof off</i>	3, 0	1, 1

Figure 14.1 Working on a joint project.

I am *not* claiming that a situation in which two people pursue a joint project *necessarily* has the structure of the *Prisoner’s Dilemma*, only that the players’ preferences in such a situation *may* be the same as in the *Prisoner’s Dilemma*! If, for example, each person prefers to work hard than to goof off when the other person works hard, then the *Prisoner’s Dilemma* does *not* model the situation: the players’ preferences are different from those given in Figure 14.1.

- ② EXERCISE 14.1 (Working on a joint project) Formulate a strategic game that models a situation in which two people work on a joint project in the case that their preferences are the same as those in the game in Figure 14.1 except that each person prefers to work hard than to goof off when the other person works hard. Present your game in a table like the one in Figure 14.1.

2.2.2 Duopoly

In a simple model of a duopoly, two firms produce the same good, for which each firm charges either a low price or a high price. Each firm wants to achieve the highest possible profit. If both firms choose *High* then each earns a profit of \$1000. If one firm chooses *High* and the other chooses *Low* then the firm choosing *High* obtains no customers and makes a loss of \$200, whereas the firm choosing *Low* earns a profit of \$1200 (its unit profit is low, but its volume is high). If both firms choose *Low* then each earns a profit of \$600. Each firm cares only about its profit, so we can represent its preferences by the profit it obtains, yielding the game in Figure 14.2.

	<i>High</i>	<i>Low</i>
<i>High</i>	1000, 1000	−200, 1200
<i>Low</i>	1200, −200	600, 600

Figure 14.2 A simple model of a price-setting duopoly.

Bearing in mind that what matters are the players’ preferences, not the partic-

ular payoff functions that we use to represent them, we see that this game, like the previous one, differs from the *Prisoner's Dilemma* only in the names of the actions. The action *High* plays the role of *Quiet*, and the action *Low* plays the role of *Fink*; firm 1 prefers $(Low, High)$ to $(High, High)$ to (Low, Low) to $(High, Low)$, and firm 2 prefers $(High, Low)$ to $(High, High)$ to (Low, Low) to $(Low, High)$.

As in the previous example, I do not claim that the incentives in a duopoly are necessarily those in the *Prisoner's Dilemma*; different assumptions about the relative sizes of the profits in the four cases generate a different game. Further, in this case one of the abstractions incorporated into the model—that each firm has only two prices to choose between—may not be harmless; if the firms may choose among many prices then the structure of the interaction may change. (A richer model is studied in Section 3.2.)

2.2.3 *The arms race*

Under some assumptions about the countries' preferences, an arms race can be modeled as the *Prisoner's Dilemma*. (The *Prisoner's Dilemma* was first studied in the early 1950s, when the USA and USSR were involved in a nuclear arms race, so you might suspect that US nuclear strategy was influenced by game theory; the evidence suggests that it was not.) Assume that each country can build an arsenal of nuclear bombs, or can refrain from doing so. Assume also that each country's favorite outcome is that it has bombs and the other country does not; the next best outcome is that neither country has any bombs; the next best outcome is that both countries have bombs (what matters is relative strength, and bombs are costly to build); and the worst outcome is that only the other country has bombs. In this case the situation is modeled by the *Prisoner's Dilemma*, in which the action *Don't build bombs* corresponds to *Quiet* in Figure 13.1 and the action *Build bombs* corresponds to *Fink*. However, once again the assumptions about preferences necessary for the *Prisoner's Dilemma* to model the situation may not be satisfied: a country may prefer *not* to build bombs if the other country does not, for example (bomb-building may be very costly), in which case the situation is modeled by a different game.

2.2.4 *Common property*

Two farmers are deciding how much to allow their sheep to graze on the village common. Each farmer prefers that her sheep graze a lot than a little, regardless of the other farmer's action, but prefers that both farmers' sheep graze a little than both farmers' sheep graze a lot (in which case the common is ruined for future use). Under these assumptions the game is the *Prisoner's Dilemma*. (A richer model is studied in Section 3.1.5.)

2.2.5 Other situations modeled as the Prisoner's Dilemma

A huge number of other situations have been modeled as the *Prisoner's Dilemma*, from mating hermaphroditic fish to tariff wars between countries.

- ? EXERCISE 16.1 (Hermaphroditic fish) Members of some species of hermaphroditic fish choose, in each mating encounter, whether to play the role of a male or a female. Each fish has a preferred role, which uses up fewer resources and hence allows more future mating. A fish obtains a payoff of H if it mates in its preferred role and L if it mates in the other role, where $H > L$. (Payoffs are measured in terms of number of offspring, which fish are evolved to maximize.) Consider an encounter between two fish whose preferred roles are the same. Each fish has two possible actions: mate in either role, and insist on its preferred role. If both fish offer to mate in either role, the roles are assigned randomly, and each fish's payoff is $\frac{1}{2}(H + L)$ (the average of H and L). If each fish insists on its preferred role, the fish do not mate; each goes off in search of another partner, and obtains the payoff S . The higher the chance of meeting another partner, the larger is S . Formulate this situation as a strategic game and determine the range of values of S , for any given values of H and L , for which the game differs from the *Prisoner's Dilemma* only in the names of the actions.

2.3 Example: Bach or Stravinsky?

In the *Prisoner's Dilemma* the main issue is whether or not the players will cooperate (choose *Quiet*). In the following game the players agree that it is better to cooperate than not to cooperate, but disagree about the best outcome.

- ◆ EXAMPLE 16.2 (*Bach or Stravinsky?*) Two people wish to go out together. Two concerts are available: one of music by Bach, and one of music by Stravinsky. One person prefers Bach and the other prefers Stravinsky. If they go to different concerts, each of them is equally unhappy listening to the music of either composer.

We may model this situation as the two-player strategic game in Figure 16.1, in which the person who prefers Bach chooses a row and the person who prefers Stravinsky chooses a column.

	<i>Bach</i>	<i>Stravinsky</i>
<i>Bach</i>	2, 1	0, 0
<i>Stravinsky</i>	0, 0	1, 2

Figure 16.1 *Bach or Stravinsky?* (BoS) (Example 16.2).

This game is also referred to as the “Battle of the Sexes” (though the conflict it models surely occurs no more frequently between people of the opposite sex than it does between people of the same sex). I call the game *BoS*, an acronym that fits both names. (I assume that each player is indifferent between listening to Bach and listening to Stravinsky when she is alone only for consistency with

the standard specification of the game. As we shall see, the analysis of the game remains the same in the absence of this indifference.)

Like the *Prisoner's Dilemma*, *BoS* models a wide variety of situations. Consider, for example, two officials of a political party deciding the stand to take on an issue. Suppose that they disagree about the best stand, but are both better off if they take the same stand than if they take different stands; both cases in which they take different stands, in which case voters do not know what to think, are equally bad. Then *BoS* captures the situation they face. Or consider two merging firms that currently use different computer technologies. As two divisions of a single firm they will both be better off if they both use the same technology; each firm prefers that the common technology be the one it used in the past. *BoS* models the choices the firms face.

2.4 Example: Matching Pennies

Aspects of both conflict and cooperation are present in both the *Prisoner's Dilemma* and *BoS*. The next game is purely conflictual.

- ◆ EXAMPLE 17.1 (*Matching Pennies*) Two people choose, simultaneously, whether to show the Head or the Tail of a coin. If they show the same side, person 2 pays person 1 a dollar; if they show different sides, person 1 pays person 2 a dollar. Each person cares only about the amount of money she receives, and (naturally!) prefers to receive more than less. A strategic game that models this situation is shown in Figure 17.1. (In this representation of the players' preferences, the payoffs are equal to the amounts of money involved. We could equally well work with another representation—for example, 2 could replace each 1, and 1 could replace each -1 .)

	<i>Head</i>	<i>Tail</i>
<i>Head</i>	1, -1	-1, 1
<i>Tail</i>	-1, 1	1, -1

Figure 17.1 *Matching Pennies* (Example 17.1).

In this game the players' interests are diametrically opposed (such a game is called "strictly competitive"): player 1 wants to take the same action as the other player, whereas player 2 wants to take the opposite action.

This game may, for example, model the choices of appearances for new products by an established producer and a new firm in a market of fixed size. Suppose that each firm can choose one of two different appearances for the product. The established producer prefers the newcomer's product to look different from its own (so that its customers will not be tempted to buy the newcomer's product), whereas the newcomer prefers that the products look alike. Or the game could

model a relationship between two people in which one person wants to be like the other, whereas the other wants to be different.

- ⑦ EXERCISE 18.1 (Games without conflict) Give some examples of two-player strategic games in which each player has two actions and the players have the same preferences, so that there is no conflict between their interests. (Present your games as tables like the one in Figure 17.1.)

2.5 Example: the Stag Hunt

A sentence in *Discourse on the origin and foundations of inequality among men* (1755) by the philosopher Jean-Jacques Rousseau discusses a group of hunters who wish to catch a stag. They will succeed if they all remain sufficiently attentive, but each is tempted to desert her post and catch a hare. One interpretation of the sentence is that the interaction between the hunters may be modeled as the following strategic game.

- ◆ EXAMPLE 18.2 (*Stag Hunt*) Each of a group of hunters has two options: she may remain attentive to the pursuit of a stag, or catch a hare. If all hunters pursue the stag, they catch it and share it equally; if any hunter devotes her energy to catching a hare, the stag escapes, and the hare belongs to the defecting hunter alone. Each hunter prefers a share of the stag to a hare.

The strategic game that corresponds to this specification is:

Players The hunters.

Actions Each player's set of actions is {*Stag*, *Hare*}.

Preferences For each player, the action profile in which all players choose *Stag* (resulting in her obtaining a share of the stag) is ranked highest, followed by any profile in which she chooses *Hare* (resulting in her obtaining a hare), followed by any profile in which she chooses *Stag* and one or more of the other players chooses *Hare* (resulting in her leaving empty-handed).

Like other games with many players, this game cannot easily be presented in a table like that in Figure 17.1. For the case in which there are two hunters, the game is shown in Figure 18.1.

	<i>Stag</i>	<i>Hare</i>
<i>Stag</i>	2, 2	0, 1
<i>Hare</i>	1, 0	1, 1

Figure 18.1 The *Stag Hunt* (Example 18.2) for the case of two hunters.

The variant of the two-player *Stag Hunt* shown in Figure 19.1 has been suggested as an alternative to the *Prisoner's Dilemma* as a model of an arms race, or, more generally, of the "security dilemma" faced by a pair of countries. The game differs from the *Prisoner's Dilemma* in that a country prefers the outcome in which

both countries refrain from arming themselves to the one in which it alone arms itself: the cost of arming outweighs the benefit if the other country does not arm itself.

	<i>Refrain</i>	<i>Arm</i>
<i>Refrain</i>	3, 3	0, 2
<i>Arm</i>	2, 0	1, 1

Figure 19.1 A variant of the two-player *Stag Hunt* that models the “security dilemma”.

2.6 Nash equilibrium

What actions will be chosen by the players in a strategic game? We wish to assume, as in the theory of a rational decision-maker (Section 1.2), that each player chooses the best available action. In a game, the best action for any given player depends, in general, on the other players’ actions. So when choosing an action a player must have in mind the actions the other players will choose. That is, she must form a *belief* about the other players’ actions.

On what basis can such a belief be formed? The assumption underlying the analysis in this chapter and the next two chapters is that each player’s belief is derived from her past experience playing the game, and that this experience is sufficiently extensive that she *knows* how her opponents will behave. No one tells her the actions her opponents will choose, but her previous involvement in the game leads her to be sure of these actions. (The question of *how* a player’s experience can lead her to the correct beliefs about the other players’ actions is addressed briefly in Section 4.9.)

Although we assume that each player has experience playing the game, we assume that she views each play of the game in isolation. She does not become familiar with the behavior of specific opponents and consequently does not condition her action on the opponent she faces; nor does she expect her current action to affect the other players’ future behavior.

It is helpful to think of the following idealized circumstances. For each player in the game there is a population of many decision-makers who may, on any occasion, take that player’s role. In each play of the game, players are selected randomly, one from each population. Thus each player engages in the game repeatedly, against ever-varying opponents. Her experience leads her to beliefs about the actions of “typical” opponents, not any specific set of opponents.

As an example, think of the interaction between buyers and sellers. Buyers and sellers repeatedly interact, but to a first approximation many of the pairings may be modeled as random. In many cases a buyer transacts only once with any given seller, or interacts repeatedly but anonymously (when the seller is a large store, for example).

In summary, the solution theory we study has two components. First, each

player chooses her action according to the model of rational choice, given her belief about the other players' actions. Second, every player's belief about the other players' actions is correct. These two components are embodied in the following definition.

A *Nash equilibrium* is an action profile a^* with the property that no player i can do better by choosing an action different from a_i^* , given that every other player j adheres to a_j^* .

In the idealized setting in which the players in any given play of the game are drawn randomly from a collection of populations, a Nash equilibrium corresponds to a *steady state*. If, whenever the game is played, the action profile is the same Nash equilibrium a^* , then no player has a reason to choose any action different from her component of a^* ; there is no pressure on the action profile to change. Expressed differently, a Nash equilibrium embodies a stable "social norm": if everyone else adheres to it, no individual wishes to deviate from it.

The second component of the theory of Nash equilibrium—that the players' beliefs about each other's actions are correct—implies, in particular, that two players' beliefs about a third player's action are the same. For this reason, the condition is sometimes said to be that the players' "expectations are coordinated".

The situations to which we wish to apply the theory of Nash equilibrium do not in general correspond exactly to the idealized setting described above. For example, in some cases the players do not have much experience with the game; in others they do not view each play of the game in isolation. Whether or not the notion of Nash equilibrium is appropriate in any given situation is a matter of judgment. In some cases, a poor fit with the idealized setting may be mitigated by other considerations. For example, inexperienced players may be able to draw conclusions about their opponents' likely actions from their experience in other situations, or from other sources. (One aspect of such reasoning is discussed in the box on page 30). Ultimately, the test of the appropriateness of the notion of Nash equilibrium is whether it gives us insights into the problem at hand.

With the aid of an additional piece of notation, we can state the definition of a Nash equilibrium precisely. Let a be an action profile, in which the action of each player i is a_i . Let a'_i be any action of player i (either equal to a_i , or different from it). Then (a'_i, a_{-i}) denotes the action profile in which every player j *except* i chooses her action a_j as specified by a , whereas player i chooses a'_i . (The $-i$ subscript on a stands for "except i ".) That is, (a'_i, a_{-i}) is the action profile in which all the players other than i adhere to a while i "deviates" to a'_i . (If $a'_i = a_i$ then of course $(a'_i, a_{-i}) = (a_i, a_{-i}) = a$.) If there are three players, for example, then (a'_2, a_{-2}) is the action profile in which players 1 and 3 adhere to a (player 1 chooses a_1 , player 3 chooses a_3) and player 2 deviates to a'_2 .

Using this notation, we can restate the condition for an action profile a^* to be a Nash equilibrium: no player i has any action a_i for which she prefers (a_i, a_{-i}^*) to a^* . Equivalently, for every player i and every action a_i of player i , the action profile a^* is at least as good for player i as the action profile (a_i, a_{-i}^*) .

JOHN F. NASH, JR.



A few of the ideas of John F. Nash Jr., developed while he was a graduate student at Princeton from 1948 to 1950, transformed game theory. Nash was born in 1928 in Bluefield, West Virginia, USA, where he grew up. He was an undergraduate mathematics major at Carnegie Institute of Technology from 1945 to 1948. In 1948 he obtained both a B.S. and an M.S., and began graduate work in the Department of Mathematics at Princeton University. (One of his letters of recommendation, from a professor at Carnegie Institute of Technology, was a single sentence: “This man is a genius” (Kuhn et al. 1995, 282).) A paper containing the main result of his thesis was sub-

mitted to the *Proceedings of the National Academy of Sciences* in November 1949, fourteen months after he started his graduate work. (“A fine goal to set ... graduate students”, to quote Harold Kuhn! (See Kuhn et al. 1995, 282.)) He completed his PhD the following year, graduating on his 22nd birthday. His thesis, 28 pages in length, introduces the equilibrium notion now known as “Nash equilibrium” and delineates a class of strategic games that have Nash equilibria (Proposition 117.1 in this book). The notion of Nash equilibrium vastly expanded the scope of game theory, which had previously focussed on two-player “strictly competitive” games (in which the players’ interests are directly opposed). While a graduate student at Princeton, Nash also wrote the seminal paper in bargaining theory, Nash (1950b) (the ideas of which originated in an elective class in international economics he took as an undergraduate). He went on to take an academic position in the Department of Mathematics at MIT, where he produced “a remarkable series of papers” (Milnor 1995, 15); he has been described as “one of the most original mathematical minds of [the twentieth] century” (Kuhn 1996). He shared the 1994 Nobel prize in economics with the game theorists John C. Harsanyi and Reinhard Selten.

- DEFINITION 21.1 (*Nash equilibrium of strategic game with ordinal preferences*) The action profile a^* in a strategic game with ordinal preferences is a **Nash equilibrium** if, for every player i and every action a_i of player i , a^* is at least as good according to player i ’s preferences as the action profile (a_i, a_{-i}^*) in which player i chooses a_i while every other player j chooses a_j^* . Equivalently, for every player i ,

$$u_i(a^*) \geq u_i(a_i, a_{-i}^*) \text{ for every action } a_i \text{ of player } i, \quad (21.2)$$

where u_i is a payoff function that represents player i ’s preferences.

This definition implies neither that a strategic game necessarily has a Nash equilibrium, nor that it has at most one. Examples in the next section show that

some games have a single Nash equilibrium, some possess no Nash equilibrium, and others have many Nash equilibria.

The definition of a Nash equilibrium is designed to model a steady state among experienced players. An alternative approach to understanding players' actions in strategic games assumes that the players know each others' preferences, and considers what each player can deduce about the other players' actions from their rationality and their knowledge of each other's rationality. This approach is studied in Chapter 12. For many games, it leads to a conclusion different from that of Nash equilibrium. For games in which the conclusion is the same the approach offers us an alternative interpretation of a Nash equilibrium, as the outcome of rational calculations by players who do not necessarily have any experience playing the game.

STUDYING NASH EQUILIBRIUM EXPERIMENTALLY

The theory of strategic games lends itself to experimental study: arranging for subjects to play games and observing their choices is relatively straightforward. A few years after game theory was launched by von Neumann and Morgenstern's (1944) book, reports of laboratory experiments began to appear. Subsequently a huge number of experiments have been conducted, illuminating many issues relevant to the theory. I discuss selected experimental evidence throughout the book.

The theory of Nash equilibrium, as we have seen, has two components: the players act in accordance with the theory of rational choice, given their beliefs about the other players' actions, and these beliefs are correct. If every subject understands the game she is playing and faces incentives that correspond to the preferences of the player whose role she is taking, then a divergence between the observed outcome and a Nash equilibrium can be blamed on a failure of one or both of these two components. Experimental evidence has the potential of indicating the types of games for which the theory works well and, for those in which the theory does not work well, of pointing to the faulty component and giving us hints about the characteristics of a better theory. In designing an experiment that cleanly tests the theory, however, we need to confront several issues.

The model of rational choice takes preferences as given. Thus to test the theory of Nash equilibrium experimentally, we need to ensure that each subject's preferences are those of the player whose role she is taking in the game we are examining. The standard way of inducing the appropriate preferences is to pay each subject an amount of money directly related to the payoff given by a payoff function that represents the preferences of the player whose role the subject is taking. Such remuneration works if each subject likes money and cares only about the amount of money she receives, ignoring the amounts received by her opponents. The assumption that people like receiving money is reasonable in many cultures, but the assumption that people care only about their own monetary rewards—are "selfish"—may, in some contexts at least, not be reasonable. Unless we check

whether our subjects are selfish in the context of our experiment, we will jointly test two hypotheses: that humans are selfish—a hypothesis not part of game theory—and that the notion of Nash equilibrium models their behavior. In some cases we may indeed wish to test these hypotheses jointly. But in order to test the theory of Nash equilibrium alone we need to ensure that we induce the preferences we wish to study.

Assuming that better decisions require more effort, we need also to ensure that each subject finds it worthwhile to put in the extra effort required to obtain a higher payoff. If we rely on monetary payments to provide incentives, the amount of money a subject can obtain must be sufficiently sensitive to the quality of her decisions to compensate her for the effort she expends (paying a flat fee, for example, is inappropriate). In some cases, monetary payments may not be necessary: under some circumstances, subjects drawn from a highly competitive culture like that of the USA may be sufficiently motivated by the possibility of obtaining a high score, even if that score does not translate into a monetary payoff.

The notion of Nash equilibrium models action profiles compatible with steady states. Thus to study the theory experimentally we need to collect observations of subjects' behavior when they have experience playing the game. But they should not have obtained that experience while knowingly facing the same opponents repeatedly, for the theory assumes that the players consider each play of the game in isolation, not as part of an ongoing relationship. One option is to have each subject play the game against many different opponents, gaining experience about how the other subjects on average play the game, but not about the choices of any other given player. Another option is to describe the game in terms that relate to a situation in which the subjects already have experience. A difficulty with this second approach is that the description we give may connote more than simply the payoff numbers of our game. If we describe the *Prisoner's Dilemma* in terms of cooperation on a joint project, for example, a subject may be biased toward choosing the action she has found appropriate when involved in joint projects, even if the structures of those interactions were significantly different from that of the *Prisoner's Dilemma*. As she plays the experimental game repeatedly she may come to appreciate how it differs from the games in which she has been involved previously, but her biases may disappear only slowly.

Whatever route we take to collect data on the choices of subjects experienced in playing the game, we confront a difficult issue: how do we know when the outcome has converged? Nash's theory concerns only equilibria; it has nothing to say about the path players' choices will take on the way to an equilibrium, and so gives us no guide as to whether 10, 100, or 1,000 plays of the game are enough to give a chance for the subjects' expectations to become coordinated.

Finally, we can expect the theory of Nash equilibrium to correspond to reality only approximately: like all useful theories, it definitely is not *exactly* correct. How do we tell whether the data are close enough to the theory to support it? One possibility is to compare the theory of Nash equilibrium with some other theory. But for many games there is no obvious alternative theory—and certainly not one with

the generality of Nash equilibrium. Statistical tests can sometimes aid in deciding whether the data is consistent with the theory, though ultimately we remain the judge of whether or not our observations persuade us that the theory enhances our understanding of human behavior in the game.

2.7 Examples of Nash equilibrium

2.7.1 Prisoner's Dilemma

By examining the four possible pairs of actions in the *Prisoner's Dilemma* (reproduced in Figure 24.1), we see that $(Fink, Fink)$ is the unique Nash equilibrium.

	<i>Quiet</i>	<i>Fink</i>
<i>Quiet</i>	2, 2	0, 3
<i>Fink</i>	3, 0	1, 1

Figure 24.1 The *Prisoner's Dilemma*.

The action pair $(Fink, Fink)$ is a Nash equilibrium because (i) given that player 2 chooses *Fink*, player 1 is better off choosing *Fink* than *Quiet* (looking at the right column of the table we see that *Fink* yields player 1 a payoff of 1 whereas *Quiet* yields her a payoff of 0), and (ii) given that player 1 chooses *Fink*, player 2 is better off choosing *Fink* than *Quiet* (looking at the bottom row of the table we see that *Fink* yields player 2 a payoff of 1 whereas *Quiet* yields her a payoff of 0).

No other action profile is a Nash equilibrium:

- $(Quiet, Quiet)$ does not satisfy (21.2) because when player 2 chooses *Quiet*, player 1's payoff to *Fink* exceeds her payoff to *Quiet* (look at the first components of the entries in the left column of the table). (Further, when player 1 chooses *Quiet*, player 2's payoff to *Fink* exceeds her payoff to *Quiet*: player 2, as well as player 1, wants to deviate. To show that a pair of actions is not a Nash equilibrium, however, it is not necessary to study player 2's decision once we have established that player 1 wants to deviate: it is enough to show that *one* player wishes to deviate to show that a pair of actions is not a Nash equilibrium.)
- $(Fink, Quiet)$ does not satisfy (21.2) because when player 1 chooses *Fink*, player 2's payoff to *Fink* exceeds her payoff to *Quiet* (look at the second components of the entries in the bottom row of the table).
- $(Quiet, Fink)$ does not satisfy (21.2) because when player 2 chooses *Fink*, player 1's payoff to *Fink* exceeds her payoff to *Quiet* (look at the first components of the entries in the right column of the table).

In summary, in the only Nash equilibrium of the *Prisoner's Dilemma* both players choose *Fink*. In particular, the incentive to free ride eliminates the possibility that the mutually desirable outcome (*Quiet, Quiet*) occurs. In the other situations discussed in Section 2.2 that may be modeled as the *Prisoner's Dilemma*, the outcomes predicted by the notion of Nash equilibrium are thus as follows: both people goof off when working on a joint project; both duopolists charge a low price; both countries build bombs; both farmers graze their sheep a lot. (The overgrazing of a common thus predicted is sometimes called the "tragedy of the commons". The intuition that some of these dismal outcomes may be avoided if the same pair of people play the game repeatedly is explored in Chapter 14.)

In the *Prisoner's Dilemma*, the Nash equilibrium action of each player (*Fink*) is the best action for each player not only if the other player chooses her equilibrium action (*Fink*), but also if she chooses her other action (*Quiet*). The action pair (*Fink, Fink*) is a Nash equilibrium because if a player believes that her opponent will choose *Fink* then it is optimal for her to choose *Fink*. But in fact it is optimal for a player to choose *Fink* regardless of the action she expects her opponent to choose. In most of the games we study, a player's Nash equilibrium action does not satisfy this condition: the action is optimal if the other players choose their Nash equilibrium actions, but some other action is optimal if the other players choose non-equilibrium actions.

- ⊛ EXERCISE 25.1 (Variant of *Prisoner's Dilemma* with altruistic preferences) Each of two players has two possible actions, *Quiet* and *Fink*; each action pair results in the players' receiving amounts of *money* equal to the numbers corresponding to that action pair in Figure 24.1. (For example, if player 1 chooses *Quiet* and player 2 chooses *Fink*, then player 1 receives nothing, whereas player 2 receives \$3.) The players are not "selfish"; rather, the preferences of each player i are represented by the payoff function $m_i(a) + \alpha m_j(a)$, where $m_i(a)$ is the amount of money received by player i when the action profile is a , j is the other player, and α is a given non-negative number. Player 1's payoff to the action pair (*Quiet, Quiet*), for example, is $2 + 2\alpha$.
- Formulate a strategic game that models this situation in the case $\alpha = 1$. Is this game the *Prisoner's Dilemma*?
 - Find the range of values of α for which the resulting game is the *Prisoner's Dilemma*. For values of α for which the game is not the *Prisoner's Dilemma*, find its Nash equilibria.
- ⊛ EXERCISE 25.2 (Selfish and altruistic social behavior) Two people enter a bus. Two adjacent cramped seats are free. Each person must decide whether to sit or stand. Sitting alone is more comfortable than sitting next to the other person, which is more comfortable than standing.
- Suppose that each person cares only about her own comfort. Model the situation as a strategic game. Is this game the *Prisoner's Dilemma*? Find its Nash equilibrium (equilibria?).

- b. Suppose that each person is altruistic, ranking the outcomes according to the *other* person's comfort, but, out of politeness, prefers to stand than to sit if the other person stands. Model the situation as a strategic game. Is this game the *Prisoner's Dilemma*? Find its Nash equilibrium (equilibria?).
- c. Compare the people's comfort in the equilibria of the two games.

EXPERIMENTAL EVIDENCE ON THE *Prisoner's Dilemma*

The *Prisoner's Dilemma* has attracted a great deal of attention by economists, psychologists, sociologists, and biologists. A huge number of experiments have been conducted with the aim of discovering how people behave when playing the game. Almost all these experiments involve each subject's playing the game repeatedly against an unchanging opponent, a situation that calls for an analysis significantly different from the one in this chapter (see Chapter 14).

The evidence on the outcome of isolated plays of the game is inconclusive. No experiment of which I am aware carefully induces the appropriate preferences and is specifically designed to elicit a steady state action profile (see the box on page 22). Thus in each case the choice of *Quiet* by a player could indicate that she is not "selfish" or that she is not experienced in playing the game, rather than providing evidence against the notion of Nash equilibrium.

In two experiments with very low payoffs, each subject played the game a small number of times against different opponents; between 50% and 94% of subjects chose *Fink*, depending on the relative sizes of the payoffs and some details of the design (Rapoport, Guyer, and Gordon 1976, 135–137, 211–213, and 223–226). In a more recent experiment, 78% of subjects chose *Fink* in the last 10 of 20 rounds of play against different opponents (Cooper, DeJong, Forsythe, and Ross 1996). In face-to-face games in which communication is allowed, the incidence of the choice of *Fink* tends to be lower: from 29% to 70% depending on the nature of the communication allowed (Deutsch 1958, and Frank, Gilovich, and Regan 1993, 163–167). (In all these experiments, the subjects were college students in the USA or Canada.)

One source of the variation in the results seems to be that some designs induce preferences that differ from those of the *Prisoner's Dilemma*; no clear answer emerges to the question of whether the notion of Nash equilibrium is relevant to the *Prisoner's Dilemma*. If, nevertheless, one interprets the evidence as showing that some subjects in the *Prisoner's Dilemma* systematically choose *Quiet* rather than *Fink*, one must fault the rational choice component of Nash equilibrium, not the coordinated expectations component. Why? Because, as noted in the text, *Fink* is optimal *no matter* what a player thinks her opponent will choose, so that any model in which the players act according to the model of rational choice, whether or not their expectations are coordinated, predicts that each player chooses *Fink*.

2.7.2 *BoS*

To find the Nash equilibria of *BoS* (Figure 16.1), we can examine each pair of actions in turn:

- (*Bach, Bach*): If player 1 switches to *Stravinsky* then her payoff decreases from 2 to 0; if player 2 switches to *Stravinsky* then her payoff decreases from 1 to 0. Thus a deviation by either player decreases her payoff. Thus (*Bach, Bach*) is a Nash equilibrium.
- (*Bach, Stravinsky*): If player 1 switches to *Stravinsky* then her payoff increases from 0 to 1. Thus (*Bach, Stravinsky*) is not a Nash equilibrium. (Player 2 can increase her payoff by deviating, too, but to show the pair is not a Nash equilibrium it suffices to show that one player can increase her payoff by deviating.)
- (*Stravinsky, Bach*): If player 1 switches to *Bach* then her payoff increases from 0 to 2. Thus (*Stravinsky, Bach*) is not a Nash equilibrium.
- (*Stravinsky, Stravinsky*): If player 1 switches to *Bach* then her payoff decreases from 1 to 0; if player 2 switches to *Bach* then her payoff decreases from 2 to 0. Thus a deviation by either player decreases her payoff. Thus (*Stravinsky, Stravinsky*) is a Nash equilibrium.

We conclude that the game has two Nash equilibria: (*Bach, Bach*) and (*Stravinsky, Stravinsky*). That is, both of these outcomes are compatible with a steady state; both outcomes are stable social norms. If, in every encounter, both players choose *Bach*, then no player has an incentive to deviate; if, in every encounter, both players choose *Stravinsky*, then no player has an incentive to deviate. If we use the game to model the choices of men when matched with women, for example, then the notion of Nash equilibrium shows that two social norms are stable: both players choose the action associated with the outcome preferred by women, and both players choose the action associated with the outcome preferred by men.

2.7.3 *Matching Pennies*

By checking each of the four pairs of actions in *Matching Pennies* (Figure 17.1) we see that the game has no Nash equilibrium. For the pairs of actions (*Head, Head*) and (*Tail, Tail*), player 2 is better off deviating; for the pairs of actions (*Head, Tail*) and (*Tail, Head*), player 1 is better off deviating. Thus for this game the notion of Nash equilibrium isolates no steady state. In Chapter 4 we return to this game; an extension of the notion of a Nash equilibrium gives us an understanding of the likely outcome.

2.7.4 *The Stag Hunt*

Inspection of Figure 18.1 shows that the two-player *Stag Hunt* has two Nash equilibria: (*Stag, Stag*) and (*Hare, Hare*). If one player remains attentive to the pursuit

of the stag, then the other player prefers to remain attentive; if one player chases a hare, the other one prefers to chase a hare (she cannot catch a stag alone). (The equilibria of the variant of the game in Figure 19.1 are analogous: $(Refrain, Refrain)$ and (Arm, Arm) .)

Unlike the Nash equilibria of *BoS*, one of these equilibria is better for both players than the other: each player prefers $(Stag, Stag)$ to $(Hare, Hare)$. This fact has no bearing on the equilibrium status of $(Hare, Hare)$, since the condition for an equilibrium is that a *single* player cannot gain by deviating, *given* the other player's behavior. Put differently, an equilibrium is immune to any *unilateral* deviation; coordinated deviations by groups of players are not contemplated. However, the existence of two equilibria raises the possibility that one equilibrium might more likely be the outcome of the game than the other. I return to this issue in Section 2.7.6.

I argue that the many-player *Stag Hunt* (Example 18.2) also has two Nash equilibria: the action profile $(Stag, \dots, Stag)$ in which every player joins in the pursuit of the stag, and the profile $(Hare, \dots, Hare)$ in which every player catches a hare.

- $(Stag, \dots, Stag)$ is a Nash equilibrium because each player prefers this profile to that in which she alone chooses *Hare*. (A player is better off remaining attentive to the pursuit of the stag than running after a hare if all the other players remain attentive.)
- $(Hare, \dots, Hare)$ is a Nash equilibrium because each player prefers this profile to that in which she alone pursues the stag. (A player is better off catching a hare than pursuing the stag if no one else pursues the stag.)
- No other profile is a Nash equilibrium, because in any other profile at least one player chooses *Stag* and at least one player chooses *Hare*, so that any player choosing *Stag* is better off switching to *Hare*. (A player is better off catching a hare than pursuing the stag if at least one other person chases a hare, since the stag can be caught only if everyone pursues it.)

? EXERCISE 28.1 (Variants of the *Stag Hunt*) Consider two variants of the n -hunter *Stag Hunt* in which only m hunters, with $2 \leq m < n$, need to pursue the stag in order to catch it. (Continue to assume that there is a single stag.) Assume that a captured stag is shared only by the hunters who catch it.

- a. Assume, as before, that each hunter prefers the fraction $1/n$ of the stag to a hare. Find the Nash equilibria of the strategic game that models this situation.
- b. Assume that each hunter prefers the fraction $1/k$ of the stag to a hare, but prefers the hare to any smaller fraction of the stag, where k is an integer with $m \leq k \leq n$. Find the Nash equilibria of the strategic game that models this situation.

The following more difficult exercise enriches the hunters' choices in the *Stag Hunt*. This extended game has been proposed as a model that captures Keynes' basic insight about the possibility of multiple economic equilibria, some of which are undesirable (Bryant 1983, 1994).

- ⊗ EXERCISE 29.1 (Extension of the *Stag Hunt*) Extend the n -hunter *Stag Hunt* by giving each hunter K (a positive integer) units of effort, which she can allocate between pursuing the stag and catching hares. Denote the effort hunter i devotes to pursuing the stag by e_i , a nonnegative integer equal to at most K . The chance that the stag is caught depends on the smallest of all the hunters' efforts, denoted $\min_j e_j$. ("A chain is as strong as its weakest link.") Hunter i 's payoff to the action profile (e_1, \dots, e_n) is $2 \min_j e_j - e_i$. (She is better off the more likely the stag is caught, and worse off the more effort she devotes to pursuing the stag, which means she catches fewer hares.) Is the action profile (e, \dots, e) , in which every hunter devotes the same effort to pursuing the stag, a Nash equilibrium for any value of e ? (What is a player's payoff to this profile? What is her payoff if she deviates to a lower or higher effort level?) Is any action profile in which not all the players' effort levels are the same a Nash equilibrium? (Consider a player whose effort exceeds the minimum effort level of all players. What happens to her payoff if she reduces her effort level to the minimum?)

2.7.5 Hawk–Dove

The game in the next exercise captures a basic feature of animal conflict.

- ⊗ EXERCISE 29.2 (Hawk–Dove) Two animals are fighting over some prey. Each can be passive or aggressive. Each prefers to be aggressive if its opponent is passive, and passive if its opponent is aggressive; given its own stance, it prefers the outcome in which its opponent is passive to that in which its opponent is aggressive. Formulate this situation as a strategic game and find its Nash equilibria.

2.7.6 A coordination game

Consider two people who wish to go out together, but who, unlike the dissidents in *BoS*, agree on the more desirable concert—say they both prefer *Bach*. A strategic game that models this situation is shown in Figure 29.1; it is an example of a *coordination game*. By examining the four action pairs, we see that the game has two Nash equilibria: $(Bach, Bach)$ and $(Stravinsky, Stravinsky)$. In particular, the action pair $(Stravinsky, Stravinsky)$ in which both people choose their less-preferred concert is a Nash equilibrium.

	<i>Bach</i>	<i>Stravinsky</i>
<i>Bach</i>	2, 2	0, 0
<i>Stravinsky</i>	0, 0	1, 1

Figure 29.1 A coordination game.

Is the equilibrium in which both people choose *Stravinsky* plausible? People who argue that the technology of Apple computers originally dominated that of IBM computers, and that the Beta format for video recording is better than VHS,

FOCAL POINTS

In games with many Nash equilibria, the theory isolates more than one pattern of behavior compatible with a steady state. In some games, some of these equilibria seem more likely to attract the players' attentions than others. To use the terminology of Schelling (1960), some equilibria are *focal*. In the coordination game in Figure 29.1, where the players agree on the more desirable Nash equilibrium and obtain the same payoff to every nonequilibrium action pair, the preferable equilibrium seems more likely to be focal (though two examples are given in the text of steady states involving the inferior equilibrium). In the variant of this game in which the two equilibria are equally good (i.e. $(2, 2)$ is replaced by $(1, 1)$), nothing in the structure of the game gives any clue as to which steady state might occur. In such a game, the names or nature of the actions, or other information, may predispose the players to one equilibrium rather than the other.

Consider, for example, voters in an election. Pre-election polls may give them information about each other's intended actions, pointing them to one of many Nash equilibria. Or consider a situation in which two players independently divide \$100 into two piles, each receiving \$10 if they choose the same divisions and nothing otherwise. The strategic game that models this situation has many Nash equilibria, in each of which both players choose the same division. But the equilibrium in which both players choose the $(\$50, \$50)$ division seems likely to command the players' attentions, possibly for esthetic reasons (it is an appealing division), and possibly because it is a steady state in an unrelated game in which the chosen division determines the players' payoffs.

The theory of Nash equilibrium is neutral about the equilibrium that will occur in a game with many equilibria. If features of the situation not modeled by the notion of a strategic game make some equilibria focal then those equilibria may be more likely to emerge as steady states, and the rate at which a steady state is reached may be higher than it otherwise would have been.

would say "yes". In both cases users had a strong interest in adopting the same standard, and one standard was better than the other; in the steady state that emerged in each case, the inferior technology was adopted by a large majority of users.

If two people played this game in a laboratory it seems likely that the outcome would be *(Bach, Bach)*. Nevertheless, *(Stravinsky, Stravinsky)* also corresponds to a steady state: if either action pair is reached, there is no reason for either player to deviate from it.

2.7.7 Provision of a public good

The model in the next exercise captures an aspect of the provision of a “public good”, like a park or a swimming pool, whose use by one person does not diminish its value to another person (at least, not until it is overcrowded). (Other aspects of public good provision are studied in Section 2.8.4.)

- EXERCISE 31.1 (Contributing to a public good) Each of n people chooses whether or not to contribute a fixed amount toward the provision of a public good. The good is provided if and only if at least k people contribute, where $2 \leq k \leq n$; if it is not provided, contributions are not refunded. Each person ranks outcomes from best to worst as follows: (i) any outcome in which the good is provided and she does not contribute, (ii) any outcome in which the good is provided and she contributes, (iii) any outcome in which the good is not provided and she does not contribute, (iv) any outcome in which the good is not provided and she contributes. Formulate this situation as a strategic game and find its Nash equilibria. (Is there a Nash equilibrium in which more than k people contribute? One in which k people contribute? One in which fewer than k people contribute? (Be careful!))

2.7.8 Strict and nonstrict equilibria

In all the Nash equilibria of the games we have studied so far a deviation by a player leads to an outcome *worse* for that player than the equilibrium outcome. The definition of Nash equilibrium (21.1), however, requires only that the outcome of a deviation be *no better* for the deviant than the equilibrium outcome. And, indeed, some games have equilibria in which a player is indifferent between her equilibrium action and some other action, given the other players' actions.

Consider the game in Figure 31.1. This game has a unique Nash equilibrium, namely (T, L) . (For every other pair of actions, one of the players is better off changing her action.) When player 2 chooses L , as she does in this equilibrium, player 1 is equally happy choosing T or B (her payoff is 1 in each case); if she deviates to B then she is no worse off than she is in the equilibrium. We say that the Nash equilibrium (T, L) is not a *strict equilibrium*.

	L	M	R
T	1, 1	1, 0	0, 1
B	1, 0	0, 1	1, 0

Figure 31.1 A game with a unique Nash equilibrium, which is not a strict equilibrium.

For a general game, an equilibrium is strict if each player's equilibrium action is *better* than all her other actions, given the other players' actions. Precisely, an action profile a^* is a **strict Nash equilibrium** if for every player i we have $u_i(a^*) > u_i(a_i, a_{-i}^*)$ for every action $a_i \neq a_i^*$ of player i . (Contrast the strict inequality in this definition with the weak inequality in (21.2).)

2.7.9 Additional examples

The following exercises are more difficult than most of the previous ones. In the first two, the number of actions of each player is arbitrary, so you cannot mechanically examine each action profile individually, as we did for games in which each player has two actions. Instead, you can consider groups of action profiles that have features in common, and show that all action profiles in any given group are or are not equilibria. Deciding how best to group the profiles into types calls for some intuition about the character of a likely equilibrium; the exercises contain suggestions on how to proceed.

- ⊗ EXERCISE 32.1 (Guessing two-thirds of the average) Each of three people announces an integer from 1 to K . If the three integers are different, the person whose integer is closest to $\frac{2}{3}$ of the average of the three integers wins \$1. If two or more integers are the same, \$1 is split equally between the people whose integer is closest to $\frac{2}{3}$ of the average integer. Is there any integer k such that the action profile (k, k, k) , in which every person announces the same integer k , is a Nash equilibrium? (If $k \geq 2$, what happens if a person announces a smaller number?) Is any other action profile a Nash equilibrium? (What is the payoff of a person whose number is the highest of the three? Can she increase this payoff by announcing a different number?)

Game theory is used widely in political science, especially in the study of elections. The game in the following exercise explores citizens' costly decisions to vote.

- ⊗ EXERCISE 32.2 (Voter participation) Two candidates, A and B , compete in an election. Of the n citizens, k support candidate A and $m (= n - k)$ support candidate B . Each citizen decides whether to vote, at a cost, for the candidate she supports, or to abstain. A citizen who abstains receives the payoff of 2 if the candidate she supports wins, 1 if this candidate ties for first place, and 0 if this candidate loses. A citizen who votes receives the payoffs $2 - c$, $1 - c$, and $-c$ in these three cases, where $0 < c < 1$.
- For $k = m = 1$, is the game the same (except for the names of the actions) as any considered so far in this chapter?
 - For $k = m$, find the set of Nash equilibria. (Is the action profile in which everyone votes a Nash equilibrium? Is there any Nash equilibrium in which the candidates tie and not everyone votes? Is there any Nash equilibrium in which one of the candidates wins by one vote? Is there any Nash equilibrium in which one of the candidates wins by two or more votes?)
 - What is the set of Nash equilibria for $k < m$?

If, when sitting in a traffic jam, you have ever thought about the time you might save if another road were built, the next exercise may lead you to think again.

Ⓣ EXERCISE 33.1 (Choosing a route) Four people must drive from A to B at the same time. Each of them must choose which route to take. Two routes are available, one via X and one via Y. (Refer to the left panel of Figure 33.1.) The roads from A to X, and from Y to B are both short and narrow; in each case, one car takes 6 minutes, and each additional car increases the travel time *per car* by 3 minutes. (If two cars drive from A to X, for example, *each car* takes 9 minutes.) The roads from A to Y, and from X to B are long and wide; on A to Y one car takes 20 minutes, and each additional car increases the travel time *per car* by 1 minute; on X to B one car takes 20 minutes, and each additional car increases the travel time *per car* by 0.9 minutes. Formulate this situation as a strategic game and find the Nash equilibria. (If all four people take one of the routes, can any of them do better by taking the other route? What if three take one route and one takes the other route, or if two take each route?)




Figure 33.1 Getting from A to B: the road networks in Exercise 33.1. The numbers beside each road are the travel times *per car* when 1, 2, 3, or 4 cars take that road.

Now suppose that a relatively short, wide road is built from X to Y, giving each person four options for travel from A to B: A-X-B, A-Y-B, A-X-Y-B, and A-Y-X-B. Assume that a person who takes A-X-Y-B travels the A-X portion at the same time as someone who takes A-X-B, and the Y-B portion at the same time as someone who takes A-Y-B. (Think of there being constant flows of traffic.) On the road between X and Y, one car takes 7 minutes and each additional car increases the travel time *per car* by 1 minute. Find the Nash equilibria in this new situation. Compare each person’s travel time with her travel time in the equilibrium before the road from X to Y was built.

2.8 Best response functions

2.8.1 Definition

We can find the Nash equilibria of a game in which each player has only a few actions by examining each action profile in turn to see if it satisfies the conditions for equilibrium. In more complicated games, it is often better to work with the players’ “best response functions”.

Consider a player, say player i . For any given actions of the players other than i , player i 's actions yield her various payoffs. We are interested in the best actions—those that yield her the highest payoff. In *BoS*, for example, *Bach* is the best action for player 1 if player 2 chooses *Bach*; *Stravinsky* is the best action for player 1 if player 2 chooses *Stravinsky*. In particular, in *BoS*, player 1 has a single best action for each action of player 2. By contrast, in the game in Figure 31.1, both T and B are best actions for player 1 if player 2 chooses L : they both yield the payoff of 1, and player 1 has no action that yields a higher payoff (in fact, she has no other action).

We denote the set of player i 's best actions when the list of the other players' actions is a_{-i} by $B_i(a_{-i})$. Thus in *BoS* we have $B_1(\text{Bach}) = \{\text{Bach}\}$ and $B_1(\text{Stravinsky}) = \{\text{Stravinsky}\}$; in the game in Figure 31.1 we have $B_1(L) = \{T, B\}$.

Precisely, we define the function B_i by

$$B_i(a_{-i}) = \{a_i \text{ in } A_i : u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \text{ for all } a'_i \text{ in } A_i\} :$$

any action in $B_i(a_{-i})$ is at least as good for player i as every other action of player i when the other players' actions are given by a_{-i} . We call B_i the **best response function** of player i .

The function B_i is *set-valued*: it associates a set of actions with any list of the other players' actions. Every member of the set $B_i(a_{-i})$ is a **best response** of player i to a_{-i} : if each of the other players adheres to a_{-i} then player i can do no better than choose a member of $B_i(a_{-i})$. In some games, like *BoS*, the set $B_i(a_{-i})$ consists of a single action for every list a_{-i} of actions of the other players: no matter what the other players do, player i has a *single* optimal action. In other games, like the one in Figure 31.1, $B_i(a_{-i})$ contains more than one action for some lists a_{-i} of actions of the other players.

2.8.2 Using best response functions to define Nash equilibrium

A Nash equilibrium is an action profile with the property that no player can do better by changing her action, given the other players' actions. Using the terminology just developed, we can alternatively define a Nash equilibrium to be an action profile for which every player's action is a best response to the other players' actions. That is, we have the following result.

- **PROPOSITION 34.1** *The action profile a^* is a Nash equilibrium of a strategic game with ordinal preferences if and only if every player's action is a best response to the other players' actions:*

$$a_i^* \text{ is in } B_i(a_{-i}^*) \text{ for every player } i. \quad (34.2)$$

If each player i has a single best response to each list a_{-i} of the other players' actions, we can write the conditions in (34.2) as equations. In this case, for each player i and each list a_{-i} of the other players' actions, denote the single member of $B_i(a_{-i})$ by $b_i(a_{-i})$ (that is, $B_i(a_{-i}) = \{b_i(a_{-i})\}$). Then (34.2) is equivalent to

$$a_i^* = b_i(a_{-i}^*) \text{ for every player } i, \quad (34.3)$$

a collection of n equations in the n unknowns a_i^* , where n is the number of players in the game. For example, in a game with two players, say 1 and 2, these equations are

$$\begin{aligned} a_1^* &= b_1(a_2^*) \\ a_2^* &= b_2(a_1^*). \end{aligned}$$

That is, in a two-player game in which each player has a single best response to every action of the other player, (a_1^*, a_2^*) is a Nash equilibrium if and only if player 1's action a_1^* is her best response to player 2's action a_2^* , and player 2's action a_2^* is her best response to player 1's action a_1^* .

2.8.3 Using best response functions to find Nash equilibria

The definition of a Nash equilibrium in terms of best response functions suggests a method for finding Nash equilibria:

- find the best response function of each player
- find the action profiles that satisfy (34.2) (which reduces to (34.3) if each player has a single best response to each list of the other players' actions).

To illustrate this method, consider the game in Figure 35.1. First find the best response of player 1 to each action of player 2. If player 2 chooses L , then player 1's best response is M (2 is the highest payoff for player 1 in this column); indicate the best response by attaching a star to player 1's payoff to (M, L) . If player 2 chooses C , then player 1's best response is T , indicated by the star attached to player 1's payoff to (T, C) . And if player 2 chooses R , then both T and B are best responses for player 1; both are indicated by stars. Second, find the best response of player 2 to each action of player 1 (for each row, find highest payoff of player 2); these best responses are indicated by attaching stars to player 2's payoffs. Finally, find the boxes in which both players' payoffs are starred. Each such box is a Nash equilibrium: the star on player 1's payoff means that player 1's action is a best response to player 2's action, and the star on player 2's payoff means that player 2's action is a best response to player 1's action. Thus we conclude that the game has two Nash equilibria: (M, L) and (B, R) .

	L	C	R
T	1, 2*	2*, 1	1*, 0
M	2*, 1*	0, 1*	0, 0
B	0, 1	0, 0	1*, 2*

Figure 35.1 Using best response functions to find Nash equilibria in a two-player game in which each player has three actions.

EXERCISE 35.1 (Finding Nash equilibria using best response functions)

- a. Find the players' best response functions in the *Prisoner's Dilemma* (Figure 13.1), *BoS* (Figure 16.1), *Matching Pennies* (Figure 17.1), and the two-player *Stag Hunt* (Figure 18.1) (and verify the Nash equilibria of these games).
- b. Find the Nash equilibria of the game in Figure 36.1 by finding the players' best response functions.

	L	C	R
T	2, 2	1, 3	0, 1
M	3, 1	0, 0	0, 0
B	1, 0	0, 0	0, 0

Figure 36.1 The game in Exercise 35.1b.

The players' best response functions for the game in Figure 35.1 are presented in a different format in Figure 36.2. In this figure, player 1's actions are on the horizontal axis and player 2's are on the vertical axis. (Thus the columns correspond to choices of player 1, and the rows correspond to choices of player 2, whereas the reverse is true in Figure 35.1. I choose this orientation for Figure 36.2 for consistency with the convention for figures of this type.) Player 1's best responses are indicated by circles, and player 2's by dots. Thus the circle at (T, C) reflects the fact that T is player 1's best response to player 2's choice of C , and the circles at (T, R) and (B, R) reflect the fact that T and B are both best responses of player 1 to player 2's choice of R . Any action pair marked by both a circle and a dot is a Nash equilibrium: the circle means that player 1's action is a best response to player 2's action, and the dot indicates that player 2's action is a best response to player 1's action.




Figure 36.2 The players' best response functions for the game in Figure 35.1. Player 1's best responses are indicated by circles, and player 2's by dots. The action pairs for which there is both a circle and a dot are the Nash equilibria.

Ⓜ EXERCISE 36.1 (Constructing best response functions) Draw the analogue of Figure 36.2 for the game in Exercise 35.1b.

- ② EXERCISE 37.1 (Dividing money) Two people have \$10 to divide between themselves. They use the following procedure. Each person names a number of dollars (a nonnegative integer), at most equal to 10. If the sum of the amounts that the people name is at most 10 then each person receives the amount of money she names (and the remainder is destroyed). If the sum of the amounts that the people name exceeds 10 and the amounts named are different then the person who names the smaller amount receives that amount and the other person receives the remaining money. If the sum of the amounts that the people name exceeds 10 and the amounts named are the same then each person receives \$5. Determine the best response of each player to each of the other player's actions, plot them in a diagram like Figure 36.2, and thus find the Nash equilibria of the game.

A diagram like Figure 36.2 is a convenient representation of the players' best response functions also in a game in which each player's set of actions is an interval of numbers, as the next example illustrates.

- ◆ EXAMPLE 37.2 (A synergistic relationship) Two individuals are involved in a synergistic relationship. If both individuals devote more effort to the relationship, they are both better off. For any given effort of individual j , the return to individual i 's effort first increases, then decreases. Specifically, an effort level is a nonnegative number, and individual i 's preferences (for $i = 1, 2$) are represented by the payoff function $a_i(c + a_j - a_i)$, where a_i is i 's effort level, a_j is the other individual's effort level, and $c > 0$ is a constant.

The following strategic game models this situation.

Players The two individuals.

Actions Each player's set of actions is the set of effort levels (nonnegative numbers).

Preferences Player i 's preferences are represented by the payoff function $a_i(c + a_j - a_i)$, for $i = 1, 2$.

In particular, each player has infinitely many actions, so that we cannot present the game in a table like those used previously (Figure 36.1, for example).

To find the Nash equilibria of the game, we can construct and analyze the players' best response functions. Given a_j , individual i 's payoff is a quadratic function of a_i that is zero when $a_i = 0$ and when $a_i = c + a_j$, and reaches a maximum in between. The symmetry of quadratic functions (see Section 17.3) implies that the best response of each individual i to a_j is

$$b_i(a_j) = \frac{1}{2}(c + a_j).$$

(If you know calculus, you can reach the same conclusion by setting the derivative of player i 's payoff with respect to a_i equal to zero.)

The best response functions are shown in Figure 38.1. Player 1's actions are plotted on the horizontal axis and player 2's actions are plotted on the vertical axis. Player 1's best response function associates an action for player 1 with every action

for player 2. Thus to interpret the function b_1 in the diagram, take a point a_2 on the vertical axis, and go across to the line labeled b_1 (the steeper of the two lines), then read down to the horizontal axis. The point on the horizontal axis that you reach is $b_1(a_2)$, the best action for player 1 when player 2 chooses a_2 . Player 2's best response function, on the other hand, associates an action for player 2 with every action of player 1. Thus to interpret this function, take a point a_1 on the horizontal axis, and go up to b_2 , then across to the vertical axis. The point on the vertical axis that you reach is $b_2(a_1)$, the best action for player 2 when player 1 chooses a_1 .




Figure 38.1 The players' best response functions for the game in Example 37.2. The game has a unique Nash equilibrium, $(a_1^*, a_2^*) = (c, c)$.

At a point (a_1, a_2) where the best response functions intersect in the figure, we have $a_1 = b_1(a_2)$, because (a_1, a_2) is on the graph of b_1 , player 1's best response function, and $a_2 = b_2(a_1)$, because (a_1, a_2) is on the graph of b_2 , player 2's best response function. Thus any such point (a_1, a_2) is a Nash equilibrium. In this game the best response functions intersect at a single point, so there is one Nash equilibrium. In general, they may intersect more than once; every point at which they intersect is a Nash equilibrium.

To find the point of intersection of the best response functions precisely, we can solve the two equations in (34.3):

$$\begin{aligned} a_1 &= \frac{1}{2}(c + a_2) \\ a_2 &= \frac{1}{2}(c + a_1). \end{aligned}$$

Substituting the second equation in the first, we get $a_1 = \frac{1}{2}(c + \frac{1}{2}(c + a_1)) = \frac{3}{4}c + \frac{1}{4}a_1$, so that $a_1 = c$. Substituting this value of a_1 into the second equation, we get $a_2 = c$. We conclude that the game has a unique Nash equilibrium $(a_1, a_2) = (c, c)$. (To reach this conclusion, it suffices to solve the two equations; we do not have to draw Figure 38.1. However, the diagram shows us at once that the game has a

unique equilibrium, in which both players' actions exceed $\frac{1}{2}c$, facts that serve to check the results of our algebra.)

In the game in this example, each player has a unique best response to every action of the other player, so that the best response functions are lines. If a player has many best responses to some of the other players' actions, then her best response function is "thick" at some points; several examples in the next chapter have this property (see, for example, Figure 64.1). Example 37.2 is special also because the game has a unique Nash equilibrium—the best response functions cross once. As we have seen, some games have more than one equilibrium, and others have no equilibrium. A pair of best response functions that illustrates some of the possibilities is shown in Figure 39.1. In this figure the shaded area of player 1's best response function indicates that for a_2 between \bar{a}_2 and \underline{a}_2 , player 1 has a range of best responses. For example, all actions of player 1 greater than a_1^{**} and at most a_1^{***} are best responses to the action a_2^{***} of player 2. For a game with these best response functions, the set of Nash equilibria consists of the pair of actions (a_1^*, a_2^*) , all the pairs of actions on player 2's best response function between (a_1^{**}, a_2^{**}) and (a_1^{***}, a_2^{***}) , and (a_1^{***}, a_2^{***}) .




Figure 39.1 An example of the best response functions of a two-player game in which each player's set of actions is an interval of numbers. The set of Nash equilibria of the game consists of the pair of actions (a_1^*, a_2^*) , all the pairs of actions on player 2's best response function between (a_1^{**}, a_2^{**}) and (a_1^{***}, a_2^{***}) , and (a_1^{***}, a_2^{***}) .

Ⓣ EXERCISE 39.1 (Strict and nonstrict Nash equilibria) Which of the Nash equilibria of the game whose best response functions are given in Figure 39.1 are strict (see the definition on page 31)?

Another feature that differentiates the best response functions in Figure 39.1 from those in Figure 38.1 is that the best response function b_1 of player 1 is not

continuous. When player 2's action is \bar{a}_2 , player 1's best response is a_1^{**} (indicated by the small disk at (a_1^{**}, \bar{a}_2)), but when player 2's action is slightly greater than \bar{a}_2 , player 1's best response is significantly less than a_1^{**} . (The small circle indicates a point excluded from the best response function.) Again, several examples in the next chapter have this feature. From Figure 39.1 we see that if a player's best response function is discontinuous, then depending on where the discontinuity occurs, the best response functions may not intersect at all—the game may, like *Matching Pennies*, have no Nash equilibrium.

- ? EXERCISE 40.1 (Finding Nash equilibria using best response functions) Find the Nash equilibria of the two-player strategic game in which each player's set of actions is the set of nonnegative numbers and the players' payoff functions are $u_1(a_1, a_2) = a_1(a_2 - a_1)$ and $u_2(a_1, a_2) = a_2(1 - a_1 - a_2)$.
- ? EXERCISE 40.2 (A joint project) Two people are engaged in a joint project. If each person i puts in the effort x_i , a nonnegative number equal to at most 1, which costs her $c(x_i)$, the outcome of the project is worth $f(x_1, x_2)$. The worth of the project is split equally between the two people, regardless of their effort levels. Formulate this situation as a strategic game. Find the Nash equilibria of the game when (a) $f(x_1, x_2) = 3x_1x_2$ and $c(x_i) = x_i^2$ for $i = 1, 2$, and (b) $f(x_1, x_2) = 4x_1x_2$ and $c(x_i) = x_i$ for $i = 1, 2$. In each case, is there a pair of effort levels that yields both players higher payoffs than the Nash equilibrium effort levels?

2.8.4 Illustration: contributing to a public good

Exercise 31.1 models decisions on whether to contribute to the provision of a "public good". We now study a model in which two people decide not only whether to contribute, but also *how much* to contribute.

Denote person i 's wealth by w_i , and the amount she contributes to the public good by c_i ($0 \leq c_i \leq w_i$); she spends her remaining wealth $w_i - c_i$ on "private goods" (like clothes and food, whose consumption by one person precludes their consumption by anyone else). The amount of the public good is equal to the sum of the contributions. Each person cares both about the amount of the public good and her consumption of private goods.

Suppose that person i 's preferences are represented by the payoff function $v_i(c_1 + c_2) + w_i - c_i$. Because w_i is a constant, person i 's preferences are alternatively represented by the payoff function

$$u_i(c_1, c_2) = v_i(c_1 + c_2) - c_i.$$

This situation is modeled by the following strategic game.

Players The two people.

Actions Player i 's set of actions is the set of her possible contributions (nonnegative numbers less than or equal to w_i), for $i = 1, 2$.

Preferences Player i 's preferences are represented by the payoff function $u_i(c_1, c_2) = v_i(c_1 + c_2) - c_i$, for $i = 1, 2$.

To find the Nash equilibria of this strategic game, consider the players' best response functions. Player 1's best response to the contribution c_2 of player 2 is the value of c_1 that maximizes $v_1(c_1 + c_2) - c_1$. Without specifying the form of the function v_1 we cannot explicitly calculate this optimal value. However, we can determine how it varies with c_2 .

First consider player 1's best response to $c_2 = 0$. Suppose that the form of the function v_1 is such that the function $u_1(c_1, 0)$ increases up to its maximum, then decreases (as in Figure 41.1). Then player 1's best response to $c_2 = 0$, which I denote $b_1(0)$, is unique. This best response is the value of c_1 that maximizes $u_1(c_1, 0) = v_1(c_1) - c_1$ subject to $0 \leq c_1 \leq w_1$. Assume that $0 < b_1(0) < w_1$: player 1's optimal contribution to the public good when player 2 makes no contribution is positive and less than her entire wealth.

Now consider player 1's best response to $c_2 = k > 0$. This best response is the value of c_1 that maximizes $u_1(c_1, k) = v_1(c_1 + k) - c_1$. Now, we have

$$u_1(c_1, k) = u_1(c_1 + k, 0) + k.$$

That is, the graph of $u_1(c_1, k)$ as a function of c_1 is the translation to the left k units and up k units of the graph of $u_1(c_1, 0)$ as a function of c_1 (refer to Figure 41.1). Thus if $k \leq b_1(0)$ then $b_1(k) = b_1(0) - k$: if player 2's contribution increases from 0 to k then player 1's best response decreases by k . If $k > b_1(0)$ then, given the form of $u_1(c_1, 0)$, we have $b_1(k) = 0$.




Figure 41.1 The relation between player 1's best responses $b_1(0)$ and $b_1(k)$ to $c_2 = 0$ and $c_2 = k$ in the game of contributing to a public good.

We conclude that if player 2 increases her contribution by k then player 1's best response is to reduce her contribution by k (or to zero, if k is larger than player 1's original contribution)!

The same analysis applies to player 2: for every unit more that player 1 contributes, player 2 contributes a unit less, so long as her contribution is nonnegative.




Figure 42.1 The best response functions for the game of contributing to a public good in Section 2.8.4 in a case in which $b_1(0) > b_2(0)$. The best response function of player 1 is the black line; that of player 2 is the gray line.

The function v_2 may be different from the function v_1 , so that player 1's best contribution $b_1(0)$ when $c_2 = 0$ may be different from player 2's best contribution $b_2(0)$ when $c_1 = 0$. But both best response functions have the same character: the slope of each function is -1 where the value of the function is positive. They are shown in Figure 42.1 for a case in which $b_1(0) > b_2(0)$.

We deduce that if $b_1(0) > b_2(0)$ then the game has a unique Nash equilibrium, $(b_1(0), 0)$: player 2 contributes nothing. Similarly, if $b_1(0) < b_2(0)$ then the unique Nash equilibrium is $(0, b_2(0))$: player 1 contributes nothing. That is, the person who contributes more when the other person contributes nothing is the only one to make a contribution in a Nash equilibrium. Only if $b_1(0) = b_2(0)$, which is not likely if the functions v_1 and v_2 differ, is there an equilibrium in which both people contribute. In this case the downward-sloping parts of the best response functions coincide, so that any pair of contributions (c_1, c_2) with $c_1 + c_2 = b_1(0)$ and $c_i \geq 0$ for $i = 1, 2$ is a Nash equilibrium.

In summary, the notion of Nash equilibrium predicts that, except in unusual circumstances, only one person contributes to the provision of the public good when each person's payoff function takes the form $v_i(c_1 + c_2) + w_i - c_i$, each function $v_i(c_i) - c_i$ increases to a maximum, then decreases, and each person optimally contributes less than her entire wealth when the other person does not contribute. The person who contributes is the one who wishes to contribute more when the other person does not contribute. In particular, the identity of the person who contributes does not depend on the distribution of wealth; any distribution in which each person optimally contributes less than her entire wealth when the other person does not contribute leads to the same outcome.

The next exercise asks you to consider a case in which the amount of the public good affects each person's enjoyment of the private good. (The public good might be clean air, which improves each person's enjoyment of her free time.)

- ? EXERCISE 43.1 (Contributing to a public good) Consider the model in this section when $u_i(c_1, c_2)$ is the sum of three parts: the amount $c_1 + c_2$ of the public good provided, the amount $w_i - c_i$ person i spends on private goods, and a term $(w_i - c_i)(c_1 + c_2)$ that reflects an interaction between the amount of the public good and her private consumption—the greater the amount of the public good, the more she values her private consumption. In summary, suppose that person i 's payoff is $c_1 + c_2 + w_i - c_i + (w_i - c_i)(c_1 + c_2)$, or

$$w_i + c_j + (w_i - c_i)(c_1 + c_2),$$

where j is the other person. Assume that $w_1 = w_2 = w$, and that each player i 's contribution c_i may be any number (positive or negative, possibly larger than w). Find the Nash equilibrium of the game that models this situation. (You can calculate the best responses explicitly. Imposing the sensible restriction that c_i lie between 0 and w complicates the analysis, but does not change the answer.) Show that in the Nash equilibrium both players are worse off than they are when they both contribute one half of their wealth to the public good. If you can, extend the analysis to the case of n people. As the number of people increases, how does the total amount contributed in a Nash equilibrium change? Compare the players' equilibrium payoffs with their payoffs when each contributes half her wealth to the public good, as n increases without bound. (The game is studied further in Exercise 388.2.)

2.9 Dominated actions

2.9.1 Strict domination

You drive up to a red traffic light. The left lane is free; in the right lane there is a car that may turn right when the light changes to green, in which case it will have to wait for a pedestrian to cross the side street. Assuming you wish to progress as quickly as possible, the action of pulling up in the left lane “strictly dominates” that of pulling up in the right lane. If the car in the right lane turns right then you are much better off in the left lane, where your progress will not be impeded; and even if the car in the right lane does not turn right, you are still better off in the left lane, rather than behind the other car.

In any game, a player's action “strictly dominates” another action if it is superior, no matter what the other players do.

- DEFINITION 43.2 (*Strict domination*) In a strategic game with ordinal preferences, player i 's action a_i'' **strictly dominates** her action a_i' if

$$u_i(a_i'', a_{-i}) > u_i(a_i', a_{-i}) \text{ for every list } a_{-i} \text{ of the other players' actions,}$$

where u_i is a payoff function that represents player i 's preferences. We say that the action a_i' is **strictly dominated**.

In the *Prisoner's Dilemma*, for example, the action *Fink* strictly dominates the action *Quiet*: regardless of her opponent's action, a player prefers the outcome when she chooses *Fink* to the outcome when she chooses *Quiet*. In *BoS*, on the other hand, neither action strictly dominates the other: *Bach* is better than *Stravinsky* if the other player chooses *Bach*, but is worse than *Stravinsky* if the other player chooses *Stravinsky*.

A strictly dominated action is not a best response to any actions of the other players: whatever the other players do, some other action is better. Since a player's Nash equilibrium action is a best response to the other players' Nash equilibrium actions,

a strictly dominated action is not used in any Nash equilibrium.

When looking for the Nash equilibria of a game, we can thus eliminate from consideration all strictly dominated actions. For example, we can eliminate *Quiet* for each player in the *Prisoner's Dilemma*, leaving $(Fink, Fink)$ as the only candidate for a Nash equilibrium. (As we know, this action pair is indeed a Nash equilibrium.)

The fact that the action a_i'' strictly dominates the action a_i' of course does *not* imply that a_i'' strictly dominates *all* actions. Indeed, a_i'' may itself be strictly dominated. In the left-hand game in Figure 44.1, for example, *M* strictly dominates *T*, but *B* is better than *M* if player 2 chooses *R*. (I give only the payoffs of player 1 in the figure, because those of player 2 are not relevant.) Since *T* is strictly dominated, the game has no Nash equilibrium in which player 1 uses it; but the game may also not have any equilibrium in which player 1 uses *M*. In the right-hand game, *M* strictly dominates *T*, but is itself strictly dominated by *B*. In this case, in any Nash equilibrium player 1's action is *B* (her only action that is not strictly dominated).

	<i>L</i>	<i>R</i>
<i>T</i>	1	0
<i>M</i>	2	1
<i>B</i>	1	3

	<i>L</i>	<i>R</i>
<i>T</i>	1	0
<i>M</i>	2	1
<i>B</i>	3	2

Figure 44.1 Two games in which player 1's action *T* is strictly dominated by *M*. (Only player 1's payoffs are given.) In the left-hand game, *B* is better than *M* if player 2 chooses *R*; in the right-hand game, *M* itself is strictly dominated, by *B*.

A strictly dominated action is incompatible not only with a steady state, but also with rational behavior by a player who confronts a game for the first time. This fact is the first step in a theory different from Nash equilibrium, explored in Chapter 12.

2.9.2 Weak domination

As you approach the red light in the situation at the start of the previous section, there is a car in *each* lane. The car in the right lane may, or may not, be turning

right; if it is, it may be delayed by a pedestrian crossing the side street. The car in the left lane cannot turn right. In this case your pulling up in the left lane “weakly dominates”, though does not strictly dominate, your pulling up in the right lane. If the car in the right lane does not turn right, then both lanes are equally good; if it does, then the left lane is better.

In any game, a player’s action “weakly dominates” another action if the first action is at least as good as the second action, no matter what the other players do, and is better than the second action for some actions of the other players.

- DEFINITION 45.1 (*Weak domination*) In a strategic game with ordinal preferences, player i ’s action a_i'' **weakly dominates** her action a_i' if

$$u_i(a_i'', a_{-i}) \geq u_i(a_i', a_{-i}) \text{ for every list } a_{-i} \text{ of the other players' actions}$$

and

$$u_i(a_i'', a_{-i}) > u_i(a_i', a_{-i}) \text{ for some list } a_{-i} \text{ of the other players' actions,}$$

where u_i is a payoff function that represents player i ’s preferences. We say that the action a_i' is **weakly dominated**.

For example, in the game in Figure 45.1 (in which, once again, only player 1’s payoffs are given), M weakly dominates T , and B weakly dominates M ; B strictly dominates T .

	L	R
T	1	0
M	2	0
B	2	1

Figure 45.1 A game illustrating weak domination. (Only player 1’s payoffs are given.) The action M weakly dominates T ; B weakly dominates M . The action B strictly dominates T .

In a *strict* Nash equilibrium (Section 2.7.8) no player’s equilibrium action is weakly dominated: every non-equilibrium action for a player yields her a payoff less than does her equilibrium action, and hence does not weakly dominate the equilibrium action.

Can an action be weakly dominated in a nonstrict Nash equilibrium? Definitely. Consider the games in Figure 46.1. In both games B weakly (but not strictly) dominates C for both players. But in both games (C, C) is a Nash equilibrium: *given* that player 2 chooses C , player 1 cannot do better than choose C , and *given* that player 1 chooses C , player 2 cannot do better than choose C . Both games also have a Nash equilibrium, (B, B) , in which neither player’s action is weakly dominated. In the left-hand game this equilibrium is better for both players than the equilibrium (C, C) in which both players’ actions are weakly dominated, whereas in the right-hand game it is worse for both players than (C, C) .

	<i>B</i>	<i>C</i>
<i>B</i>	1, 1	0, 0
<i>C</i>	0, 0	0, 0

	<i>B</i>	<i>C</i>
<i>B</i>	1, 1	2, 0
<i>C</i>	0, 2	2, 2

Figure 46.1 Two strategic games with a Nash equilibrium (C, C) in which both players' actions are weakly dominated.

- ⑦ EXERCISE 46.1 (Strict equilibria and dominated actions) For the game in Figure 46.2, determine, for each player, whether any action is strictly dominated or weakly dominated. Find the Nash equilibria of the game; determine whether any equilibrium is strict.

	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	0, 0	1, 0	1, 1
<i>M</i>	1, 1	1, 1	3, 0
<i>B</i>	1, 1	2, 1	2, 2

Figure 46.2 The game in Exercise 46.1.

- ⑦ EXERCISE 46.2 (Nash equilibrium and weakly dominated actions) Give an example of a two-player strategic game in which each player has finitely many actions and in the only Nash equilibrium both players' actions are weakly dominated.

2.9.3 Illustration: voting

Two candidates, A and B , vie for office. Each of an odd number of citizens may vote for either candidate. (Abstention is not possible.) The candidate who obtains the most votes wins. (Because the number of citizens is odd, a tie is impossible.) A majority of citizens prefer A to win than B to win.

The following strategic game models the citizens' voting decisions in this situation.

Players The citizens.

Actions Each player's set of actions consists of voting for A and voting for B .

Preferences All players are indifferent among all action profiles in which a majority of players vote for A ; all players are also indifferent among all action profiles in which a majority of players vote for B . Some players (a majority) prefer an action profile of the first type to one of the second type, and the others have the reverse preference.

I claim that a citizen's voting for her less preferred candidate is weakly dominated by her voting for her favorite candidate. Suppose that citizen i prefers candidate A ; fix the votes of all citizens other than i . If citizen i switches from voting for B to voting for A then, depending on the other citizens' votes, either the outcome

does not change, or A wins rather than B ; such a switch cannot cause the winner to change from A to B . That is, citizen i 's switching from voting for B to voting for A either has no effect on the outcome, or makes her better off; it cannot make her worse off.

The game has Nash equilibria in which some, or all, citizens' actions are weakly dominated. For example, the action profile in which all citizens vote for B is a Nash equilibrium (no citizen's switching her vote has any effect on the outcome).

- ⓧ EXERCISE 47.1 (Voting) Find all the Nash equilibria of the game. (First consider action profiles in which the winner obtains one more vote than the loser and at least one citizen who votes for the winner prefers the loser to the winner, then profiles in which the winner obtains one more vote than the loser and all citizens who vote for the winner prefer the winner to the loser, and finally profiles in which the winner obtains three or more votes more than the loser.) Is there any equilibrium in which no player uses a weakly dominated action?

Consider a variant of the game in which the number of candidates is greater than two. A variant of the argument above shows that a citizen's action of voting for her least preferred candidate is weakly dominated by her other two actions. The next exercise asks you to show that no other action is weakly dominated.

- ⓧ EXERCISE 47.2 (Voting between three candidates) Suppose there are three candidates, A , B , and C , and no citizen is indifferent between any two of them. A tie for first place is possible in this case; assume that a citizen who prefers a win by x to a win by y ranks a tie between x and y between an outright win for x and an outright win for y . Show that a citizen's only weakly dominated action is a vote for her least preferred candidate. Find a Nash equilibrium in which some citizen does not vote for her favorite candidate, but the action she takes is not weakly dominated.
- ⓧ EXERCISE 47.3 (Approval voting) In the system of "approval voting", a citizen may vote for as many candidates as she wishes. If there are two candidates, say A and B , for example, a citizen may vote for neither candidate, for A , for B , or for both A and B . As before, the candidate who obtains the most votes wins. Show that any action that includes a vote for a citizen's least preferred candidate is weakly dominated, as is any action that does not include a vote for her most preferred candidate. More difficult: show that if there are k candidates then for a citizen who prefers candidate 1 to candidate 2 to \dots to candidate k , the action that consists of votes for candidates 1 and $k - 1$ is *not* weakly dominated.

2.9.4 Illustration: collective decision-making

The members of a group of people are affected by a policy, modeled as a number. Each person i has a favorite policy, denoted x_i^* ; she prefers the policy y to the policy z if and only if y is closer to x_i^* than is z . The number n of people is odd. The following mechanism is used to choose a policy: each person names a policy, and the policy chosen is the median of those named. (That is, the policies named

are put in order, and the one in the middle is chosen. If, for example, there are five people, and they name the policies $-2, 0, 0.6, 5,$ and $10,$ then the policy 0.6 is chosen.)

What outcome does this mechanism induce? Does anyone have an incentive to name her favorite policy, or are people induced to distort their preferences? We can answer these questions by studying the following strategic game.

Players The n people.

Actions Each person's set of actions is the set of policies (numbers).

Preferences Each person i prefers the action profile a to the action profile a' if and only if the median policy named in a is closer to x_i^* than is the median policy named in a' .

I claim that for each player $i,$ the action of naming her favorite policy x_i^* weakly dominates *all* her other actions. The reason is that relative to the situation in which she names $x_i^*,$ she can change the median only by naming a policy *further* from her favorite policy than the current median; no change in the policy she names moves the median closer to her favorite policy.

Precisely, I show that for each action $x_i \neq x_i^*$ of player $i,$ (a) for *all* actions of the other players, player i is at least as well off naming x_i^* as she is naming $x_i,$ and (b) for *some* actions of the other players she is better off naming x_i^* than she is naming $x_i.$ Take $x_i > x_i^*.$

a. For any list of actions of the players *other than* player $i,$ denote the value of the $\frac{1}{2}(n-1)$ th highest action by \underline{a} and the value of the $\frac{1}{2}(n+1)$ th highest action by \bar{a} (so that half of the remaining players' actions are at most \underline{a} and half of them are at least \bar{a}).

- If $\bar{a} \leq x_i^*$ or $\underline{a} \geq x_i$ then the median policy is the same whether player i names x_i^* or $x_i.$
- If $\bar{a} > x_i^*$ and $\underline{a} < x_i$ then when player i names x_i^* the median policy is at most the greater of x_i^* and \underline{a} and when player i names x_i the median policy is at least the lesser of x_i and $\bar{a}.$ Thus player i is worse off naming x_i than she is naming $x_i^*.$

b. Suppose that half of the remaining players name policies less than x_i^* and half of them name policies greater than $x_i.$ Then the outcome is x_i^* if player i names $x_i^*,$ and x_i if she names $x_i.$ Thus she is better off naming x_i^* than she is naming $x_i.$

A symmetric argument applies when $x_i < x_i^*.$

If we think of the mechanism as asking the players to name their favorite policies, then the result is that telling the truth weakly dominates all other actions.

An implication of the fact that player i 's naming her favorite policy x_i^* weakly dominates *all* her other actions is that the action profile in which every player

names her favorite policy is a Nash equilibrium. That is, truth-telling is a Nash equilibrium, in the interpretation of the previous paragraph.

- ⑦ EXERCISE 49.1 (Other Nash equilibria of the game modeling collective decision-making) Find two Nash equilibria in which the outcome is the median favorite policy, and one in which it is not.
- ⑦ EXERCISE 49.2 (Another mechanism for collective decision-making) Consider the variant of the mechanism for collective decision-making described above in which the policy chosen is the *mean*, rather than the median, of the policies named by the players. Does a player's action of naming her favorite policy weakly dominate all her other actions?

2.10 Equilibrium in a single population: symmetric games and symmetric equilibria

A Nash equilibrium of a strategic game corresponds to a steady state of an interaction between the members of several populations, one for each player in the game, each play of the game involving one member of each population. Sometimes we want to model a situation in which the members of a *single* homogeneous population are involved anonymously in a symmetric interaction. Consider, for example, pedestrians approaching each other on a sidewalk or car drivers arriving simultaneously at an intersection from different directions. In each case, the members of each encounter are drawn from the same population: pairs from a single population of pedestrians meet each other, and groups from a single population of car drivers simultaneously approach intersections. And in each case, every participant's role is the same.

I restrict attention here to cases in which each interaction involves two participants. Define a two-player game to be "symmetric" if each player has the same set of actions and each player's evaluation of an outcome depends only on her action and that of her opponent, not on whether she is player 1 or player 2. That is, player 1 feels the same way about the outcome (a_1, a_2) , in which her action is a_1 and her opponent's action is a_2 , as player 2 feels about the outcome (a_2, a_1) , in which her action is a_1 and her opponent's action is a_2 . In particular, the players' preferences may be represented by payoff functions in which both players' payoffs are the same whenever the players choose the same action: $u_1(a, a) = u_2(a, a)$ for every action a .

- ▶ DEFINITION 49.3 (*Symmetric two-player strategic game with ordinal preferences*) A two-player strategic game with ordinal preferences is **symmetric** if the players' sets of actions are the same and the players' preferences are represented by payoff functions u_1 and u_2 for which $u_1(a_1, a_2) = u_2(a_2, a_1)$ for every action pair (a_1, a_2) .

A two-player game in which each player has two actions is symmetric if the players' preferences are represented by payoff functions that take the form shown

in Figure 50.1, where $w, x, y,$ and z are arbitrary numbers. Several of the two-player games we have considered are symmetric, including the *Prisoner's Dilemma*, the two-player *Stag Hunt* (given again in Figure 50.2), and the game in Exercise 37.1. *BoS* (Figure 16.1) and *Matching Pennies* (Figure 17.1) are not symmetric.

	A	B
A	w, w	x, y
B	y, x	z, z

Figure 50.1 A two-player symmetric game.

	<i>Quiet</i>	<i>Fink</i>		<i>Stag</i>	<i>Hare</i>
<i>Quiet</i>	2, 2	0, 3		2, 2	0, 1
<i>Fink</i>	3, 0	1, 1		1, 0	1, 1

Figure 50.2 Two symmetric games: the *Prisoner's Dilemma* (left) and the two-player *Stag Hunt* (right).

- ❓ EXERCISE 50.1 (Symmetric strategic games) Which of the games in Exercises 29.2 and 40.1, Example 37.2, Section 2.8.4, and Figure 46.1 are symmetric?

When the players in a symmetric two-player game are drawn from a single population, nothing distinguishes one of the players in any given encounter from the other. We may call them “player 1” and “player 2”, but these labels are only for our convenience. There is only one role in the game, so that a steady state is characterized by a *single* action used by every participant whenever playing the game. An action a^* corresponds to such a steady state if no player can do better by using any other action, given that all the other players use a^* . An action a^* has this property if and only if (a^*, a^*) is a Nash equilibrium of the game. In other words, the solution that corresponds to a steady state of pairwise interactions between the members of a single population is “symmetric Nash equilibrium”: a Nash equilibrium in which both players take the same action. The idea of this notion of equilibrium does not depend on the game’s having only two players, so I give a definition for a game with any number of players.

- DEFINITION 50.2 (*Symmetric Nash equilibrium*) An action profile a^* in a strategic game with ordinal preferences in which each player has the same set of actions is a **symmetric Nash equilibrium** if it is a Nash equilibrium and a_i^* is the same for every player i .

As an example, consider a model of approaching pedestrians. Each participant in any given encounter has two possible actions—to step to the right, and to step to the left—and is better off when participants both step in the same direction than when they step in different directions (in which case a collision occurs). The resulting symmetric strategic game is given in Figure 51.1. The game has two symmetric Nash equilibria, namely $(Left, Left)$ and $(Right, Right)$. That is, there

are two steady states, in one of which every pedestrian steps to the left as she approaches another pedestrian, and in another of which both participants step to the right. (The latter steady state seems to prevail in the USA and Canada.)

	<i>Left</i>	<i>Right</i>
<i>Left</i>	1, 1	0, 0
<i>Right</i>	0, 0	1, 1

Figure 51.1 Approaching pedestrians.

A symmetric game may have no symmetric Nash equilibrium. Consider, for example, the game in Figure 51.2. This game has two Nash equilibria, (X, Y) and (Y, X) , neither of which is symmetric. You may wonder if, in such a situation, there is a steady state in which each player does not always take the same action in every interaction. This question is addressed in Section 4.7.

	<i>X</i>	<i>Y</i>
<i>X</i>	0, 0	1, 1
<i>Y</i>	1, 1	0, 0

Figure 51.2 A symmetric game with no symmetric Nash equilibrium.

- ⑦ EXERCISE 51.1 (Equilibrium for pairwise interactions in a single population) Find all the Nash equilibria of the game in Figure 51.3. Which of the equilibria, if any, correspond to a steady state if the game models pairwise interactions between the members of a single population?

	<i>A</i>	<i>B</i>	<i>C</i>
<i>A</i>	1, 1	2, 1	4, 1
<i>B</i>	1, 2	5, 5	3, 6
<i>C</i>	1, 4	6, 3	0, 0

Figure 51.3 The game in Exercise 51.1.

Notes

The notion of a strategic game originated in the work of Borel (1921) and von Neumann (1928). The notion of Nash equilibrium (and its interpretation) is due to Nash (1950a). (The idea that underlies it goes back at least to Cournot 1838, Ch. 7.)

The *Prisoner's Dilemma* appears to have first been considered by Melvin Dresher and Merrill Flood, who used it in an experiment at the RAND Corporation in January 1950 (Flood 1958/59, 11–17); it is an example in Nash's PhD thesis (see Kuhn and Nasar 2001), submitted in May 1950. The story associated with it is due to Tucker (1950) (see Straffin 1980). O'Neill (1994, 1010–1013) argues that there is no

evidence that game theory (and in particular the *Prisoner's Dilemma*) influenced US nuclear strategists in the 1950s. The precise analysis of the idea that common property will be overused was initiated by Gordon (1954). Hardin (1968) coined the phrase "tragedy of the commons".

BoS, like the *Prisoner's Dilemma*, is an example in Nash's PhD thesis; Luce and Raiffa (1957, 90–91) name it and associate a story with it. *Matching Pennies* was first considered by von Neumann (1928). Rousseau's sentence about hunting stags is interpreted as a description of a game by Ullmann-Margalit (1977, 121) and Jervis (1977/78), following discussion by Waltz (1959, 167–169) and Lewis (1969, 7, 47).

The information about John Nash in the box on p. 21 comes from Leonard (1994), Kuhn et al. (1995), Kuhn (1996), Myerson (1996), Nasar (1998), and Nash (1995). Nash's PhD thesis is reprinted in Kuhn and Nasar (2001). *Hawk–Dove* is known also as "Chicken" (two drivers approach each other on a narrow road; the one who pulls over first is "chicken"). It was first suggested (in a more complicated form) as a model of animal conflict by Maynard Smith and Price (1973). The discussion of focal points in the box on p. 30 draws on Schelling (1960, 54–58).

Games modeling voluntary contributions to a public good were first considered by Olson (1965, Section I.D). The game in Exercise 31.1 is studied in detail by Palfrey and Rosenthal (1984). The result in Section 2.8.4 is due to Warr (1983) and Bergstrom, Blume, and Varian (1986).

Game theory was first used to study voting behavior by Farquharson (1969) (whose book was completed in 1958; see also Niemi 1983). The system of "approval voting" in Exercise 47.3 was first studied formally by Brams and Fishburn (1978, 1983).

Exercise 16.1 is based on Leonard (1990). Exercise 25.2 is based on Ullmann-Margalit (1977, 48). The game in Exercise 29.1 is taken from Van Huyck, Battalio, and Beil (1990). The game in Exercise 32.1 is taken from Moulin (1986, 72). The game in Exercise 32.2 was first studied by Palfrey and Rosenthal (1983). Exercise 33.1 is based on Braess (1968); see also Murchland (1970). The game in Exercise 37.1 is taken from Brams (1993).