GAMES OF STRATEGY

THIRD EDITION

Avinash Dixit
Princeton University

Susan Skeath
Wellesley College

David Reiley
University of Arizona and Yahoo! Research

W. W. Norton & Company
New York • London

Contents

PART FOUR
Applications to Specific Strategic Situations

1

Basic Ideas and Examples

All introductory textbooks begin by attempting to convince the student readers that the subject is of great importance in the world and therefore merits their attention. The physical sciences and engineering claim to be the basis of modern technology and therefore of modern life; the social sciences discuss big issues of governance—for example, democracy and taxation; the humanities claim that they revive your soul after it has been deadened by exposure to the physical and social sciences and to engineering. Where does the subject games of strategy, often called game theory, fit into this picture, and why should you study it?

We offer a practical motivation much more individual and closer to your personal concerns than most other subjects. You play games of strategy all the time: with your parents, siblings, friends, and enemies, and even with your professors. You have probably acquired a lot of instinctive expertise, and we hope you will recognize in what follows some of the lessons that you have already learned. We will build on this experience, systematize it, and develop it to the point where you will be able to improve your strategic skills and use them more methodically. Opportunities for such uses will appear throughout the rest of your life; you will go on playing such games with your employers, employees, spouses, children, and even strangers.

Not that the subject lacks wider importance. Similar games are played in business, politics, diplomacy, and wars—in fact, whenever people interact to strike mutually agreeable deals or to resolve conflicts. Being able to recognize such games will enrich your understanding of the world around you and will make you a better participant in all its affairs.

It will also have a more immediate payoff in your study of many other subjects. Economics and business courses already use a great deal of game-theoretic thinking. Political science is rapidly catching up. Biology has been importantly influenced by the concepts of evolutionary games and has in turn exported these ideas to economics. Psychology and philosophy also interact with the study of games of strategy. Game theory has become a provider of concepts and techniques of analysis for many disciplines, one might say all disciplines except those dealing with completely inanimate objects.

1 WHAT IS A GAME OF STRATEGY?

The word *game* may convey an impression that the subject is frivolous or unimportant in the larger scheme of things—that it deals with trivial pursuits such as gambling and sports when the world is full of weightier matters such as war and business and your education, career, and relationships. Actually, games of strategy are not "just a game"; all of these weighty matters are instances of games, and game theory helps us understand them all. But it will not hurt to start with gambling or sports.

Most games include chance, skill, and strategy in varying proportions. Playing double or nothing on the toss of a coin is a game of pure chance, unless you have exceptional skill in doctoring or tossing coins. A hundred-yard dash is a game of pure skill, although some chance elements can creep in; for example, a runner may simply have a slightly off day for no clear reason.

Strategy is a skill of a different kind. In the context of sports, it is a part of the mental skill needed to play well; it is the calculation of how best to use your physical skill. For example, in tennis, you develop physical skill by practicing your serves (first serves hard and flat, second serves with spin or kick) and passing shots (hard, low, and accurate). The strategic skill is knowing where to put your serve (wide, or on the T) or passing shot (crosscourt, or down the line). In football, you develop such physical skills as blocking and tackling, running and catching, and throwing. Then the coach, knowing the physical skills of his own team and those of the opposing team, calls the plays that best exploit his team's skills and the other team's weaknesses. The coach's calculation constitutes the strategy. The physical game of football is played on the gridiron by jocks; the strategic game is played in the offices and on the sidelines by coaches and by nerdy assistants.

A hundred-yard dash is a matter of exercising your physical skill as best you can; it offers no opportunities to observe and react to what other runners in the race are doing and therefore no scope for strategy. Longer races do entail strategy—whether you should lead to set the pace, how soon before the finish you should try to break away, and so on.

Strategic thinking is essentially about your interactions with others: someone else is also doing similar thinking at the same time and about the same situation. Your opponents in a marathon may try to frustrate or facilitate your attempts to lead, as they think best suits their interests. Your opponent in tennis tries to guess where you will put your serve or passing shot; the opposing coach in football calls the play that will best counter what he thinks you will call. Of course, just as you must take into account what the other player is thinking, he is taking into account what you are thinking. Game theory is the analysis, or science, if you like, of such interactive decision making.

When you think carefully before you act—when you are aware of your objectives or preferences and of any limitations or constraints on your actions and choose your actions in a calculated way to do the best according to your own criteria—you are said to be behaving rationally. Game theory adds another dimension to rational behavior—namely, interaction with other equally rational decision makers. In other words, game theory is the science of rational behavior in interactive situations.

We do not claim that game theory will teach you the secrets of perfect play or ensure that you will never lose. For one thing, your opponent can read the same book, and both of you cannot win all the time. More importantly, many games are complex and subtle enough, and most actual situations include enough idiosyncratic or chance elements, that game theory cannot hope to offer surefire recipes for action. What it does is to provide some general principles for thinking about strategic interactions. You have to supplement these ideas and some methods of calculation with many details specific to your situation before you can devise a successful strategy for it. Good strategists mix the science of game theory with their own experience; one might say that game playing is as much art as science. We will develop the general ideas of the science but will also point out its limitations and tell you when the art is more important.

You may think that you have already acquired the art from your experience or instinct, but you will find the study of the science useful nonetheless. The science systematizes many general principles that are common to several contexts or applications. Without general principles, you would have to figure out from scratch each new situation that requires strategic thinking. That would be especially difficult to do in new areas of application—for example, if you learned your art by playing games against parents and siblings and must now practice strategy against business competitors. The general principles of game theory provide you with a ready reference point. With this foundation in place, you can proceed much more quickly and confidently to acquire and add the situation-specific features or elements of the art to your thinking and action.

2 SOME EXAMPLES AND STORIES OF STRATEGIC GAMES

With the aims announced in Section 1, we will begin by offering you some simple examples, many of them taken from situations that you have probably encountered in your own lives, where strategy is of the essence. In each case we will point out the crucial strategic principle. Each of these principles will be discussed more fully in a later chapter, and after each example we will tell you where the details can be found. But don't jump to them right away; for a while, just read all the examples to get a preliminary idea of the whole scope of strategy and of strategic games.

A. Which Passing Shot?

Tennis at its best consists of memorable duels between top players: John McEnroe versus Ivan Lendl, Pete Sampras versus Andre Agassi, and Martina Navratilova versus Chris Evert. Picture the 1983 U.S. Open final between Evert and Navratilova.[1] Navratilova at the net has just volleyed to Evert on the baseline. Evert is about to hit a passing shot. Should she go down the line or crosscourt? And should Navratilova expect a down-the-line shot and lean slightly that way or expect a crosscourt shot and lean the other way?

Conventional wisdom favors the down-the-line shot. The ball has a shorter distance to travel to the net, so the other player has less time to react. But this does not mean that Evert should use that shot all of the time. If she did, Navratilova would confidently come to expect it and prepare for it, and the shot would not be so successful. To improve the success of the down-the-line passing shot, Evert has to use the crosscourt shot often enough to keep Navratilova guessing on any single instance.

Similarly in football, with a yard to go on third down, a run up the middle is the percentage play—that is, the one used most often—but the offense must throw a pass occasionally in such situations "to keep the defense honest."

Thus the most important general principle of such situations is not what Evert *should* do but what she *should not* do: she should not do the same thing all the time or systematically. If she did, then Navratilova would learn to cover that, and Evert's chances of success would fall.

Not doing any one thing systematically means more than not playing the same shot in every situation of this kind. Evert should not even mechanically switch back and forth between the two shots. Navratilova would spot and exploit this *pattern* or indeed any other detectable system. Evert must make the choice on each particular occasion *at random* to prevent this guessing.

[1] Chris Evert won her first title at the U.S. Open in 1975. Navratilova claimed her first title in the 1983 final.

This general idea of "mixing one's plays" is well known, even to sports commentators on television. But there is more to the idea, and these further aspects require analysis in greater depth. Why is down-the-line the percentage shot? Should one play it 80% of the time or 90% or 99%? Does it make any difference if the occasion is particularly big; for example, does one throw that pass on third down in the regular season but not in the Super Bowl? In actual practice, just how does one mix one's plays? What happens when a third possibility (the lob) is introduced? We will examine and answer such questions in Chapters 7 and 8.

The movie *The Princess Bride* (1987) illustrates the same idea in the "battle of wits" between the hero (Westley) and a villain (Vizzini). Westley is to poison one of two wineglasses out of Vizzini's sight, and Vizzini is to decide who will drink from which glass. Vizzini goes through a number of convoluted arguments as to why Westley should poison one glass. But all of the arguments are innately contradictory, because Westley can anticipate Vizzini's logic and choose to put the poison in the other glass. Conversely, if Westley uses any specific logic or system to choose one glass, Vizzini can anticipate that and drink from the other glass, leaving Westley to drink from the poisoned one. Thus, Westley's strategy has to be random or unsystematic.

The scene illustrates something else as well. In the film, Vizzini loses the game and with it his life. But it turns out that Westley had poisoned both glasses; over the last several years, he had built up immunity to the poison. So Vizzini was actually playing the game under a fatal information disadvantage. Players can sometimes cope with such asymmetries of information; Chapter 9 examines when and how they can do so.

B. The GPA Rat Race

You are enrolled in a course that is graded on a curve. No matter how well you do in absolute terms, only 40% of the students will get As, and only 40% will get Bs. Therefore you must work hard, not just in absolute terms, but relative to how hard your classmates (actually, "class enemies" seems a more fitting term in this context) work. All of you recognize this, and after the first lecture you hold an impromptu meeting in which all students agree not to work too hard. As weeks pass by, the temptation to get an edge on the rest of the class by working just that little bit harder becomes overwhelming. After all, the others are not able to observe your work in any detail; nor do they have any real hold over you. And the benefits of an improvement in your grade point average are substantial. So you hit the library more often and stay up a little longer.

The trouble is, everyone else is doing the same. Therefore your grade is no better than it would have been if you and everyone else had abided by the agreement. The only difference is that all of you have spent more time working than you would have liked.

This is an example of the prisoners' dilemma.[2] In the original story, two suspects are being separately interrogated and invited to confess. One of them, say A, is told, "If the other suspect, B, does not confess, then you can cut a very good deal for yourself by confessing. But if B does confess, then you would do well to confess, too; otherwise the court will be especially tough on you. So you should confess no matter what the other does." B is told to confess, with the use of similar reasoning. Faced with this choice, both A and B confess. But it would have been better for both if neither had confessed, because the police had no really compelling evidence against them.

Your situation is similar. If the others slack off, then you can get a much better grade by working hard; if the others work hard, then you had better do the same or else you will get a very bad grade. You may even think that the label "prisoner" is very fitting for a group of students trapped in a required course.

There is a prisoners' dilemma for professors and schools, too. Each professor can make his course look good or attractive by grading it slightly more liberally, and each school can place its students in better jobs or attract better applicants by grading all of its courses a little more liberally. Of course, when all do this, none has any advantage over the others; the only result is rampant grade inflation, which compresses the spectrum of grades and therefore makes it difficult to distinguish abilities.

People often think that in every game there must be a winner and a loser. The prisoners' dilemma is different—both or all players can come out losers. People play (and lose) such games every day, and the losses can range from minor inconveniences to potential disasters. Spectators at a sports event stand up to get a better view but, when all stand, no one has a better view than when they were all sitting. Superpowers acquire more weapons to get an edge over their rivals but, when both do so, the balance of power is unchanged; all that has happened is that both have spent economic resources that they could have used for better purposes, and the risk of accidental war has escalated. The magnitude of the potential cost of such games to all players makes it important to understand the ways in which mutually beneficial cooperation can be achieved and sustained. All of Chapter 11 deals with the study of this game.

Just as the prisoners' dilemma is potentially a lose-lose game, there are win-win games, too. International trade is an example; when each country produces more of what it can do relatively best, all share in the fruits of this international

[2]There is some disagreement regarding the appropriate grammatical placement of the apostrophe in the term *prisoners' dilemma*. Our placement acknowledges the facts that there must be at least two prisoners in order for there to be any dilemma at all and that the (at least two) prisoners therefore jointly possess the dilemma.

division of labor. But successful bargaining about the division of the pie is needed if the full potential of trade is to be realized. The same applies to many other bargaining situations. We will study these in Chapter 18.

C. "We Can't Take the Exam, Because We Had a Flat Tire"

Here is a story, probably apocryphal, that circulates on the undergraduate e-mail networks; each of us independently received it from our students:

> There were two friends taking chemistry at Duke. Both had done pretty well on all of the quizzes, the labs, and the midterm, so that going into the final they each had a solid A. They were so confident the weekend before the final that they decided to go to a party at the University of Virginia. The party was so good that they overslept all day Sunday, and got back too late to study for the chemistry final that was scheduled for Monday morning. Rather than take the final unprepared, they went to the professor with a sob story. They said they each had gone up to UVA and had planned to come back in good time to study for the final but had had a flat tire on the way back. Because they didn't have a spare, they had spent most of the night looking for help. Now they were really too tired, so could they please have a makeup final the next day? The professor thought it over and agreed.
>
> The two studied all of Monday evening and came well prepared on Tuesday morning. The professor placed them in separate rooms and handed the test to each. The first question on the first page, worth 10 points, was very easy. Each of them wrote a good answer, and greatly relieved, turned the page. It had just one question, worth 90 points. It was: "Which tire?"

The story has two important strategic lessons for future party goers. The first is to recognize that the professor may be an intelligent game player. He may suspect some trickery on the part of the students and may use some device to catch them. Given their excuse, the question was the likeliest such device. They should have foreseen it and prepared their answer in advance. This idea that one should look ahead to future moves in the game and then reason backward to calculate one's best current action is a very general principle of strategy, which we will elaborate on in Chapter 3. We will also use it, most notably, in Chapter 10.

But it may not be possible to foresee all such professorial countertricks; after all, professors have much more experience of seeing through students' excuses than students have of making up such excuses. If the pair are unprepared, can they independently produce a mutually consistent lie? If each picks a tire at random, the chances are only 25% that the two will pick the same one. (Why?) Can they do better?

You may think that the front tire on the passenger side is the one most likely to suffer a flat, because a nail or a shard of glass is more likely to lie closer to the side of the road than the middle and the front tire on that side will encounter it first. You may think this is good logic, but that is not enough to make it a good choice. What matters is not the logic of the choice but making the same choice as your friend does. Therefore you have to think about whether your friend would use the same logic and would consider that choice equally obvious. But even that is not the end of the chain of reasoning. Would your friend think that the choice would be equally obvious to you? And so on. The point is not whether a choice is obvious or logical, but whether it is obvious to the other that it is obvious to you that it is obvious to the other. . . . In other words, what is needed is a convergence of expectations about what should be chosen in such circumstances. Such a commonly expected strategy on which the players can successfully coordinate is called a focal point.

There is nothing general or intrinsic to the structure of all such games that creates such convergence. In some games, a focal point may exist because of chance circumstances about the labeling of strategies or because of some experience or knowledge shared by the players. For example, if the passenger's front side of a car were for some reason called the Duke's side, then two Duke students would be very likely to choose it without any need for explicit prior understanding. Or, if the driver's front side of all cars were painted orange (for safety, to be easily visible to oncoming cars), then two Princeton students would be very likely to choose that tire, because orange is the Princeton color. But without some such clue, tacit coordination might not be possible at all.

We will study focal points in more detail in Chapter 4. Here in closing we merely point out that when asked in classrooms, more than 50% of students choose the driver's front side. They are generally unable to explain why, except to say that it seems the obvious choice.

D. Why Are Professors So Mean?

Many professors have inflexible rules not to give makeup exams and never to accept late submission of problem sets or term papers. Students think the professors must be really hardhearted to behave in this way. The true strategic reason is often exactly the opposite. Most professors are kindhearted and would like to give their students every reasonable break and accept any reasonable excuse. The trouble lies in judging what is reasonable. It is hard to distinguish between similar excuses and almost impossible to verify their truth. The professor knows that on each occasion he will end up by giving the student the benefit of the doubt. But the professor also knows that this is a slippery slope. As the students come to know that the professor is a soft touch, they will procrastinate more and produce ever-flimsier excuses. Deadlines will cease to mean anything, and examinations will become a chaotic mix of postponements and makeup tests.

Often the only way to avoid this slippery slope is to refuse to take even the first step down it. Refusal to accept any excuses at all is the only realistic alternative to accepting them all. By making an advance commitment to the "no excuses" strategy, the professor avoids the temptation to give in to all.

But how can a softhearted professor maintain such a hardhearted commitment? He must find some way to make a refusal firm and credible. The simplest way is to hide behind an administrative procedure or university-wide policy. "I wish I could accept your excuse, but the university won't let me" not only puts the professor in a nicer light, but removes the temptation by genuinely leaving him no choice in the matter. Of course, the rules may be made by the same collectivity of professors as hides behind them but, once they are made, no individual professor can unmake the rules in any particular instance.

If the university does not provide such a general shield, then the professor can try to make up commitment devices of his own. For example, he can make a clear and firm announcement of the policy at the beginning of the course. Any time an individual student asks for an exception, he can invoke a fairness principle, saying, "If I do this for you, I would have to do it for everyone." Or the professor can acquire a reputation for toughness by acting tough a few times. This may be an unpleasant thing for him to do and it may run against his true inclination, but it helps in the long run over his whole career. If a professor is believed to be tough, few students will try excuses on him, so he will actually suffer less pain in denying them.

We will study commitments, and related strategies, such as threats and promises, in considerable detail in Chapter 10.

E. Roommates and Families on the Brink

You are sharing an apartment with one or more other students. You notice that the apartment is nearly out of dishwasher detergent, paper towels, cereal, beer, and other items. You have an agreement to share the actual expenses, but the trip to the store takes time. Do you spend your time or do you hope that someone else will spend his, leaving you more time to study or relax? Do you go and buy the soap or stay in and watch TV to catch up on the soap operas?[3]

In many situations of this kind, the waiting game goes on for quite a while before someone who is really impatient for one of the items (usually beer) gives in and spends the time for the shopping trip. Things may deteriorate to the point of serious quarrels or even breakups among the roommates.

This game of strategy can be viewed from two perspectives. In one, each of the roommates is regarded as having a simple binary choice—to do

[3]This example comes from Michael Grunwald's "At Home" column, "A Game of Chicken," in the *Boston Globe Magazine*, April 28, 1996.

the shopping or not. The best outcome for you is where someone else does the shopping and you stay at home; the worst is where you do the shopping while the others get to use their time better. If both do the shopping (unknown to each other, on the way home from school or work), there is unnecessary duplication and perhaps some waste of perishables; if neither does, there can be serious inconvenience or even disaster if the toilet paper runs out at a crucial time.

This is analogous to the game of chicken that used to be played by American teenagers. Two of them drove their cars toward each other. The first to swerve to avoid a collision was the loser (chicken); the one who kept driving straight was the winner. We will analyze the game of chicken further in Chapter 4 and in Chapters 7 and 13.

A more interesting dynamic perspective on the same situation regards it as a "war of attrition," where each roommate tries to wait out the others, hoping that someone else's patience will run out first. In the meantime, the risk escalates that the apartment will run out of something critical, leading to serious inconvenience or a blowup. Each player lets the risk escalate to the point of his own tolerance; the one revealed to have the least tolerance loses. Each sees how close to the brink of disaster the others will let the situation go. Hence the name "brinkmanship" for this strategy and this game. It is a dynamic version of chicken, offering richer and more interesting possibilities.

One of us (Dixit) was privileged to observe a brilliant example of brinkmanship at a dinner party one Saturday evening. Before dinner, the company was sitting in the living room when the host's fifteen-year-old daughter appeared at the door and said, "Bye, Dad." The father asked, "Where are you going?" and the daughter replied, "Out." After a pause that was only a couple of seconds but seemed much longer, the host said, "All right, bye."

Your strategic observer of this scene was left thinking how it might have gone differently. The host might have asked, "With whom?" and the daughter might have replied, "Friends." The father could have refused permission unless the daughter told him exactly where and with whom she would be. One or the other might have capitulated at some such later stage of this exchange or it could have led to a blowup.

This was a risky game for both the father and the daughter to play. The daughter might have been punished or humiliated in front of strangers; an argument could have ruined the father's evening with his friends. Each had to judge how far to push the process, without being fully sure whether and when the other might give in or whether there would be an unpleasant scene. The risk of an explosion would increase as the father tried harder to force the daughter to answer and as she defied each successive demand.

In this respect the game played by the father and the daughter was just like that between a union and a company's management who are negotiating a labor

contract or between two superpowers who are encroaching on each other's sphere of influence in the world. Neither side can be fully sure of the other's intentions, so each side explores them through a succession of small incremental steps, each of which escalates the risk of mutual disaster. The daughter in our story was exploring previously untested limits of her freedom; the father was exploring previously untested—and perhaps unclear even to himself—limits of his authority.

This was an example of brinkmanship, a game of escalating mutual risk, par excellence. Such games can end in one of two ways. In the first way, one of the players reaches the limit of his own tolerance for risk and concedes. (The father in our story conceded quickly, at the very first step. Other fathers might be more successful strict disciplinarians, and their daughters might not even initiate a game like this.) In the second way, before either has conceded, the risk that they both fear comes about, and the blowup (the strike or the war) occurs. The feud in our host's family ended "happily"; although the father conceded and the daughter won, a blowup would have been much worse for both.

We will analyze the strategy of brinkmanship more fully in Chapter 10; in Chapter 15, we will examine a particularly important instance of it—namely, the Cuban missile crisis of 1962.

F. The Dating Game

When you are dating, you want to show off the best attributes of your personality to your date and to conceal the worst ones. Of course, you cannot hope to conceal them forever if the relationship progresses; but you are resolved to improve or hope that by that stage the other person will accept the bad things about you with the good ones. And you know that the relationship will not progress at all unless you make a good first impression; you won't get a second chance to do so.

Of course, you want to find out everything, good and bad, about the other person. But you know that, if the other is as good at the dating game as you are, he or she will similarly try to show the best side and hide the worst. You will think through the situation more carefully and try to figure out which signs of good qualities are real and which ones can easily be put on for the sake of making a good impression. Even the worst slob can easily appear well groomed for a big date; ingrained habits of courtesy and manners that are revealed in a hundred minor details may be harder to simulate for a whole evening. Flowers are relatively cheap; more expensive gifts may have value, not for intrinsic reasons, but as credible evidence of how much the other person is willing to sacrifice for you. And the "currency" in which the gift is given may have different significance, depending on the context; from a millionaire, a diamond may be worth less in this regard than the act of giving up valuable time for your company or time spent on some activity at your request.

You should also recognize that your date will similarly scrutinize your actions for their information content. Therefore you should take actions that are credible signals of your true good qualities, and not just the ones that anyone can imitate.

This is important not just on a first date; revealing, concealing, and eliciting information about the other person's deepest intentions remain important throughout a relationship. Here is a story to illustrate that.

Once upon a time in New York City there lived a man and a woman who had separate rent-controlled apartments, but their relationship had reached the point at which they were using only one of them. The woman suggested to the man that they give up the other apartment. The man, an economist, explained to her a fundamental principle of his subject: it is always better to have more choice available. The probability of their splitting up might be small but, given even a small risk, it would be useful to retain the second low-rent apartment. The woman took this very badly and promptly ended the relationship!

Economists who hear this story say that it just confirms their principle that greater choice is better. But strategic thinking offers a very different and more compelling explanation. The woman was not sure of the man's commitment to the relationship, and her suggestion was a brilliant strategic device to elicit the truth. Words are cheap; anyone can say, "I love you." If the man had put his property where his mouth was and had given up his rent-controlled apartment, that would have been concrete evidence of his love. The fact that he refused to do so constituted hard evidence of the opposite, and the woman did right to end the relationship.

These are examples, designed to appeal to your immediate experience, of a very important class of games—namely, those where the real strategic issue is manipulation of information. Strategies that convey good information about yourself are called signals; strategies that induce others to act in ways that will credibly reveal their private information, good or bad, are called screening devices. Thus the woman's suggestion of giving up one of the apartments was a screening device, which put the man in the situation of offering to give up his apartment or else revealing his lack of commitment. We will study games of information, as well as signaling and screening, in Chapters 9 and 14.

3 OUR STRATEGY FOR STUDYING GAMES OF STRATEGY

We have chosen several examples that relate to your experiences as amateur strategists in real life to illustrate some basic concepts of strategic thinking and strategic games. We could continue, building a whole stock of dozens of similar stories. The hope would be that, when you face an actual strategic situation,

you might recognize a parallel with one of these stories, which would help you decide the appropriate strategy for your own situation. This is the *case study* approach taken by most business schools. It offers a concrete and memorable vehicle for the underlying concepts. However, each new strategic situation typically consists of a unique combination of so many variables that an intolerably large stock of cases is needed to cover all of them.

An alternative approach focuses on the general principles behind the examples and so constructs a *theory* of strategic action—namely, formal game theory. The hope here is that, facing an actual strategic situation, you might recognize which principle or principles apply to it. This is the route taken by the more academic disciplines, such as economics and political science. A drawback to this approach is that the theory is presented in a very abstract and mathematical manner, without enough cases or examples. This makes it difficult for most beginners to understand or remember the theory and to connect the theory with reality afterward.

But knowing some general theory has an overwhelming compensating advantage. It gives you a deeper understanding of games and of *why* they have the outcomes they do. This helps you play better than you would if you merely read some cases and knew the recipes for *how* to play some specific games. With the knowledge of why, you can think through new and unexpected situations where a mechanical follower of a "how" recipe would be lost. A world champion of checkers, Tom Wiswell, has expressed this beautifully: "The player who knows how will usually draw, the player who knows why will usually win."[4] This is not to be taken literally for all games; some games may be hopeless situations for one of the players no matter how knowledgable he may be. But the statement contains the germ of an important general truth—knowing why gives you an advantage beyond what you can get if you merely know how. For example, knowing the why of a game can help you foresee a hopeless situation and avoid getting into such a game in the first place.

Therefore we will take an intermediate route that combines some of the advantages of both approaches—case studies (how) and theory (why). We will organize the subject around its general principles, generally one in each of the chapters to follow. Therefore you don't have to figure them out on your own from the cases. But we will develop the general principles through illustrative cases rather than abstractly, so the context and scope of each idea will be clear and evident. In other words, we will focus on theory but build it up through cases, not abstractly.

Of course, such an approach requires some compromises of its own. Most important, you should remember that each of our examples serves the purpose of conveying some general idea or principle of game theory. Therefore we will

[4] Quoted in Victor Niederhoffer, *The Education of a Speculator* (New York: Wiley, 1997), p. 169. We thank Austin Jaffe of Pennsylvania State University for bringing this aphorism to our attention.

leave out many details of each case that are incidental to the principle at stake. If some examples seem somewhat artificial, please bear with us; we have generally considered the omitted details and left them out for good reasons.

A word of reassurance. Although the examples that motivate the development of our conceptual or theoretical frameworks are deliberately selected for that purpose (even at the cost of leaving out some other features of reality), once the theory has been constructed, we pay a lot of attention to its connection with reality. Throughout the book, we examine factual and experimental evidence in regard to how well the theory explains reality. The frequent answer—very well in some respects and less well in others—should give you cautious confidence in using the theory and should be a spur to contributing to the formulation of better theories. In appropriate places, we examine in great detail how institutions evolve in practice to solve some problems pointed out by the theories; note in particular the discussion in Chapter 11 of how prisoners' dilemmas arise and are solved in reality and a similar discussion of more general collective-action problems in Chapter 12. Finally, in Chapter 15 we will examine the use of brinkmanship in the Cuban missile crisis. Such theory-based case studies, which take rich factual details of a situation and subject them to an equally detailed theoretical analysis, are becoming common in such diverse fields as business studies, political science, and economic history; we hope our original study of an important episode in the diplomatic and military areas will give you an interesting introduction to this genre.

To pursue our approach, in which examples lead to general theories that are then tested against reality and used to interpret reality, we must first identify the general principles that serve to organize the discussion. We will do so in Chapter 2 by classifying or dichotomizing games along several key dimensions of different strategic matters or concepts. Along each dimension, we will identify two extreme pure types. For example, one such dimension concerns the order of moves, and the two pure types are those in which the players take turns making moves (sequential games) and those in which all players act at once (simultaneous games). Actual games rarely correspond to exactly one of these conceptual categories; most partake of some features of each extreme type. But each game can be located in our classification by considering which concepts or dimensions bear on it and how it mixes the two pure types in each dimension. To decide how to act in a specific situation, one then combines in appropriate ways the lessons learned for the pure types.

Once this general framework has been constructed in Chapter 2, the chapters that follow will build on it, developing several general ideas and principles for each player's strategic choice and the interaction of all players' strategies in games.

2

How to Think About Strategic Games

CHAPTER 1 GAVE SOME simple examples of strategic games and strategic thinking. In this chapter, we begin a more systematic and analytical approach to the subject. We choose some crucial conceptual categories or dimensions, in each of which there is a dichotomy of types of strategic interactions. For example, one such dimension concerns the timing of the players' actions, and the two pure types are games where the players act in strict turns (sequential moves) and where they act at the same time (simultaneous moves). We consider some matters that arise in thinking about each pure type in this dichotomy, as well as in similar dichotomies, with respect to other matters, such as whether the game is played only once or repeatedly and what the players know about each other.

In the chapters that follow, we will examine each of these categories or dimensions in more detail and show how the analysis can be used in several specific applications. Of course, most actual applications are not of a pure type but rather a mixture. Moreover, in each application, two or more of the categories have some relevance. The lessons learned from the study of the pure types must therefore be combined in appropriate ways. We will show how to do this by using the context of our applications.

In this chapter, we state some basic concepts and terminology—such as strategies, payoffs, and equilibrium—that are used in the analysis and briefly describe solution methods. We also provide a brief discussion of the uses of game theory and an overview of the structure of the remainder of the book.

1 DECISIONS VERSUS GAMES

When a person (or team or firm or government) decides how to act in dealings with other people (or teams or firms or governments), there must be some cross-effect of their actions; what one does must affect the outcome for the other. When George Pickett (of Pickett's Charge at the battle of Gettysburg) was asked to explain the Confederacy's defeat in the Civil War, he responded, "I think the Yankees had something to do with it."[1]

For the interaction to become a strategic game, however, we need something more—namely, the participants' mutual awareness of this cross-effect. What the other person does affects you; if you know this, you can react to his actions, or take advance actions to forestall the bad effects his future actions may have on you and to facilitate any good effects, or even take advance actions so as to alter his future reactions to your advantage. If you know that the other person knows that what you do affects him, you know that he will be taking similar actions. And so on. It is this mutual awareness of the cross-effects of actions and the actions taken as a result of this awareness that constitute the most interesting aspects of strategy.

This distinction is captured by reserving the label **strategic games** (or sometimes just **games**, because we are not concerned with other types of games, such as those of pure chance or pure skill) for interactions between mutually aware players and **decisions** for action situations where each person can choose without concern for reaction or response from others. If Robert E. Lee (who ordered Pickett to lead the ill-fated Pickett's Charge) had thought that the Yankees had been weakened by his earlier artillery barrage to the point that they no longer had any ability to resist, his choice to attack would have been a decision; if he was aware that the Yankees were prepared and waiting for his attack, then the choice became a part of a (deadly) game. The simple rule is that unless there are two or more players, each of whom responds to what others do (or what each thinks the others might do), it is not a game.

Strategic games arise most prominently in head-to-head confrontations of two participants: the arms race between the United States and the Soviet Union from the 1950s through the 1980s; wage negotiations between General Motors and the United Auto Workers; or a Super Bowl matchup between two "pirates," the Tampa Bay Buccaneers and the Oakland Raiders. In contrast, interactions among a large number of participants seem less susceptible to the issues raised by mutual awareness. Because each farmer's output is an insignificant part of

[1]James M. McPherson, "American Victory, American Defeat," in *Why the Confederacy Lost*, ed. Gabor S. Boritt (New York: Oxford University Press, 1993), p. 19.

the whole nation's or the world's output, the decision of one farmer to grow more or less corn has almost no effect on the market price, and not much appears to hinge on thinking of agriculture as a strategic game. This was indeed the view prevalent in economics for many years. A few confrontations between large companies—as in the U.S. auto market, which was once dominated by GM, Ford, and Chrysler—were usefully thought of as strategic games, but most economic interactions were supposed to be governed by the impersonal forces of supply and demand.

In fact, game theory has a much greater scope. Many situations that start out as impersonal markets with thousands of participants turn into strategic interactions of two or just a few. This happens for one of two broad classes of reasons—mutual commitments or private information.

Consider commitment first. When you are contemplating building a house, you can choose one of several dozen contractors in your area; the contractor can similarly choose from several potential customers. There appears to be an impersonal market. Once each side has made a choice, however, the customer pays an initial installment, and the builder buys some materials for the plan of this particular house. The two become tied to each other, separately from the market. Their relationship becomes *bilateral*. The builder can try to get away with a somewhat sloppy job or can procrastinate, and the client can try to delay payment of the next installment. Strategy enters the picture. Their initial contract has to anticipate their individual incentives and specify a schedule of installments of payments that are tied to successive steps in the completion of the project. Even then, some adjustments have to be made after the fact, and these adjustments bring in new elements of strategy.

Next, consider private information. Thousands of farmers seek to borrow money for their initial expenditures on machinery, seed, fertilizer, and so forth, and hundreds of banks exist to lend to them. Yet the market for such loans is not impersonal. A borrower with good farming skills who puts in a lot of effort will be more likely to be successful and will repay the loan; a less-skilled or lazy borrower may fail at farming and default on the loan. The risk of default is highly personalized. It is not a vague entity called "the market" that defaults, but individual borrowers who do so. Therefore each bank will have to view its lending relation with each individual borrower as a separate game. It will seek collateral from each borrower or will investigate each borrower's creditworthiness. The farmer will look for ways to convince the bank of his quality as a borrower; the bank will look for effective ways to ascertain the truth of the farmer's claims.

Similarly, an insurance company will make some efforts to determine the health of individual applicants and will check for any evidence of arson when a claim for a fire is made; an employer will inquire into the qualifications of individual employees and monitor their performance. More generally, when participants in a transaction possess some private information bearing on the outcome, each bilateral deal becomes a game of strategy, even though the larger picture may have thousands of very similar deals going on.

To sum up, when each participant is significant in the interaction, either because each is a large player to start with or because commitments or private information narrow the scope of the relationship to a point where each is an important player *within* the relationship, we must think of the interaction as a strategic game. Such situations are the rule rather than the exception in business, in politics, and even in social interactions. Therefore the study of strategic games forms an important part of all fields that analyze these matters.

2 CLASSIFYING GAMES

Games of strategy arise in many different contexts and accordingly have many different features that require study. This task can be simplified by grouping these features into a few categories or dimensions, along each of which we can identify two pure types of games and then recognize any actual game as a mixture of the pure types. We develop this classification by asking a few questions that will be pertinent for thinking about the actual game that you are playing or studying.

A. Are the Moves in the Game Sequential or Simultaneous?

Moves in chess are sequential: White moves first, then Black, then White again, and so on. In contrast, participants in an auction for an oil-drilling lease or a part of the airwave spectrum make their bids simultaneously, in ignorance of competitors' bids. Most actual games combine aspects of both. In a race to research and develop a new product, the firms act simultaneously, but each competitor has partial information about the others' progress and can respond. During one play in football, the opposing offensive and defensive coaches simultaneously send out teams with the expectation of carrying out certain plays but, after seeing how the defense has set up, the quarterback can change the play at the line of scrimmage or call a time-out so that the coach can change the play.

The distinction between **sequential** and **simultaneous moves** is important because the two types of games require different types of interactive thinking. In a sequential-move game, each player must think: if I do this, how will my opponent react? Your current move is governed by your calculation of its *future* consequences. With simultaneous moves, you have the trickier task of trying to figure out what your opponent is going to do *right now*. But you must recognize that, in making his own calculation, the opponent is also trying to figure out your current move, while at the same time recognizing that you are doing the same with him. . . . Both of you have to think your way out of this circle.

In the next three chapters, we will study the two pure cases. In Chapter 3, we examine sequential-move games, where you must look ahead to act now; in Chapters 4 and 5, the subject is simultaneous-move games, where you must square the circle of "He thinks that I think that he thinks . . ." In each case, we will devise some simple tools for such thinking—trees and payoff tables—and obtain some simple rules to guide actions.

The study of sequential games also tells us when it is an advantage to move first and when second. Roughly speaking, this depends on the relative importance of commitment and flexibility in the game in question. For example, the game of economic competition among rival firms in a market has a first-mover advantage if one firm, by making a firm commitment to compete aggressively, can get its rivals to back off. But, in political competition, a candidate who has taken a firm stand on an issue may give his rivals a clear focus for their attack ads, and the game has a second-mover advantage.

Knowledge of the balance of these considerations can also help you devise ways to manipulate the order of moves to your own advantage. That in turn leads to the study of strategic moves, such as threats and promises, which we will take up in Chapter 10.

B. Are the Players' Interests in Total Conflict, or Is There Some Commonality?

In simple games such as chess or football, there is a winner and a loser. One player's gain is the other's loss. Similarly, in gambling games, one player's winnings are the others' losses, so the total is zero. This is why such situations are called *zero-sum games*. More generally, the idea is that the players' interests are in complete conflict. Such conflict arises when players are dividing up any fixed amount of possible gain, whether it be measured in yards, dollars, acres, or scoops of ice cream. Because the available gain need not always be exactly zero, the term *constant-sum game* is often substituted for zero-sum; we will use the two interchangeably.

Most economic and social games are not zero-sum. Trade, or economic activity more generally, offers scope for deals that benefit everyone. Joint ventures can combine the participants' different skills and generate synergy to produce more than the sum of what they could have produced separately. But the interests are not completely aligned either; the partners can cooperate to create a larger total pie, but they will clash when it comes to deciding how to split this pie among them.

Even wars and strikes are not zero-sum games. A nuclear war is the most striking example of a situation where there can be only losers, but the concept is far older. Pyrrhus, the king of Epirus, defeated the Romans at Heraclea in 280 B.C. but at such great cost to his own army that he exclaimed, "Another such victory and we are lost!" Hence the phrase "Pyrrhic victory." In the 1980s, at the height of the frenzy of business takeovers, the battles among rival bidders led to such costly escalation that the successful bidder's victory was often similarly Pyrrhic.

Most games in reality have this tension between conflict and cooperation, and many of the most interesting analyses in game theory come from the need to handle it. The players' attempts to resolve their conflict—distribution of territory or profit—are influenced by the knowledge that, if they fail to agree, the outcome will be bad for all of them. One side's threat of a war or a strike is its attempt to frighten the other side into conceding its demands.

Even when a game is constant-sum for all players, when there are three (or more) players, we have the possibility that two of them will cooperate at the expense of the third; this leads to the study of alliances and coalitions. We will examine and illustrate these ideas later, especially in Chapter 18 on bargaining.

C. Is the Game Played Once or Repeatedly, and with the Same or Changing Opponents?

A game played just once is in some respects simpler and in others more complicated than one with a longer interaction. You can think about a one-shot game without worrying about its repercussions on other games you might play in the future against the same person or against others who might hear of your actions in this one. Therefore actions in one-shot games are more likely to be unscrupulous or ruthless. For example, an automobile repair shop is much more likely to overcharge a passing motorist than a regular customer.

In one-shot encounters, each player doesn't know much about the others; for example, what their capabilities and priorities are, whether they are good at calculating their best strategies or have any weaknesses that can be exploited, and so on. Therefore in one-shot games, secrecy or surprise is likely to be an important component of good strategy.

Games with ongoing relationships require the opposite considerations. You have an opportunity to build a reputation (for toughness, fairness, honesty, reliability, and so forth, depending on the circumstances) and to find out more about your opponent. The players together can better exploit mutually beneficial prospects by arranging to divide the spoils over time (taking turns to "win") or to punish a cheater in future plays (an eye for an eye or tit-for-tat). We will consider these possibilities in Chapter 11 on the prisoners' dilemma.

More generally, a game may be zero-sum in the short run but have scope for mutual benefit in the long run. For example, each football team likes to win, but they all recognize that close competition generates more spectator interest, which benefits all teams in the long run. That is why they agree to a drafting scheme where teams get to pick players in reverse order of their current standing, thereby reducing the inequality of talent. In long-distance races, the runners or cyclists often develop a lot of cooperation; two or more of them can help one another by taking turns following in one another's slipstream. Near the end of the race, the cooperation collapses as all of them dash for the finish line.

Here is a useful rule of thumb for your own strategic actions in life. In a game that has some conflict and some scope for cooperation, you will often think up a great strategy for winning big and grinding a rival into dust but have a nagging worry at the back of your mind that you are behaving like the worst 1980s yuppie. In such a situation, the chances are that the game has a repeated or ongoing aspect that you have overlooked. Your aggressive strategy may gain you a short-run advantage, but its long-run side effects will cost you even more. Therefore you should dig deeper and recognize the cooperative element and then alter your strategy accordingly. You will be surprised how often niceness, integrity, and the golden rule of doing to others as you would have them do to you turn out to be not just old nostrums, but good strategies as well, when you consider the whole complex of games that you will be playing in the course of your life.

D. Do the Players Have Full or Equal Information?

In chess, each player knows exactly the current situation and all the moves that led to it, and each knows that the other aims to win. This situation is exceptional; in most other games, the players face some limitation of information. Such limitations come in two kinds. First, a player may not know all the information that is pertinent for the choice that he has to make at every point in the game. This type of information problem arises because of the player's uncertainty about relevant variables, both internal and external to the game. For example, he may be uncertain about external circumstances, such as the weekend weather or the quality of a product he wishes to purchase; we call this situation one of **external uncertainty.** Or he may be uncertain about exactly what moves his opponent has made in the past or is making at the same time he makes his own move; we call this **strategic uncertainty.** If a game has neither external nor strategic uncertainty, we say that the game is one of **perfect information;** otherwise the game has **imperfect information.** We will give a more precise technical definition of perfect information in Chapter 6, Section 3.A, after we have introduced the concept of an information set. We will develop the theory of games with imperfect information (uncertainty) in three future chapters. In Chapter 4, we discuss games with contemporaneous (simultaneous) actions, which entail strategic uncertainty, and we analyze methods for making choices under external uncertainty in the Appendix to Chapter 7 and in Chapter 9.

Trickier strategic situations arise when one player knows more than another does; they are called situations of **incomplete** or, better, **asymmetric information.** In such situations, the players' attempts to infer, conceal, or sometimes convey their private information become an important part of the game and the strategies. In bridge or poker, each player has only partial knowledge of the cards held by the others. Their actions (bidding and play in bridge, the number of cards taken and the betting behavior in poker) give information to opponents. Each player tries to manipulate his actions to mislead the opponents (and, in bridge, to inform his partner truthfully), but in doing so each must be aware that the opponents know this and that they will use strategic thinking to interpret that player's actions.

You may think that if you have superior information, you should always conceal it from other players. But that is not true. For example, suppose you are the CEO of a pharmaceutical firm that is engaged in an R&D competition to develop a new drug. If your scientists make a discovery that is a big step forward, you may want to let your competitors know, in the hope that they will give up their own searches. In war, each side wants to keep its tactics and troop deployments secret; but, in diplomacy, if your intentions are peaceful, then you desperately want other countries to know and believe this fact.

The general principle here is that you want to release your information selectively. You want to reveal the good information (the kind that will draw responses from the other players that work to your advantage) and conceal the bad (the kind that may work to your disadvantage).

This raises a problem. Your opponents in a strategic game are purposive, rational players and know that you are one, too. They will recognize your incentive to exaggerate or even to lie. Therefore they are not going to accept your unsupported declarations about your progress or capabilities. They can be convinced only by objective evidence or by actions that are credible proof of your information. Such actions on the part of the more-informed player are called **signals,** and strategies that use them are called **signaling.** Conversely, the less-informed party can create situations in which the more-informed player will have to take some action that credibly reveals his information; such strategies are called **screening,** and the methods they use are called **screening devices.** The word *screening* is used here in the sense of testing in order to sift or separate, not in the sense of concealing. Recall that in the dating game in Section 2.F of Chapter 1, the woman was screening the man to test his commitment to their relationship, and her suggestion that the pair give up one of their two rent-controlled apartments was the screening device. If the man had been committed to the relationship, he might have acted first and volunteered to give up his apartment; this action would have been a signal of his commitment.

Now we see how, when different players have different information, the manipulation of information itself becomes a game, perhaps more important than the game that will be played after the information stage. Such information games are ubiquitous, and playing them well is essential for success in life. We will study more games of this kind in greater detail in Chapter 9 and also in Chapter 14.

E. Are the Rules of the Game Fixed or Manipulable?

The rules of chess, card games, or sports are given, and every player must follow them, no matter how arbitrary or strange they seem. But in games of business, politics, and ordinary life, the players can make their own rules to a greater or lesser extent. For example, in the home, parents constantly try to make the rules, and children constantly look for ways to manipulate or circumvent those rules. In legislatures, rules for the progress of a bill (including the order in which amendments and main motions are voted on) are fixed, but the game that sets the agenda—which amendments are brought to a vote first—can be manipulated; that is where political skill and power have the most scope, and we will address these matters in detail in Chapter 16.

In such situations, the real game is the "pregame" where rules are made, and your strategic skill must be deployed at that point. The actual playing out of the subsequent game can be more mechanical; you could even delegate it to someone else. However, if you "sleep" through the pregame, you might find that you have lost the game before it ever began. For many years, American firms ignored the rise of foreign competition in just this way and ultimately paid the price. Others, such as oil magnate John D. Rockefeller, Sr., adopted the strategy of limiting their participation to games in which they could also participate in making the rules.[2]

The distinction between changing rules and acting within the chosen rules will be most important for us in our study of strategic moves, such as threats and promises. Questions of how you can make your own threats and promises credible or how you can reduce the credibility of your opponent's threats basically have to do with a pregame that entails manipulating the rules of the subsequent game in which the promises or threats may have to be carried out. More generally, the strategic moves that we will study in Chapter 10 are essentially ploys for such manipulation of rules.

But if the pregame of rule manipulation is the real game, what fixes the rules of the pregame? Usually these pregame rules depend on some hard facts related to the players' innate abilities. In business competition, one firm can take preemptive actions that alter subsequent games between it and its rivals; for example, it can expand its factory or advertise in a way that twists the results of subsequent price competition more favorably to itself. Which firm can do this best or most easily depends on which one has the managerial or organizational resources to make the investments or to launch the advertising campaigns.

Players may also be unsure of their rivals' abilities. This often makes the pregame one of unequal information, requiring more subtle strategies and occasionally resulting in some big surprises. We will comment on all these matters in the appropriate places in the chapters that follow.

[2]For more on the methods used in Rockefeller's rise to power, see Ron Chernow, *Titan* (New York: Random House, 1998).

F. Are Agreements to Cooperate Enforceable?

We saw that most strategic interactions consist of a mixture of conflict and common interest. Then there is a case to be made that all participants should get together and reach an agreement about what everyone should do, balancing their mutual interest in maximizing the total benefit and their conflicting interests in the division of gains. Such negotiations can take several rounds in which agreements are made on a tentative basis, better alternatives are explored, and the deal is finalized only when no group of players can find anything better. The concept of the core in Chapter 19 embodies such a process and its outcome. However, even after the completion of such a process, additional difficulties often arise in putting the final agreement into practice. For instance, all the players must perform, in the end, the actions that were stipulated for them in the agreement. When all others do what they are supposed to do, any one participant can typically get a better outcome for himself by doing something different. And, if each one suspects that the others may cheat in this way, he would be foolish to adhere to his stipulated cooperative action.

Agreements to cooperate can succeed if all players act immediately and in the presence of the whole group, but agreements with such immediate implementation are quite rare. More often the participants disperse after the agreement has been reached and then take their actions in private. Still, if these actions are observable to the others and a third party—for example, a court of law—can enforce compliance, then the agreement of joint action can prevail.

However, in many other instances individual actions are neither directly observable nor enforceable by external forces. Without enforceability, agreements will stand only if it is in all participants' individual interests to abide by them. Games among sovereign countries are of this kind, as are many games with private information or games where the actions are either outside the law or too trivial or too costly to enforce in a court of law. In fact, games where agreements for joint action are not enforceable constitute a vast majority of strategic interactions.

Game theory uses a special terminology to capture the distinction between situations in which agreements are enforceable and those in which they are not. Games in which joint-action agreements are enforceable are called **cooperative;** those in which such enforcement is not possible, and individual participants must be allowed to act in their own interests, are called **noncooperative.** This has become standard terminology, but it is somewhat unfortunate because it gives the impression that the former will produce cooperative outcomes and the latter will not. In fact, individual action can be compatible with the achievement of a lot of mutual gain, especially in repeated interactions. The important distinction is that in so-called noncooperative games, cooperation will emerge only if it is in the participants' separate and individual interests to continue to take the prescribed actions. This emergence of cooperative outcomes from noncooperative behavior is one of the most interesting findings of game theory, and we will develop the idea in Chapters 11, 12, and 13.

We will adhere to the standard usage, but emphasize that the terms *cooperative* and *noncooperative* refer to the way in which actions are implemented or enforced—collectively in the former mode and individually in the latter—and not to the nature of the outcomes.

As we said earlier, most games in practice do not have adequate mechanisms for external enforcement of joint-action agreements. Therefore most of our analytical development will proceed in the noncooperative mode. The few exceptions include the discussion of bargaining in Chapter 18 and a brief treatment of markets and competition in Chapter 19.

3 SOME TERMINOLOGY AND BACKGROUND ASSUMPTIONS

When one thinks about a strategic game, the logical place to begin is by specifying its structure. This includes all the strategies available to all the players, their information, and their objectives. The first two aspects will differ from one game to another along the dimensions discussed in the preceding section, and one must locate one's particular game within that framework. The objectives raise some new and interesting considerations. Here we consider aspects of all these matters.

A. Strategies

Strategies are simply the choices available to the players, but even this basic notion requires some further study and elaboration. If a game has purely simultaneous moves made only once, then each player's strategy is just the action taken on that single occasion. But if a game has sequential moves, then the actions of a player who moves later in the game can respond to what other players have done (or what he himself has done) at earlier points. Therefore each such player must make a complete plan of action, for example: "If the other does A, then I will do X but, if the other does B, then I will do Y." This complete plan of action constitutes the strategy in such a game.

There is a very simple test to determine whether your strategy is complete. It should specify how you would play the game in such full detail—describing your action in every contingency—that, if you were to write it all down, hand it to someone else, and go on vacation, this other person acting as your representative could play the game just as you would have played it. He would know what to do on each occasion that could conceivably arise in the course of play, without ever needing to disturb your vacation for instructions on how to deal with some situation that you had not foreseen.

This test will become clearer in Chapter 3, when we develop and apply it in some specific contexts. For now, you should simply remember that a strategy is a complete plan of action.

This notion is similar to the common usage of the word *strategy* to denote a longer-term or larger-scale plan of action, as distinct from tactics that pertain to a shorter term or a smaller scale. For example, the military makes strategic plans for a war or a large-scale battle, while tactics for a smaller skirmish or a particular theater of battle are often left to be devised by lower-level officers to suit the local conditions. But game theory does not use the term *tactics* at all. The term *strategy* covers all the situations, meaning a complete plan of action when necessary and meaning a single move if that is all that is needed in the particular game being studied.

The word *strategy* is also commonly used to describe a person's decisions over a fairly long time span and a sequence of choices, even though there is no game in our sense of purposive and aware interaction with other people. Thus you have probably already chosen a career strategy. When you start earning an income, you will make saving and investment strategies and eventually plan a retirement strategy. This usage of the term *strategy* has the same sense as ours—a plan for a succession of actions in response to evolving circumstances. The only difference is that we are reserving it for a situation—namely, a game—where the circumstances evolve because of actions taken by other purposive players.

B. Payoffs

When asked what a player's objective in a game is, most newcomers to strategic thinking respond that it is "to win"; but matters are not always so simple. Sometimes the margin of victory matters; for example, in R&D competition, if your product is only slightly better than the nearest rival's, your patent may be more open to challenge. Sometimes there may be smaller prizes for several participants, so winning isn't everything. Most important, very few games of strategy are purely zero-sum or win-lose; they combine some common interest and some conflict among the players. Thinking about such mixed-motive games requires more refined calculations than the simple dichotomy of winning and losing—for example, comparisons of the gains from cooperating versus cheating.

We will give each player a complete numerical scale with which to compare all logically conceivable outcomes of the game, corresponding to each available combination of choices of strategies by all the players. The number associated with each possible outcome will be called that player's **payoff** for that outcome. Higher payoff numbers attach to outcomes that are better in this player's rating system.

Sometimes the payoffs will be simple numerical ratings of the outcomes, the worst labeled 1, the next worst 2, and so on, all the way to the best. In other games, there may be more natural numerical scales—for example, money income or profit for firms, viewer-share ratings for television networks, and so on. In many situations, the payoff numbers are only educated guesses; then we should do some sensitivity tests by checking that the results of our analysis do not change significantly if we vary these guesses within some reasonable margin of error.

Two important points about the payoffs need to be understood clearly. First, the payoffs for one player capture everything in the outcomes of the game that he cares about. In particular, the player need not be selfish, but his concern about others should be already included in his numerical payoff scale. Second, we will suppose that, if the player faces a random prospect of outcomes, then the number associated with this prospect is the average of the payoffs associated with each component outcome, each weighted by its probability. Thus, if in one player's ranking, outcome A has payoff 0 and outcome B has payoff 100, then the prospect of a 75% probability of A and a 25% probability of B should have the payoff $0.75 \times 0 + 0.25 \times 100 = 25$. This is often called the **expected payoff** from the random prospect. The word *expected* has a special connotation in the jargon of probability theory. It does not mean what you think you will get or expect to get; it is the mathematical or probabilistic or statistical expectation, meaning an average of all possible outcomes, where each is given a weight proportional to its probability.

The second point creates a potential difficulty. Consider a game where players get or lose money and payoffs are measured simply in money amounts. In reference to the preceding example, if a player has a 75% chance of getting nothing and a 25% chance of getting $100, then the expected payoff as calculated in that example is $25. That is also the payoff that the player would get from a simple nonrandom outcome of $25. In other words, in this way of calculating payoffs, a person should be indifferent to whether he receives $25 for sure or faces a risky prospect of which the average is $25. One would think that most people would be averse to risk, preferring a sure $25 to a gamble that yields only $25 on the average.

A very simple modification of our payoff calculation gets around this difficulty. We measure payoffs not in money sums but by using a nonlinear rescaling of the dollar amounts. This is called the expected utility approach, and we will present it in detail in the Appendix to Chapter 7. For now, please take our word that incorporating differing attitudes toward risk into our framework is a manageable task. Almost all of game theory is based on the expected utility approach, and it is indeed very useful, although not without flaws. We will adopt it in this book, but we also indicate some of the difficulties that it leaves unresolved, with the use of a simple example in Chapter 8, Section 6.

C. Rationality

Each player's aim in the game will be to achieve as high a payoff for himself as possible. But how good is each player at pursuing this aim? This question is not about whether and how other players pursuing their own interests will impede him; that is in the very nature of a game of strategic interaction. We mean how good each player is at calculating the strategy that is in his own best interests and at following this strategy in the actual course of play.

Much of game theory assumes that players are perfect calculators and flawless followers of their best strategies. This is the assumption of **rational behavior.** Observe the precise sense in which the term *rational* is being used. It means that each has a consistent set of rankings (values or payoffs) over all the logically possible outcomes and calculates the strategy that best serves these interests. Thus rationality has two essential ingredients: complete knowledge of one's own interests, and flawless calculation of what actions will best serve those interests.

It is equally important to understand what is *not* included in this concept of rational behavior. It does not mean that players are selfish; a player may rate highly the well-being of some other and incorporate this high rating into his payoffs. It does not mean that players are short-run thinkers; in fact, calculation of future consequences is an important part of strategic thinking, and actions that seem irrational from the immediate perspective may have valuable long-term strategic roles. Most important, being rational does not mean having the same value system as other players, or sensible people, or ethical or moral people would use; it means merely pursuing one's own value system consistently. Therefore, when one player carries out an analysis of how other players will respond (in a game with sequential moves) or of the successive rounds of thinking about thinking (in a game with simultaneous moves), he must recognize that the other players calculate the consequences of their choices by using their own value or rating system. You must not impute your own value systems or standards of rationality to others and assume that they would act as you would in that situation. Thus many "experts" commenting on the Persian Gulf conflict in late 1990 predicted that Saddam Hussein would back down "because he is rational"; they failed to recognize that Saddam's value system was different from the one held by most Western governments and by the Western experts.

In general, each player does not really know the other players' value systems; this is part of the reason that in reality many games have incomplete and asymmetric information. In such games, trying to find out the values of others and trying to conceal or convey one's own become important components of strategy.

Game theory assumes that all players are rational. How good is this assumption, and therefore how good is the theory that employs it? At one level, it is obvious that the assumption cannot be literally true. People often don't even have full advance knowledge of their own value systems; they don't think ahead about how they would rank hypothetical alternatives and then remember these rankings until they are actually confronted with a concrete choice. Therefore they find it very difficult to perform the logical feat of tracing all possible con-

sequences of their and other players' conceivable strategic choices and ranking the outcomes in advance in order to choose which strategy to follow. Even if they knew their preferences, the calculation would remain far from easy. Most games in real life are very complex, and most real players are limited in their thinking and computational abilities. In games such as chess, it is known that the calculation for the best strategy can be performed in a finite number of steps, but no one has succeeded in performing it, and good play remains largely an art.

The assumption of rationality may be closer to reality when the players are regulars who play the game quite often. Then they benefit from having experienced the different possible outcomes. They understand how the strategic choices of various players lead to the outcomes and how well or badly they themselves fare. Then we as analysts of the game can hope that their choices, even if not made with full and conscious computations, closely approximate the results of such computations. We can think of the players as implicitly choosing the optimal strategy or behaving as if they were perfect calculators. We will offer some experimental evidence in Chapter 5 that the experience of playing the game generates more rational behavior.

The advantage of making a complete calculation of your best strategy, taking into account the corresponding calculations of a similar strategically calculating rival, is that then you are not making mistakes that the rival can exploit. In many actual situations, you may have specific knowledge of the way in which the other players fall short of this standard of rationality, and you can exploit this in devising your own strategy. We will say something about such calculations, but very often this is a part of the "art" of game playing, not easily codifiable in rules to be followed. You must always beware of the danger that the others are merely pretending to have poor skills or strategy, losing small sums through bad play and hoping that you will then raise the stakes, when they can raise the level of their play and exploit your gullibility. When this risk is real, the safer advice to a player may be to assume that the rivals are perfect and rational calculators and to choose his own best response to them. In other words, one should play to the opponents' capabilities instead of their limitations.

D. Common Knowledge of Rules

We suppose that, at some level, the players have a common understanding of the rules of the game. In a *Peanuts* cartoon, Lucy thought that body checking was allowed in golf and decked Charlie Brown just as he was about to take his swing. We do not allow this.

The qualification "at some level" is important. We saw how the rules of the immediate game could be manipulated. But this merely admits that there is another game being played at a deeper level—namely, where the players choose the rules of the superficial game. Then the question is whether the rules of this deeper game are fixed. For example, in the legislative context, what are the rules of the agenda-setting game? They may be that the committee chairs have the power. Then how are the committees and their chairs elected? And so on. At some basic level, the rules are fixed by the constitution, by the technology of campaigning, or by general social norms of behavior. We ask that all players recognize the given rules of this basic game, and that is the focus of the analysis. Of course, that is an ideal; in practice, we may not be able to proceed to a deep enough level of analysis.

Strictly speaking, the rules of the game consist of (1) the list of players, (2) the strategies available to each player, (3) the payoffs of each player for all possible combinations of strategies pursued by all the players, and (4) the assumption that each player is a rational maximizer.

Game theory cannot properly analyze a situation where one player does not know whether another player is participating in the game, what the entire sets of actions available to the other players are from which they can choose, what their value systems are, or whether they are conscious maximizers of their own payoffs. But in actual strategic interactions, some of the biggest gains are to be made by taking advantage of the element of surprise and doing something that your rivals never thought you capable of. Several vivid examples can be found in historic military conflicts. For example, in 1967 Israel launched a preemptive attack that destroyed the Egyptian air force on the ground; in 1973 it was Egypt's turn to spring a surprise by launching a tank attack across the Suez Canal.

It would seem, then, that the strict definition of game theory leaves out a very important aspect of strategic behavior, but in fact matters are not that bad. The theory can be reformulated so that each player attaches some small probability to the situation where such dramatically different strategies are available to the other players. Of course, each player knows his own set of available strategies. Therefore the game becomes one of asymmetric information and can be handled by using the methods developed in Chapter 9.

The concept of common knowledge itself requires some explanation. For some fact or situation X to be common knowledge between two people, A and B, it is not enough for each of them separately to know X. Each should also know that the other knows X; otherwise, for example, A might think that B does not know X and might act under this misapprehension in the midst of a game. But then A should also know that B knows that A knows X, and the other way around, otherwise A might mistakenly try to exploit B's supposed ignorance of A's knowledge. Of course, it doesn't even stop there. A should know that B knows that A knows that B knows, and so on ad infinitum. Philosophers have a lot of fun exploring the fine points of this infinite regress and the intellectual paradoxes that it can generate. For us, the general notion that the players have a common understanding of the rules of their game will suffice.

E. Equilibrium

Finally, what happens when rational players' strategies interact? Our answer will generally be in the framework of **equilibrium.** This simply means that each player is using the strategy that is the best response to the strategies of the other players. We will develop game-theoretic concepts of equilibrium in Chapters 3 through 8 and then use them in subsequent chapters.

Equilibrium does not mean that things don't change; in sequential-move games the players' strategies are the complete plans of action and reaction, and the position evolves all the time as the successive moves are made and responded to. Nor does equilibrium mean that everything is for the best; the interaction of rational strategic choices by all players can lead to bad outcomes for all, as in the prisoners' dilemma. But we will generally find that the idea of equilibrium is a useful descriptive tool and organizing concept for our analysis. We will consider this idea in greater detail later, in connection with specific equilibrium concepts. We will also see how the concept of equilibrium can be augmented or modified to remove some of its flaws and to incorporate behavior that falls short of full calculating rationality.

Just as the rational behavior of individual players can be the result of experience in playing the game, the fitting of their choices into an overall equilibrium can come about after some plays that involve trial and error and nonequilibrium outcomes. We will look at this matter in Chapter 5.

Defining an equilibrium is not hard; actually finding an equilibrium in a particular game—that is, solving the game—can be a lot harder. Throughout this book we will solve many simple games in which there are two or three players, each of them having two or three strategies or one move each in turn. Many people believe this to be the limit of the reach of game theory and therefore believe that the theory is useless for the more complex games that take place in reality. That is not true.

Humans are severely limited in their speed of calculation and in their patience for performing long calculations. Therefore humans can easily solve only the simple games with two or three players and strategies. But computers are very good at speedy and lengthy calculations. Many games that are far beyond the power of human calculators are easy for computers. The level of complexity in many games in business and politics is already within the powers of computers. Even in games such as chess that are far too complex to solve completely, computers have reached a level of ability comparable to that of the best humans; we consider chess in more detail in Chapter 3.

Computer programs for solving quite complex games exist, and more are appearing rapidly. Mathematica and similar program packages contain routines for finding mixed-strategy equilibria in simultaneous-move games. Gambit, a

National Science Foundation project led by Professors Richard D. McKelvey of the California Institute of Technology and Andrew McLennan of the University of Minnesota, is producing a comprehensive set of routines for finding equilibria in sequential- and simultaneous-move games, in pure and mixed strategies, and with varying degrees of uncertainty and incomplete information. We will refer to this project again in several places in the next several chapters. The biggest advantage of the project is that its programs are open source and can easily be obtained from its Web site with the URL http://gambit.sourceforge.net.

Why then do we set up and solve several simple games in detail in this book? The reason is that understanding the concepts is an important prerequisite for making good use of the mechanical solutions that computers can deliver, and understanding comes from doing simple cases yourself. This is exactly how you learned and now use arithmetic. You came to understand the ideas of addition, subtraction, multiplication, and division by doing many simple problems mentally or using paper and pencil. With this grasp of basic concepts, you can now use calculators and computers to do far more complicated sums than you would ever have the time or patience to do manually. If you did not understand the concepts, you would make errors in using calculators; for example, you might solve $3 + 4 \times 5$ by grouping additions and multiplications incorrectly as $(3 + 4) \times 5 = 35$ instead of correctly as $3 + (4 \times 5) = 23$.

Thus the first step of understanding the concepts and tools is essential. Without it, you would never learn to set up correctly the games that you ask the computer to solve. You would not be able to inspect the solution with any feeling for whether it was reasonable and, if it was not, would not be able to go back to your original specification, improve it, and solve it again until the specification and the calculation correctly capture the strategic situation that you want to study. Therefore please pay serious attention to the simple examples that we solve and the drill exercises that we ask you to solve, especially in Chapters 3 through 8.

F. Dynamics and Evolutionary Games

The theory of games based on assumptions of rationality and equilibrium has proved very useful, but it would be a mistake to rely on it totally. When games are played by novices who do not have the necessary experience to perform the calculations to choose their best strategies, explicitly or implicitly, their choices, and therefore the outcome of the game, can differ significantly from the predictions of analysis based on the concept of equilibrium.

However, we should not abandon all notions of good choice; we should recognize the fact that even poor calculators are motivated to do better for their own sakes and will learn from experience and by observing others. We should allow for a dynamic process in which strategies that proved to be better in previous plays of the game are more likely to be chosen in later plays.

The **evolutionary** approach to games does just this. It is derived from the idea of evolution in biology. Any individual animal's genes strongly influence its behavior. Some behaviors succeed better in the prevailing environment, in the sense that the animals exhibiting those behaviors are more likely to reproduce successfully and pass their genes to their progeny. An evolutionary stable state, relative to a given environment, is the ultimate outcome of this process over several generations.

The analogy in games would be to suppose that strategies are not chosen by conscious rational maximizers, but instead that each player comes to the game with a particular strategy "hardwired" or "programmed" in. The players then confront other players who may be programmed to apply the same or different strategies. The payoffs to all the players in such games are then obtained. The strategies that fare better—in the sense that the players programmed to play them get higher payoffs in the games—multiply faster, whereas the strategies that fare worse decline. In biology, the mechanism of this growth or decay is purely genetic transmission through reproduction. In the context of strategic games in business and society, the mechanism is much more likely to be social or cultural—observation and imitation, teaching and learning, greater availability of capital for the more successful ventures, and so on.

The object of study is the dynamics of this process. Does it converge to an evolutionary stable state? Does just one strategy prevail over all others in the end, or can a few strategies coexist? Interestingly, in many games the evolutionary stable limit is the same as the equilibrium that would result if the players were consciously rational calculators. Therefore the evolutionary approach gives us a backdoor justification for equilibrium analysis.

The concept of evolutionary games has thus imported biological ideas into game theory; there has been an influence in the opposite direction, too. Biologists have recognized that significant parts of animal behavior consist of strategic interactions with other animals. Members of a given species compete with one another for space or mates; members of different species relate to one another as predators and prey along a food chain. The payoff in such games in turn contributes to reproductive success and therefore to biological evolution. Just as game theory has benefited by importing ideas from biological evolution for its analysis of choice and dynamics, biology has benefited by importing game-theoretic ideas of strategies and payoffs for its characterization of basic interactions between animals. We have a true instance of synergy or symbiosis. We provide an introduction to the study of evolutionary games in Chapter 13.

G. Observation and Experiment

All of Section 3 to this point has concerned how to think about games or how to analyze strategic interactions. This constitutes theory. This book will cover an extremely simple level of theory, developed through cases and illustrations instead of formal mathematics or theorems, but it will be theory just the same. All theory should relate to reality in two ways. Reality should help structure the theory, and reality should provide a check on the results of the theory.

We can find out the reality of strategic interactions in two ways: (1) by observing them as they occur naturally and (2) by conducting special experiments that help us pin down the effects of particular conditions. Both methods have been used, and we will mention several examples of each in the proper contexts.

Many people have studied strategic interactions—the participants' behavior and the outcomes—under experimental conditions, in classrooms among "captive" players, or in special laboratories with volunteers. Auctions, bargaining, prisoners' dilemmas, and several other games have been studied in this way. The results are a mixture. Some conclusions of the theoretical analysis are borne out; for example, in games of buying and selling, the participants generally settle quickly on the economic equilibrium. In other contexts, the outcomes differ significantly from the theoretical predictions; for example, prisoners' dilemmas and bargaining games show more cooperation than theory based on the assumption of selfish, maximizing behavior would lead us to expect, whereas auctions show some gross overbidding.

At several points in the chapters that follow, we will review the knowledge that has been gained by observation and experiments, discuss how it relates to the theory, and consider what reinterpretations, extensions, and modifications of the theory have been made or should be made in the light of this knowledge.

4 THE USES OF GAME THEORY

We began Chapter 1 by saying that games of strategy are everywhere—in your personal and working life; in the functioning of the economy, society, and polity around you; in sports and other serious pursuits; in war; and in peace. This should be motivation enough to study such games systematically, and that is what game theory is about, but your study can be better directed if you have a clearer idea of just how you can put game theory to use. We suggest a threefold method.

The first use is in *explanation*. Many events and outcomes prompt us to ask: Why did that happen? When the situation requires the interaction of decision makers with different aims, game theory often supplies the key to understanding the situation. For example, cutthroat competition in business is the result of the rivals being trapped in a prisoners' dilemma. At several points in the book we will mention actual cases where game theory helps us to understand how and why the events unfolded as they did. This includes the detailed case study of the Cuban missile crisis from the perspective of game theory.

The other two uses evolve naturally from the first. The second is in *prediction*. When looking ahead to situations where multiple decision makers will interact strategically, we can use game theory to foresee what actions they will take and what outcomes will result. Of course, prediction for a particular context depends on its details, but we will prepare you to use prediction by analyzing several broad classes of games that arise in many applications.

The third use is in *advice* or *prescription*: we can act in the service of one participant in the future interaction and tell him which strategies are likely to yield good results and which ones are likely to lead to disaster. Once again such work is context specific, and we equip you with several general principles and techniques and show you how to apply them to some general types of contexts. For example, in Chapters 7 and 8 we will show how to mix moves, in Chapter 10 we will examine how to make your commitments, threats, and promises credible, and in Chapter 11 we will examine alternative ways of overcoming prisoners' dilemmas.

The theory is far from perfect in performing any of the three functions. To explain an outcome, one must first have a correct understanding of the motives and behavior of the participants. As we saw earlier, most of game theory takes a specific approach to these matters—namely, the framework of rational choice of individual players and the equilibrium of their interaction. Actual players and interactions in a game might not conform to this framework. But the proof of the pudding is in the eating. Game-theoretic analysis has greatly improved our understanding of many phenomena, as reading this book should convince you. The theory continues to evolve and improve as the result of ongoing research. This book will equip you with the basics so that you can more easily learn and profit from the new advances as they appear.

When explaining a past event, we can often use historical records to get a good idea of the motives and the behavior of the players in the game. When attempting prediction or advice, there is the additional problem of determining what motives will drive the players' actions, what informational or other limitations they will face, and sometimes even who the players will be. Most important, if game-theoretic analysis assumes that the other player is a rational maximizer of his own objectives when in fact he is unable to do the calculations or is a clueless person acting at random, the advice based on that analysis may prove wrong. This risk is reduced as more and more players recognize the importance of strategic interaction and think through their strategic choices or get expert advice on the matter, but some risk remains. Even then, the systematic thinking made possible by the framework of game theory helps keep the errors down to this irreducible minimum, by eliminating the errors that arise from faulty logical thinking about the strategic interaction. Also, game theory can take into account many kinds of uncertainty and incomplete information, including that about the strategic possibilities and rationality of the opponent. We will consider a few examples in the chapters to come.

5 THE STRUCTURE OF THE CHAPTERS TO FOLLOW

In this chapter we introduced several considerations that arise in almost every game in reality. To understand or predict the outcome of any game, we must know in greater detail all of these ideas. We also introduced some basic concepts that will prove useful in such analysis. However, trying to cope with all of the concepts at once merely leads to confusion and a failure to grasp any of them. Therefore we will build up the theory one concept at a time. We will develop the appropriate technique for analyzing that concept and illustrate it.

In the first group of chapters, from Chapters 3 to 8, we will construct and illustrate the most important of these concepts and techniques. We will examine purely sequential-move games in Chapter 3 and introduce the techniques—game trees and rollback reasoning—that are used to analyze and solve such games. In Chapters 4 and 5, we will turn to games with simultaneous moves and develop for them another set of concepts—payoff tables, dominance, and Nash equilibrium. Both chapters will focus on games where players use pure strategies; in Chapter 4, we will restrict players to a finite set of pure strategies and, in Chapter 5, we will allow strategies that are continuous variables. Chapter 5 will also examine some mixed empirical evidence and conceptual criticisms and counterarguments on Nash equilibrium, and a prominent alternative to Nash equilibrium—namely, rationalizability. In Chapter 6, we will show how games that have some sequential moves and some simultaneous moves can be studied by combining the techniques developed in Chapters 3 through 5. In Chapters 7 and 8, we will turn to simultaneous-move games that require the use of randomization or mixed strategies. In Chapter 7, we introduce the basic ideas about mixing in two-by-two games, develop the simplest techniques for finding mixed-strategy Nash equilibria, and consider empirical evidence on mixing. Chapter 8 will then develop a little general theory of mixed strategies.

The ideas and techniques developed in Chapters 3 through 8 are the most basic ones: (1) correct forward-looking reasoning for sequential-move games and (2) equilibrium strategies—pure and mixed—for simultaneous-move games. Equipped with these concepts and tools, we can apply them to study some broad classes of games and strategies in Chapters 9 through 13.

Chapter 9 studies what happens in games when players are subject to uncertainty or when they have asymmetric information. We will examine strategies for coping with risk and even for using risk strategically. We will also study the important strategies of signaling and screening that are used for manipulating and eliciting information. We develop the appropriate generalization of Nash equilibrium in the context of uncertainty, namely Bayesian Nash equilibrium, and show the different kinds of equilibria that can arise. In Chapter 10, we will continue to examine the role of player manipulation in games as we consider strategies that players use to manipulate the rules of a game, by seizing a first-mover advantage and making a strategic move. Such moves are of three kinds—commitments, threats, and promises. In each case, credibility is essential to the success of the move, and we will outline some ways of making such moves credible.

In Chapter 11, we will move on to study the best-known game of them all—the prisoners' dilemma. We will study whether and how cooperation can be sustained, most importantly in a repeated or ongoing relationship. Then, in Chapter 12, we will turn to situations where large populations, rather than pairs or small groups of players, interact strategically, games that concern problems of collective action. Each person's actions have an effect—in some instances beneficial, in others, harmful—on the others. The outcomes are generally not the best from the aggregate perspective of the society as a whole. We will clarify the nature of these outcomes and describe some simple policies that can lead to better outcomes.

All these theories and applications are based on the supposition that the players in a game fully understand the nature of the game and deploy calculated strategies that best serve their objectives in the game. Such rationally optimal behavior is sometimes too demanding of information and calculating power to be believable as a good description of how people really act. Therefore Chapter 13 will look at games from a very different perspective. Here, the players are not calculating and do not pursue optimal strategies. Instead, each player is tied, as if genetically preordained, to a particular strategy. The population is diverse, and different players have different predetermined strategies. When players from such a population meet and act out their strategies, which strategies perform better? And if the more successful strategies proliferate better in the population, whether through inheritance or imitation, then what will the eventual structure of the population look like? It turns out that such evolutionary dynamics often favor exactly those strategies that would be used by rational optimizing players. Thus our study of evolutionary games lends useful indirect support to the theories of optimal strategic choice and equilibrium that we will have studied in the previous chapters.

In the final group, Chapters 14 through 19, we will take up specific applications to situations of strategic interactions. Here, we will use as needed the ideas and methods from all the earlier chapters. Chapter 14 uses the methods developed in Chapter 9 to analyze the strategies that people and firms have to use when dealing with others who have some private information. We will illustrate the screening mechanisms that are used for eliciting information, for example, the multiple fares with different restrictions that airlines use for separating the business travelers who are willing to pay more from the tourists who are more price sensitive. We will also develop the methods for designing incentive payments to elicit effort from workers when direct monitoring is difficult or too costly. Chapter 15 then applies the ideas from Chapter 10 to examine a particularly interesting dynamic version of a threat, known as the strategy of brinkmanship. We will elucidate its nature and apply the idea to study the Cuban missile crisis of 1962. Chapter 16 is about voting in committees and elections. We will look at the variety of voting rules available and some paradoxical results that can arise. In addition, we will address the potential for strategic behavior not only by voters but also by candidates in a variety of election types.

Chapters 17 through 19 will look at mechanisms for the allocation of valuable economic resources: Chapter 17 will treat auctions, Chapter 18 will consider bargaining processes, and Chapter 19 will look at markets. In our discussion of auctions, we will emphasize the roles of information and attitudes toward risk in the formulation of optimal strategies for both bidders and sellers. We will also take the opportunity to apply the theory to the newest type of auctions, those that take place online. Chapter 18 will present bargaining in both cooperative and noncooperative settings. Finally, Chapter 19 will consider games of market exchange, building on some of the concepts used in bargaining theory and including some theory of the core.

All of these chapters together provide a lot of material; how might readers or teachers with more specialized interests choose from it? Chapters 3 through 7 constitute the core theoretical ideas that are needed throughout the rest of the book. Chapters 10 and 11 are likewise important for the general classes of games and strategies considered therein. Beyond that, there is a lot from which to pick and choose. Section 1 of Chapter 5 and all of Chapter 8 consider some more advanced topics and go somewhat deeper into theory and mathematics. These chapters will appeal to those with more scientific and quantitative backgrounds and interests, but those who come from the social sciences or humanities and have less quantitative background can omit them without loss of continuity. Chapter 9 deals with an important topic in that most games in practice have incomplete and asymmetric information, and the players' attempts to manipulate information is a critical aspect of many strategic interactions. However, the concepts and techniques for analyzing information games are inherently somewhat more complex. Therefore some readers and teachers may choose to study just the examples that convey the basic ideas of signaling and screening and leave out the rest. We have placed this chapter early in Part Three, however, in view of the importance of the subject. Chapters 10 and 11 are key to understanding many phenomena in the real world, and most teachers will want to include them in their courses, but Section 5 of Chapter 11 is mathematically a little more advanced and can be omitted. Chapters 12 and 13 both look at games with large numbers of players. In Chapter 12, the focus is on social interactions; in Chapter 13, the

focus is on evolutionary biology. The topics in Chapter 13 will be of greatest interest to those with interests in biology, but similar themes are emerging in the social sciences, and students from that background should aim to get the gist of the ideas even if they skip the details. Chapter 14 is most important for students of business and organization theories. Chapters 15 and 16 present topics from political science—international diplomacy and elections, respectively—and Chapters 17 through 19 cover topics from economics—auctions, bargaining, and markets. Those teaching courses with more specialized audiences may choose a subset from Chapters 12 through 19, and indeed expand on the ideas considered therein.

Whether you come from mathematics, biology, economics, politics, other sciences, or from history or sociology, the theory and examples of strategic games will stimulate and challenge your intellect. We urge you to enjoy the subject even as you are studying or teaching it.

SUMMARY

Strategic *games* situations are distinguished from individual decision-making situations by the presence of significant interactions among the players. Games can be classified according to a variety of categories including the timing of play, the common or conflicting interests of players, the number of times an interaction occurs, the amount of information available to the players, the type of rules, and the feasibility of coordinated action.

Learning the terminology for a game's structure is crucial for analysis. Players have *strategies* that lead to different *outcomes* with different associated *payoffs*. Payoffs incorporate everything that is important to a player about a game and are calculated by using probabilistic averages or *expectations* if outcomes are random or include some risk. *Rationality*, or consistent behavior, is assumed of all players, who must also be aware of all of the relevant rules of conduct. *Equilibrium* arises when all players use strategies that are best responses to others' strategies; some classes of games allow learning from experience and the study of dynamic movements toward equilibrium. The study of behavior in actual game situations provides additional information about the performance of the theory.

Game theory may be used for explanation, prediction, or prescription in various circumstances. Although not perfect in any of these roles, the theory continues to evolve; the importance of strategic interaction and strategic thinking has also become more widely understood and accepted.

KEY TERMS[3]

SOLVED EXERCISES[4]

S1. Determine which of the following situations describe games and which describe decisions. In each case, indicate what specific features of the situation caused you to classify it as you did.

(a) A group of grocery shoppers in the dairy section, with each shopper choosing a flavor of yogurt to purchase

(b) A pair of teenage girls choosing dresses for their prom

(c) A college student considering what type of postgraduate education to pursue

(d) The *New York Times* and the *Wall Street Journal* choosing the prices for their online subscriptions this year

(e) A presidential candidate picking a running mate

S2. Consider the strategic games described below. In each case, state how you would classify the game according to the six dimensions outlined in the text. (i) Are moves sequential or simultaneous? (ii) Is the game zero-sum or not? (iii) Is the game repeated? (iv) Is there imperfect information, and if so, is there incomplete (asymmetric) information? (v) Are the rules fixed or not? (vi) Are cooperative agreements possible or not? If you do not have enough information to classify a game in a particular dimension, explain why not.

(a) *Rock-Paper-Scissors:* On the count of three, each player makes the shape of one of the three items with his hand. Rock beats Scissors, Scissors beats Paper, and Paper beats Rock.

[3]The number in parentheses after each key term is the page on which that term is defined or discussed.

[4]**Note to Students:** The solutions to the **Solved Exercises** are found on the following Web site, which is free and open to all: **wwnorton.com/books/games_of_strategy**.

(b) *Roll-call voting:* Voters cast their votes orally as their names are called. The choice with the most votes wins.

(c) *Sealed-bid auction:* Bidders on a bottle of wine seal their bids in envelopes. The highest bidder wins the item and pays the amount of his bid.

S3. "A game player would never prefer an outcome in which every player gets a little profit to an outcome in which he gets all the available profit." Is this statement true or false? Explain why in one or two sentences.

S4. You and a rival are engaged in a game in which there are three possible outcomes: you win, your rival wins (you lose), or the two of you tie. You get a payoff of 50 if you win, a payoff of 20 if you tie, and a payoff of zero if you lose. What is your expected payoff in each of the following situations?

(a) There is a 50% chance that the game ends in a tie, but only a 10% chance that you win. (There is thus a 40% chance that you lose.)

(b) There is a 50-50 chance that you win or lose. There are no ties.

(c) There is an 80% chance that you lose, a 10% chance that you win, and a 10% chance that you tie.

S5. Explain the difference between game theory's use as a predictive tool and its use as a prescriptive tool. In what types of real-world settings might these two uses be most important?

UNSOLVED EXERCISES

U1. Determine which of the following situations describe games and which describe decisions. In each case, indicate what specific features of the situation caused you to classify it as you did.

(a) A party nominee for president of the United States must choose whether to use private financing or public financing for her campaign.

(b) Frugal Fred receives a $20 gift card for downloadable music and must choose whether to purchase individual songs or whole albums.

(c) Beautiful Belle receives 100 replies to her online dating profile and must choose whether to reply to each of them.

(d) NBC chooses how to distribute its television shows online this season. They consider Amazon.com, iTunes, and/or NBC.com. The fee they might pay to Amazon or to iTunes is open to negotiation.

(e) China chooses a level of tariffs to apply to American imports.

U2. Consider the strategic games described below. In each case, state how you would classify the game according to the six dimensions outlined in the text. (i) Are moves sequential or simultaneous? (ii) Is the game zero-sum or not? (iii) Is the game repeated? (iv) Is there imperfect information, and if so, is there incomplete (asymmetric) information? (v) Are the rules fixed or not? (vi) Are cooperative agreements possible or not? If you do not have enough information to classify a game in a particular dimension, explain why not.

(a) Garry and Ross are sales representatives for the same company. Their manager informs them that of the two of them, whoever sells more this year wins a Cadillac.

(b) On the game show *The Price is Right*, four contestants are asked to guess the price of a television set. Play starts with the leftmost player, and each player's guess must be different from the guesses of the previous players. The person who comes closest to the real price, without going over it, wins the television set.

(c) Six thousand players each pay $10,000 to enter the World Series of Poker. Each starts the tournament with $10,000 in chips, and they play No-Limit Texas Hold 'Em (a type of poker) until someone wins all the chips. The top six hundred players each receive prize money according to the order of finish, with the winner receiving more than $8,000,000.

(d) Passengers on Desert Airlines are not assigned seats; passengers choose seats once they board. The airline assigns the order of boarding according to the time the passenger checks in, either on the Web site up to 24 hours before takeoff, or in person at the airport.

U3. "Any gain by the winner must harm the loser." Is this statement true or false? Explain your reasoning in one or two sentences.

U4. Alice, Bob, and Confucius are bored during recess, so they decide to play a new game. Each of them puts a dollar in the pot, and each tosses a quarter. Alice wins if the coins land all heads or all tails. Bob wins if two heads and one tail land, and Confucius wins if one head and two tails land. The quarters are fair, and the winner receives a net payment of $2 ($3 − $1 = $2), and the losers lose their $1.

(a) What is the probability that Alice will win and the probability that she will lose?

(b) What is Alice's expected payoff?

(c) What is the probability that Confucius will win and the probability that he will lose?

(d) What is Confucius' expected payoff?

(e) Is this a zero-sum game? Please explain your answer.

U5. "When one player surprises another, this indicates that the players did not have common knowledge of the rules." Give an example that illustrates this statement, and give a counterexample that shows that the statement is not always true.

PART TWO

Concepts and Techniques

3

Games with Sequential Moves

Sequential-move games entail strategic situations in which there is a strict order of play. Players take turns making their moves, and they know what players who have gone before them have done. To play well in such a game, participants must use a particular type of interactive thinking. Each player must consider: If I make this move, how will my opponent respond? Whenever actions are taken, players need to think about how their current actions will influence future actions, both for their rivals and for themselves. Players thus decide their current moves on the basis of calculations of future consequences.

Most actual games combine aspects of both sequential- and simultaneous-move situations. But the concepts and methods of analysis are more easily understood if they are first developed separately for the two pure cases. Therefore in this chapter we study purely sequential games. Chapters 4 and 5 deal with purely simultaneous games, and Chapter 6 and parts of Chapters 7 and 8 show how to combine the two types of analysis in more realistic mixed situations. The analysis presented here can be used whenever a game includes sequential decision making. Analysis of sequential games also provides information about when it is to a player's advantage to move first and when it is better to move second. Players can then devise ways, called *strategic moves*, to manipulate the order of play to their advantage. The analysis of such moves is the focus of Chapter 10.

1 GAME TREES

We begin by developing a graphical technique for displaying and analyzing sequential-move games, called a **game tree.** This tree is referred to as the **extensive form** of a game. It shows all the component parts of the game that we introduced in Chapter 2: players, actions, and payoffs.

You have probably come across **decision trees** in other contexts. Such trees show all the successive decision points, or nodes, for a single decision maker in a neutral environment. Decision trees also include branches corresponding to the available choices emerging from each node. Game trees are just joint decision trees for all of the players in a game. The trees illustrate all of the possible actions that can be taken by all of the players and indicate all of the possible outcomes of the game.

A. Nodes, Branches, and Paths of Play

Figure 3.1 shows the tree for a particular sequential game. We do not supply a story for this game, because we want to omit circumstantial details and to help you focus on general concepts. Our game has four players: Ann, Bob, Chris, and Deb. The rules of the game give the first move to Ann; this is shown at the leftmost point, or **node,** which is called the **initial node** or **root** of the game tree. At this node, which may also be called an **action** or **decision node,** Ann has two choices available to her. Ann's possible choices are labeled "Stop" and "Go" (remember that these labels are abstract and have no necessary significance) and are shown as **branches** emerging from the initial node.

If Ann chooses "Stop," then it will be Bob's turn to move. At his action node, he has three available choices labeled 1, 2, and 3. If Ann chooses "Go," then Chris gets the next move, with choices "Risky" and "Safe." Other nodes and branches follow successively and, rather than list them all in words, we draw your attention to a few prominent features.

If Ann chooses "Stop" and then Bob chooses 1, Ann gets another turn, with new choices, "Up" and "Down." It is quite common in actual sequential-move games for a player to get to move several times and to have her available moves differ at different turns. In chess, for example, two players make alternate moves; each move changes the board and therefore the available moves at subsequent turns.

B. Uncertainty and "Nature's Moves"

If Ann chooses "Go" and then Chris chooses "Risky," something happens at random—a fair coin is tossed and the outcome of the game is determined by whether that coin comes up "heads" or "tails." This aspect of the game is an

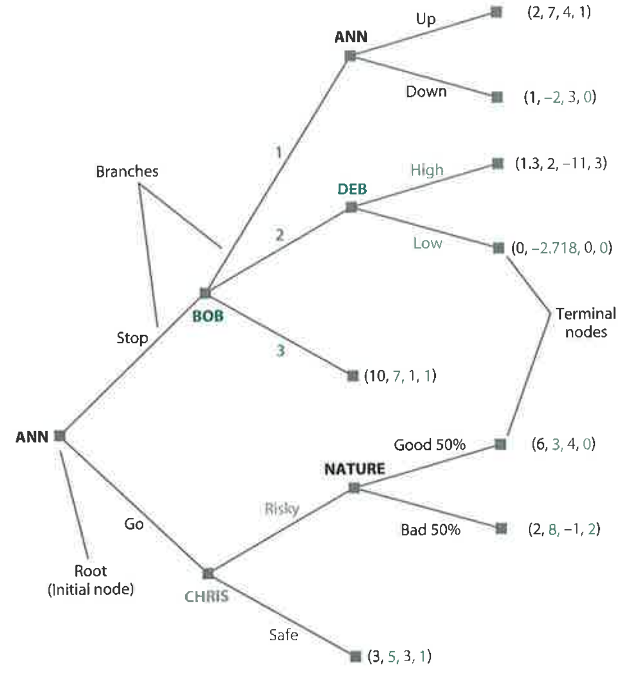

FIGURE 3.1 An Illustrative Game Tree

example of external uncertainty and is handled in the tree by introducing an outside player called "Nature." Control over the random event is ceded to the player known as Nature, who chooses, as it were, one of two branches, each with 50% probability. The probabilities here are fixed by the type of random event, a coin toss, but could vary in other circumstances; for example, with the throw of a die, Nature could specify six possible outcomes, each with $16\frac{2}{3}\%$ probability. Use of the player Nature allows us to introduce external uncertainty in a game and gives us a mechanism to allow things to happen that are outside the control of any of the actual players.

You can trace a number of different paths through the game tree by following successive branches. In Figure 3.1, each path leads you to an end point of the game after a finite number of moves. An end point is not a necessary feature of all games; some may in principle go on forever. But most applications that we will consider are finite games.

C. Outcomes and Payoffs

At the last node along each path, called a **terminal node,** no player has another move. (Note that terminal nodes are thus distinguished from *action* nodes.) Instead, we show the outcome of that particular sequence of actions, as measured by the payoffs for the players. For our four players, we list the payoffs in order (Ann, Bob, Chris, Deb). It is important to specify which payoff belongs to which player. The usual convention is to list payoffs in the order in which the players make the moves. But this method may sometimes be ambiguous; in our example, it is not clear whether Bob or Chris should be said to have the second move. Thus we have used alphabetical order. Further, we have color-coded everything so that Ann's name, choices, and payoffs are all in black; Bob's in dark green; Chris's in grey; and Deb's in light green. When drawing trees for any games that you analyze, you can choose any specific convention you like, but you should state and explain it clearly for the reader.

The payoffs are numerical, and generally for each player a higher number means a better outcome. Thus, for Ann, the outcome of the bottommost path (payoff 3) is better than that of the topmost path (payoff 2) in Figure 3.1. But there is no necessary comparability across players. Thus there is no necessary sense in which, at the end of the topmost path, Bob (payoff 7) does better than Ann (payoff 2). Sometimes, if payoffs are dollar amounts, for example, such interpersonal comparisons may be meaningful.

Players use information about payoffs when deciding among the various actions available to them. The inclusion of a random event (a choice made by Nature) means that players need to determine what they get on average when Nature moves. For example, if Ann chooses "Go" at the game's first move, Chris may then choose "Risky," giving rise to the coin toss and Nature's "choice" of "Good" or "Bad." In this situation, Ann could anticipate a payoff of 6 half the time and a payoff of 2 half the time, or a statistical average or *expected payoff* of $4 = (0.5 \times 6) + (0.5 \times 2)$.

D. Strategies

Finally, we use the tree in Figure 3.1 to explain the concept of a strategy. A single action taken by a player at a node is called a **move.** But players can, do, and should make plans for the succession of moves that they expect to make in all of the various eventualities that might arise in the course of a game. Such a plan of action is called a **strategy.**

In this tree, Bob, Chris, and Deb each get to move at most once; Chris, for example, gets a move only if Ann chooses "Go" on her first move. For them, there is no distinction between a move and a strategy. We can qualify the move by specifying the contingency in which it gets made; thus, a strategy for Bob might be,

"Choose 1 if Ann has chosen Stop." But Ann has two opportunities to move, so her strategy needs a fuller specification. One strategy for her is, "Choose Stop, and then if Bob chooses 1, choose Down."

In more complex games such as chess, where there are long sequences of moves with many choices available at each, descriptions of strategies get very complicated; we consider this aspect in more detail later in this chapter. But the general principle for constructing strategies is simple, except for one peculiarity. If Ann chooses "Go" on her first move, she never gets to make a second move. Should a strategy in which she chooses "Go" also specify what she would do in the hypothetical case in which she somehow found herself at the node of her second move? Your first instinct may be to say *no,* but formal game theory says *yes,* and for two reasons.

First, Ann's choice of "Go" at the first move may be influenced by her consideration of what she would have to do at her second move if she were to choose "Stop" originally instead. For example, if she chooses "Stop," Bob may then choose 1; then Ann gets a second move and her best choice would be "Up," giving her a payoff of 2. If she chooses "Go" on her first move, Chris would choose "Safe" (because his payoff of 3 from "Safe" is better than his expected payoff of 1.5 from "Risky"), and that outcome would yield Ann a payoff of 3. To make this thought process clearer, we state Ann's strategy as, "Choose Go at the first move, and choose Up if the next move arises."

The second reason for this seemingly pedantic specification of strategies has to do with the stability of equilibrium. When considering stability, we ask what would happen if players' choices were subjected to small disturbances. One such disturbance is that players make small mistakes. If choices are made by pressing a key, for example, Ann may intend to press the "Go" key, but there is a small probability that her hand may tremble and she may press the "Stop" key instead. In such a setting, it is important to specify how Ann will follow up when she discovers her error because Bob chooses 1 and it is Ann's turn to move again. More advanced levels of game theory require such stability analyses, and we want to prepare you for that by insisting on your specifying strategies as such complete plans of action right from the beginning.

E. Tree Construction

Now we sum up the general concepts illustrated by the tree of Figure 3.1. Game trees consist of nodes and branches. Nodes are connected to one another by the branches and come in two types. The first node type is called a decision node. Each decision node is associated with the player who chooses an action at that node; every tree has one decision node that is the game's initial node, the starting point of the game. The second type of node is called a terminal node. Each terminal node has associated with it a set of outcomes for the players taking part in the game; these outcomes are the payoffs received by each player if the game has followed the branches that lead to this particular terminal node.

The branches of a game tree represent the possible actions that can be taken from any decision node. Each branch leads from a decision node on the tree either to another decision node, generally for a different player, or to a terminal node. The tree must account for all of the possible choices that could be made by a player at each node; so some game trees include branches associated with the choice "Do nothing." There must be at least one branch leading from each decision node, but there is no maximum. Every decision node can have only one branch leading to it, however.

Game trees are often drawn from left to right across a page. However, game trees can be drawn in any orientation that best suits the game at hand: bottom up, sideways, top down, or even radially outward from a center. The tree is a metaphor, and the important feature is the idea of successive branching, as decisions are made at the tree nodes.

2 SOLVING GAMES BY USING TREES

We illustrate the use of trees in finding equilibrium outcomes of sequential-move games in a very simple context that many of you have probably confronted—whether to smoke. This situation and many other similar one-player strategic situations can be described as games if we recognize that future choices are actually made by a different player. That player is one's future self who will be subject to different influences and will have different views about the ideal outcome of the game.

Take, for example, a teenager named Carmen who is deciding whether to smoke. First, she has to decide whether to try smoking at all. If she does try it, she has the further decision of whether to continue. We illustrate this example as a simple decision in the tree of Figure 3.2.

The nodes and the branches are labeled with Carmen's available choices, but we need to explain the payoffs. Choose the outcome of never smoking at all as the standard of reference, and call its payoff 0. There is no special significance to the number zero in this context; all that matters for comparing outcomes, and thus for Carmen's decision, is whether this payoff is bigger or smaller than the others. Suppose Carmen best likes the outcome in which she tries smoking for a while but does not continue. The reason may be that she just likes to have experienced many things first-hand or so that she can more convincingly be able to say "I have been there and know it to be a bad situation" when she tries in the future to dissuade her children from smoking. Give this outcome the payoff 1. The outcome in which she tries smoking and then continues is the worst. Leav-

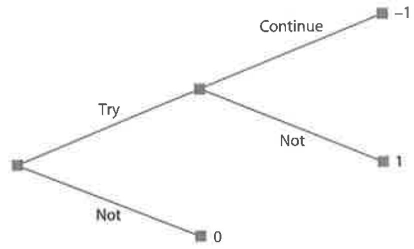

FIGURE 3.2 The Smoking Decision

ing aside the long-term health hazards, there are immediate problems—her hair and clothes will smell bad, and her friends will avoid her. Give this outcome the payoff -1. Carmen's best choice then seems clear—she should try smoking but she should not continue.

However, this analysis ignores the problem of addiction. Once Carmen has tried smoking for a while, she becomes a different person with different tastes, as well as different payoffs. The decision of whether to continue will be made not by "Today's Carmen" with today's assessment of outcomes as shown in Figure 3.2, but by a different "Future Carmen" with a different ranking of the alternatives then available. When she makes her choice today, she has to look ahead to this consequence and factor it into her current decision, which she should make on the basis of her current preferences. In other words, the choice problem concerning smoking is not really a decision in the sense explained in Chapter 2—a choice made by a single person in a neutral environment—but a game in the technical sense also explained in Chapter 2, where the other player is Carmen's future self with her own distinct preferences. When Today's Carmen makes her decision, she has to play against her future self.

We convert the decision tree of Figure 3.2 into a game tree in Figure 3.3, by distinguishing between the two players who make the choices at the two nodes. At the initial node, "Today's Carmen" decides whether to try smoking. If her decision is to try, then the addicted "Future Carmen" comes into being and chooses whether to continue. We show the healthy, non-polluting Today's Carmen, her actions, and her payoffs in green, and the addicted Future Carmen, her actions, and her payoffs in black, the color that her lungs have become. The payoffs of Today's Carmen are as before. But Future Carmen will enjoy continuation and will suffer terrible withdrawal symptoms if she does not continue. Let Future Carmen's payoff from "continue" be $+1$ and that from "not" be -1.

Given the preferences of the addicted Future Carmen, she will choose "continue" at her decision node. Today's Carmen should look ahead to this prospect and fold it into her current decision, recognizing that the choice to

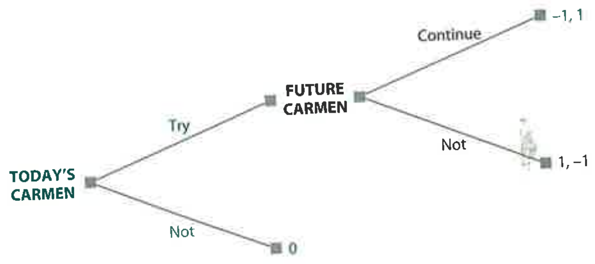

FIGURE 3.3 The Smoking Game

try smoking will inevitably lead to continuation. Even though Today's Carmen does not want continuation to happen given her preferences today, she will not be able to implement her currently preferred choice at the future time, because a different Carmen with different preferences will make that choice. So Today's Carmen should foresee that the choice "Try" will lead to "Continue" and get her the payoff -1 as judged by her today, whereas the choice "Don't Try" will get her the payoff 0. So she should choose the latter.

This argument is shown more formally and with greater visual effect in Figure 3.4. In Figure 3.4a, we cut off, or **prune,** the branch "Not" emerging from the second node. This pruning corresponds to the fact that Future Carmen, who makes the choice at that node, will not choose the action associated with that branch, given her preferences as shown in black.

The tree that remains has two branches emerging from the first node where Today's Carmen makes her choice; each of these branches now leads directly to a terminal node. The pruning allows Today's Carmen to forecast completely the eventual consequence of each of her choices. "Try" will be followed by "Continue" and yield a payoff -1, as measured in the preferences of Today's Carmen, while "Not" will yield 0. Carmen's choice today should then be "Not" rather than "Try." Therefore we can prune the "Try" branch emerging from the first node (along with its foreseeable continuation). This pruning is done in Figure 3.4b. The tree shown there is now "fully pruned," leaving only one branch emerging from the initial node and leading to a terminal node. Following the only remaining path through the tree shows what will happen in the game when all players make their best choices with correct forecasting of all future consequences.

In pruning the tree in Figure 3.4, we crossed out the branches not chosen. Another equivalent but alternative way of showing player choices is to "highlight" the branches that *are* chosen. To do so, you can place check marks or arrowheads on these branches or show them as thicker lines. Any one method will do; Figure 3.5 shows them all. You can choose whether to prune or to highlight,

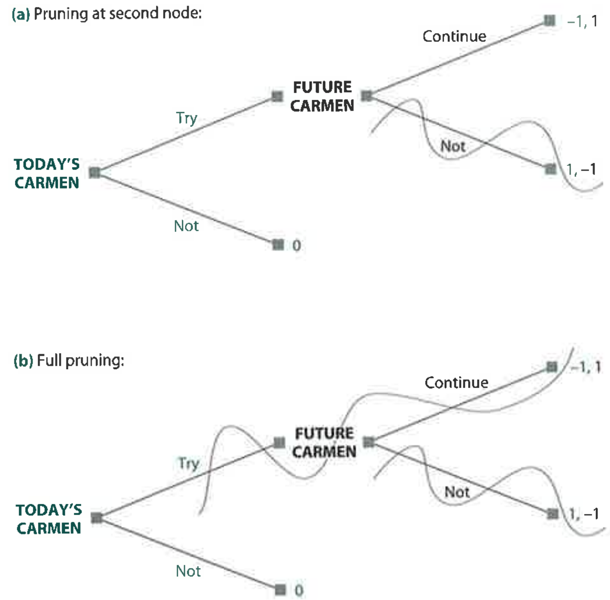

FIGURE 3.4 Pruning the Tree of the Smoking Game

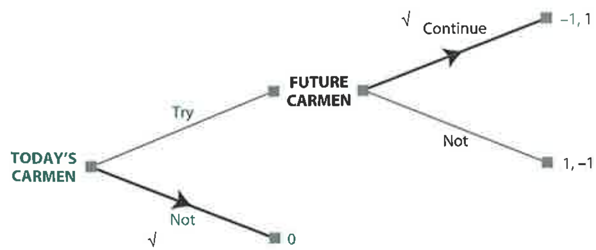

FIGURE 3.5 Showing Branch Selection on the Tree of the Smoking Game

but the latter, especially in its arrowhead form, has some advantages. First, it produces a cleaner picture. Second, the mess of the pruning picture sometimes does not clearly show the order in which various branches were cut. For example, in Figure 3.4b, a reader may get confused and incorrectly think that the "Continue" branch at the second node was cut first and that the "Try" branch at the first node followed by the "Not continue" branch at the second node were cut next. Finally, and most important, the arrowheads show the outcome of the sequence of optimal choices most visibly as a continuous link of arrows from the initial node to a terminal node. Therefore, in subsequent diagrams of this type, we generally use arrows instead of pruning. When you draw game trees, you should practice showing both methods for a while; when you are comfortable with trees, you can choose either to suit your taste.

No matter how you display your thinking in a game tree, the logic of the analysis is the same and important. You must start your analysis by considering those action nodes that lead directly to terminal nodes. The optimal choices for a player moving at such a node can be found immediately by comparing her payoffs at the relevant terminal nodes. With the use of these end-of-game choices to forecast consequences of earlier actions, the choices at nodes just preceding the final decision nodes can be determined. Then the same can be done for the nodes before them, and so on. By working backward along the tree in this way, you can solve the whole game.

This method of looking ahead and reasoning back to determine behavior in sequential-move games is known as **rollback.** As the name suggests, using rollback requires starting to think about what will happen at all the terminal nodes and literally "rolling back" through the tree to the initial node as you do your analysis. Because this reasoning requires working backward one step at a time, the method is also called **backward induction.** We use the term rollback because it is simpler and becoming more widely used, but other sources on game theory will use the older term backward induction. Just remember that the two are equivalent.

When all players choose their optimal strategies found by doing rollback analysis, we call this set of strategies the **rollback equilibrium** of the game; the outcome that arises from playing these strategies is the *rollback equilibrium outcome.* Game theory predicts this outcome as the equilibrium of a sequential game when all players are rational calculators in pursuit of their respective best payoffs. Later in this chapter, we address how well this prediction is borne out in practice. For now, you should know that all finite sequential-move games presented in this book have at least one rollback equilibrium. In fact, most have exactly one. Only in exceptional cases where a player gets equal payoffs from two or more different sets of moves, and is therefore indifferent between them, will games have more than one rollback equilibrium.

In the smoking game, the rollback equilibrium is where Today's Carmen chooses the strategy "Not" and Future Carmen chooses the strategy "Continue."

When Today's Carmen takes her optimal action, the addicted Future Carmen does not come into being at all and therefore gets no actual opportunity to move. But Future Carmen's shadowy presence and the strategy that she would choose if Today's Carmen chose "Try" and gave her an opportunity to move are important parts of the game. In fact, they are instrumental in determining the optimal move for Today's Carmen.

We introduced the ideas of the game tree and rollback analysis in a very simple example, where the solution was obvious from verbal argument. Now we proceed to use the ideas in successively more complex situations, where verbal analysis becomes harder to conduct and the visual analysis with the use of the tree becomes more important.

3 ADDING MORE PLAYERS

The techniques developed in Section 2 in the simplest setting of two players and two moves can be readily extended. The trees get more complex, with more branches, nodes, and levels, but the basic concepts and the method of rollback remain unchanged. In this section, we consider a game with three players, each of whom has two choices; with slight variations, this game reappears in many subsequent chapters.

The three players, Emily, Nina, and Talia, all live on the same small street. Each has been asked to contribute toward the creation of a flower garden at the intersection of their small street with the main highway. The ultimate size and splendor of the garden depends on how many of them contribute. Furthermore, although each player is happy to have the garden—and happier as its size and splendor increase—each is reluctant to contribute because of the cost that she must incur to do so.

Suppose that, if two or all three contribute, there will be sufficient resources for the initial planting and subsequent maintenance of the garden; it will then be quite attractive and pleasant. However, if one or none contribute, it will be too sparse and poorly maintained to be pleasant. From each player's perspective, there are thus four distinguishable outcomes:

- She does not contribute, both of the others do (pleasant garden, saves cost of own contribution)
- She contributes, and one or both of the others do (pleasant garden, incurs cost of contribution)
- She does not contribute, only one or neither of the others does (sparse garden, saves cost of own contribution)
- She contributes, but neither of the others does (sparse garden, incurs cost of own contribution)

Of these outcomes, the one listed at the top is clearly the best and the one listed at the bottom is clearly the worst. We want higher payoff numbers to indicate outcomes that are more highly regarded; so we give the top outcome the payoff 4 and the bottom one the payoff 1. (Sometimes payoffs are associated with an outcome's rank order; so, with four outcomes, 1 would be best and 4 worst, and smaller numbers would denote more preferred outcomes. When reading, you should carefully note which convention the author is using; when writing, you should carefully state which convention you are using.)

There is some ambiguity about the two middle outcomes. Let us suppose that each player regards a pleasant garden more highly than her own contribution. Then the outcome listed second gets payoff 3, and the outcome listed third gets payoff 2.

Suppose the players move sequentially. Emily has the first move, and chooses whether to contribute. Then, after observing what Emily has chosen, Nina makes her choice between contributing and not contributing. Finally, having observed what Emily and Nina have chosen, Talia makes a similar choice.[1]

Figure 3.6 shows the tree for this game. We have labeled the action nodes for easy reference. Emily moves at the initial node, *a*, and the branches corresponding to

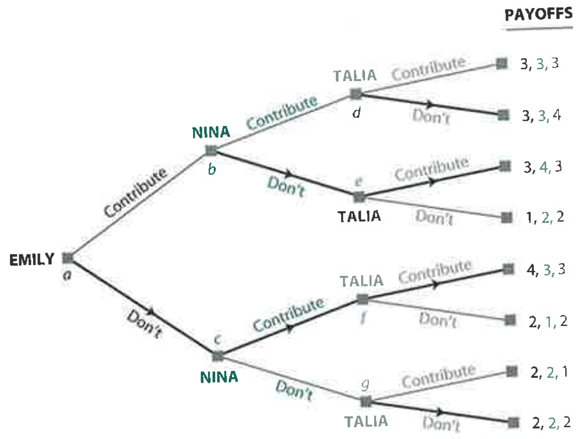

FIGURE 3.6 The Street Garden Game

[1] In later chapters, we vary the rules of this game—the order of moves and payoffs—and examine how such variation changes the outcomes.

her two choices, Contribute and Don't, respectively, lead to nodes *b* and *c*. At each of these nodes, Nina gets to move and to choose between Contribute and Don't. Her choices lead to nodes *d*, *e*, *f*, and *g*, at each of which Talia gets to move. Her choices lead to eight terminal nodes, where we show the payoffs in order (Emily, Nina, Talia).[2] For example, if Emily contributes, then Nina does not, and finally Talia does, then the garden is pleasant, and the two contributors get payoffs 3 each, while the noncontributor gets her top outcome with payoff 4; in this case, the payoff list is (3, 4, 3).

To apply rollback analysis to this game, we begin with the action nodes that come immediately before the terminal nodes—namely, *d*, *e*, *f*, and *g*. Talia moves at each of these nodes. At *d*, she faces the situation where both Emily and Nina have contributed. The garden is already assured to be pleasant; so, if Talia chooses Don't, she gets her best outcome, 4, whereas, if she chooses Contribute, she gets the next best, 3. Her preferred choice at this node is Don't. We show this preference both by thickening the branch for Don't and by adding an arrowhead; either one would suffice to illustrate Talia's choice. At node *e*, Emily has contributed and Nina has not; so Talia's contribution is crucial for a pleasant garden. Talia gets the payoff 3 if she chooses Contribute and 2 if she chooses Don't. Her preferred choice at *e* is Contribute. You can check Talia's choices at the other two nodes similarly.

Now we roll back the analysis to the preceding stage—namely, nodes *b* and *c*, where it is Nina's turn to choose. At *b*, Emily has contributed. Nina's reasoning now goes as follows: "If I choose Contribute, that will take the game to node *d*, where I know that Talia will choose Don't, and my payoff will be 3. (The garden will be pleasant, but I will have incurred the cost of my contribution.) If I choose Don't, the game will go to node *e*, where I know that Talia will choose Contribute, and I will get a payoff of 4. (The garden will be pleasant, and I will have saved the cost of my contribution.) Therefore I should choose Don't." Similar reasoning shows that at *c*, Nina will choose Contribute.

Finally, consider Emily's choice at the initial node, *a*. She can foresee the subsequent choices of both Nina and Talia. Emily knows that, if she chooses Contribute, these later choices will be Don't for Nina and Contribute for Talia. With two contributors, the garden will be pleasant but Emily will have incurred a cost; so her payoff will be 3. If Emily chooses Don't, then the subsequent choices will both be Contribute, and, with a pleasant garden and no cost of her own contribution, Emily's payoff will be 4. So her preferred choice at *a* is Don't.

The result of rollback analysis for this street garden game is now easily summarized. Emily will choose Don't, then Nina will choose Contribute, and finally

[2] Recall from the discussion of the general tree in Section 1 that the usual convention for sequential-move games is to list payoffs in the order in which the players move; however, in case of ambiguity or simply for clarity, it is good practice to specify the order explicitly.

Talia will choose Contribute. These choices trace a particular **path of play** through the tree—along the lower branch from the initial node, *a*, and then along the upper branches at each of the two subsequent nodes reached, *c* and *f*. In Figure 3.6, the path of play is easily seen as the continuous sequence of arrowheads joined tail to tip from the initial node to the terminal node fifth from the top of the tree. The payoffs that accrue to the players are shown at this terminal node.

Rollback analysis is simple and appealing. Here, we emphasize some features that emerge from it. First, notice that the **equilibrium path of play** of a sequential-move game misses most of the branches and nodes. Calculating the best actions that would be taken if these other nodes were reached, however, is an important part of determining the ultimate equilibrium. Choices made early in the game are affected by players' expectations of what would happen if they chose to do something other than their best actions and by what would happen if any opposing player chose to do something other than what was best for her. These expectations, based on predicted actions at out-of-equilibrium nodes (nodes associated with branches pruned in the process of rollback), keep players choosing optimal actions at each node. For instance, Emily's optimal choice of Don't at the first move is governed by the knowledge that, if she chose Contribute, then Nina would choose Don't, followed by Talia choosing Contribute; this sequence would give Emily the payoff 3, instead of the 4 that she can get by choosing Don't at the first move.

The rollback equilibrium gives a complete statement of all this analysis by specifying the optimal *strategy* for each player. Recall that a strategy is a complete plan of action. Emily moves first and has just two choices, so her strategy is quite simple and is effectively the same thing as her move. But Nina, moving second, acts at one of two nodes, at one if Emily has chosen Contribute and at the other if Emily has chosen Don't. Nina's complete plan of action has to specify what she would do in either case. One such plan, or strategy, might be "choose Contribute if Emily has chosen Contribute, choose Don't if Emily has chosen Don't." We know from our rollback analysis that Nina will not choose this strategy, but our interest at this point is in describing all the available strategies from which Nina can choose within the rules of the game. We can abbreviate and write C for Continue and D for Don't; then this strategy can be written as "C if Emily chooses C so that the game is at node *b*, D if Emily chooses D so that the game is at node *c*," or, more simply, "C at *b*, D at *c*," or even "CD" if the circumstances in which each of the stated actions is taken are evident or previously explained. Now it is easy to see that, because Nina has two choices available at each of the two nodes where she might be acting, she has available to her four plans, or strategies—"C at *b*, C at *c*," "C at *b*, D at *c*," "D at *b*, C at *c*," and "D at *b*, D at *c*," or "CC," "CD," "DC," and "DD." Of these strategies, the

rollback analysis and the arrows at nodes *b* and *c* of Figure 3.6 show that her optimal strategy is "DC."

Matters are even more complicated for Talia. When her turn comes, the history of play can, according to the rules of the game, be any one of four possibilities. Talia's turn to act comes at one of four nodes in the tree, one after Emily has chosen C and Nina has chosen C (node *d*), the second after Emily's C and Nina's D (node *e*), the third after Emily's D and Nina's C (node *f*), and the fourth after both Emily and Nina choose D (node *g*). Each of Talia's strategies, or complete plans of action, must specify one of her two actions for each of these four scenarios, or one of her two actions at each of her four possible action nodes. With four nodes at which to specify an action and with two actions from which to choose at each node, there are 2 times 2 times 2 times 2, or 16 possible combinations of actions. So Talia has available to her 16 possible strategies. One of them could be written as

"C at *d*, D at *e*, D at *f*, C at *g*" or "CDDC" for short,

where we have fixed the order of the four scenarios (the histories of moves by Emily and Nina) in the order of nodes *d*, *e*, *f*, and *g*. Then, with the use of the same abbreviation, the full list of 16 strategies available to Talia is

CCCC, CCCD, CCDC, CCDD, CDCC, CDCD, CDDC, CDDD,
DCCC, DCCD, DCDC, DCDD, DDCC, DDCD, DDDC, DDDD.

Of these strategies, the rollback analysis of Figure 3.6 and the arrows at nodes *d*, *e*, *f*, and *g* show that Talia's optimal strategy is DCCD.

Now we can express the findings of our rollback analysis by stating the strategy choices of each player—Emily chooses D from the two strategies available to her, Nina chooses DC from the four strategies available to her, and Talia chooses DCCD from the sixteen strategies available to her. When each player looks ahead in the tree to forecast the eventual outcomes of her current choices, she is calculating the optimal strategies of the other players. This configuration of strategies, D for Emily, DC for Nina, and DCCD for Talia, then constitutes the rollback equilibrium of the game.

We can put together the optimal strategies of the players to find the actual path of play that will result in the rollback equilibrium. Emily will begin by choosing D. Nina, following her strategy DC, chooses the action C in response to Emily's D. (Remember that Nina's DC means "choose D if Emily has played C, and choose C if Emily has played D.") According to the convention that we have adopted, Talia's actual action after Emily's D and then Nina's C—from node *f*—is the third letter in the four-letter specification of her strategies. Because Talia's optimal strategy is DCCD, her action along the path of play is C. Thus the actual path of play consists of Emily playing D, followed successively by Nina and Talia playing C.

To sum up, we have three distinct concepts:

1. The lists of available strategies for each player. The list, especially for later players, may be very long, because their actions in situations corresponding to all conceivable preceding moves by other players must be specified.
2. The optimal strategy, or complete plan of action for each player. This strategy must specify the player's best choices at each node where the rules of the game specify that she moves, even though many of these nodes will never be reached in the actual path of play. This specification is in effect the preceding movers' forecasting of what would happen if they took different actions and is therefore an important part of their calculation of their own best actions at the earlier nodes. The optimal strategies of all players together yield the rollback equilibrium.
3. The actual path of play in the rollback equilibrium, found by putting together the optimal strategies for all the players.

4 ORDER ADVANTAGES

In the rollback equilibrium of the street-garden game, Emily gets her best outcome (payoff 4), because she can take advantage of the opportunity to make the first move. When she chooses not to contribute, she puts the onus on the other two players—each can get her next-best outcome if and only if both of them choose to contribute. Most casual thinkers about strategic games have the preconception that such **first-mover advantage** should exist in all games. However, that is not the case. It is easy to think of games in which an opportunity to move second is an advantage. Consider the strategic interaction between two firms that sell similar merchandise from catalogs—say, Land's End and L.L. Bean. If one firm had to release its catalog first, and then the second firm could see what prices the first had set before printing its own catalog, then the second mover could just undercut its rival on all items and gain a tremendous competitive edge.

First-mover advantage comes from the ability to commit oneself to an advantageous position and to force the other players to adapt to it; **second-mover advantage** comes from the flexibility to adapt oneself to the others' choices. Whether commitment or flexibility is more important in a specific game depends on its particular configuration of strategies and payoffs; no generally valid rule can be laid down. We will come across examples of both kinds of advantages throughout this book. The general point that there need not be first-mover advantage, a point that runs against much common perception, is so important that we felt it necessary to emphasize at the outset.

When a game has a first- or second-mover advantage, each player may try to manipulate the order of play so as to secure for herself the advantageous position. Tactics for such manipulation are strategic moves, which we consider in Chapter 10.

5 ADDING MORE MOVES

We saw in Section 3 that adding more players increases the complexity of the analysis of sequential-play games. In this section, we consider another type of complexity that arises from adding additional moves to the game. We can do so most simply in a two-person game by allowing players to alternate moves more than once. Then the tree is enlarged in the same fashion as a multiple-player game tree would be, but later moves in the tree are made by the players who have made decisions earlier in the same game.

Many common games, such as tic-tac-toe, checkers, and chess, are two-person strategic games with such alternating sequential moves. The use of game trees and rollback should allow us to "solve" such games—to determine the rollback equilibrium outcome and the equilibrium strategies leading to that outcome. Unfortunately, as the complexity of the game grows and as strategies become more and more intricate, the search for an optimal solution becomes more and more difficult as well. In such cases, when manual solution is no longer really feasible, computer routines such as Gambit, mentioned in Chapter 2, become useful.

A. Tic-Tac-Toe

Start with the most simple of the three examples mentioned in the preceding paragraph, tic-tac-toe, and consider an easier-than-usual version in which two players (X and O) each try to be the first to get two of their symbols to fill any row, column, or diagonal of a two-by-two game board. The first player has four possible positions in which to put her X. The second player then has three possible actions at each of four decision nodes. When the first player gets to her second turn, she has two possible actions at each of 12 (4 times 3) decision nodes. As Figure 3.7 shows, even this mini-game of tic-tac-toe has a very complex game tree. This tree is actually not too complex, because the game is guaranteed to end after the first player moves a second time; but there are still 24 terminal nodes to consider.

We show this tree merely as an illustration of how complex game trees can become in even simple (or simplified) games. As it turns out, using rollback on the mini-game of tic-tac-toe leads us quickly to an equilibrium. Rollback shows

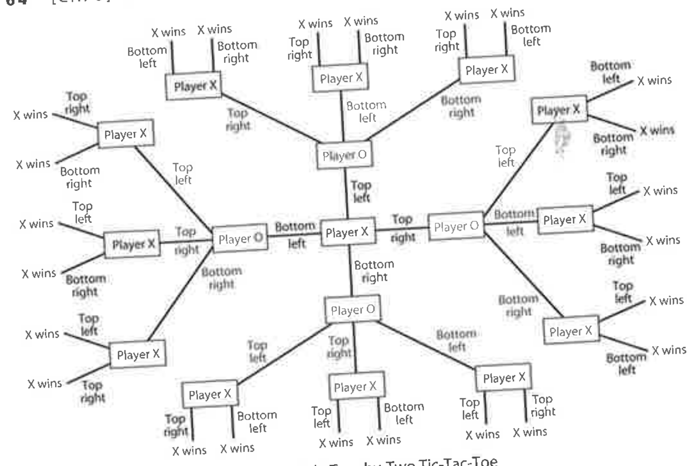

FIGURE 3.7 The Complex Tree for Simple Two-by-Two Tic-Tac-Toe

that all of the choices for the first player at her second move lead to the same outcome. There is no optimal action; any move is as good as any other move. Thus when the second player makes her first move, she also sees that each possible move yields the same outcome and she, too, is indifferent among her three choices at each of her four decision nodes. Finally, the same is true for the first player on her first move; any choice is as good as any other, so she is guaranteed to win the game.

Although this version of tic-tac-toe has an interesting tree, its solution is not as interesting. The first player always wins; so choices made by either player cannot affect the ultimate outcome. Most of us are more familiar with the three-by-three version of tic-tac-toe. To illustrate that version with a game tree, we would have to show that the first player has nine possible actions at the initial node, the second player has eight possible actions at each of nine decision nodes, and then the first player, on her second turn, has seven possible actions at each of 8 times 9 = 72 nodes, while the second player, on her second turn, has six possible actions at each of 7 times 8 times 9 = 504 nodes. This pattern continues until eventually the tree stops branching so rapidly because certain combinations of moves lead to a win for one player and the game ends. But no win is possible until at least the fifth move. Drawing the complete tree for this game requires a very large piece of paper or very tiny handwriting.

Most of you know, however, how to achieve at worst a tie when you play three-by-three tic-tac-toe. So there is a simple solution to this game that can be

found by rollback, and a learned strategic thinker can reduce the complexity of the game considerably in the quest for such a solution. It turns out that, as in the two-by-two version, many of the possible paths through the game tree are strategically identical. Of the nine possible initial moves, there are only three types; you put your X in either a corner position (of which there are four possibilities), a side position (of which there are also four possibilities), or the (one) middle position. Using this method to simplify the tree can help reduce the complexity of the problem and lead you to a description of an optimal rollback equilibrium strategy. Specifically, we could show that the player who moves second can always guarantee at least a tie with an appropriate first move and then by continually blocking the first player's attempts to get three symbols in a row.[3]

B. Chess

Although relatively small games, such as tic-tac-toe, can be solved using rollback, we showed above how rapidly the complexity of game trees can increase even in two-player games. Thus when we consider more complicated games, such as chess, finding a complete solution becomes much more difficult.

In chess, the players, White and Black, have a collection of 16 pieces in six distinct shapes, each of which is bound by specified rules of movement on the eight-by-eight game board shown in Figure 3.8.[4] White opens with a move, Black responds with one, and so on, in turns. All the moves are visible to the other player, and nothing is left to chance, as it would be in card games that include shuffling and dealing. Moreover, a chess game must end in a finite number of moves. The rules declare that a game is drawn if a given position on the board is repeated three times in the course of play. Because there are a finite number of ways to place the 32 (or fewer after captures) pieces on 64 squares, a game could not go on infinitely long without running up against this rule. Therefore in principle chess is amenable to full rollback analysis.

That rollback analysis has not been carried out, however. Chess has not been "solved" as tic-tac-toe has been. And the reason is that, for all its simplicity of rules, chess is a bewilderingly complex game. From the initial set position of

[3] If the first player puts her first symbol in the middle position, the second player must put her first symbol in a corner position. Then the second player can guarantee a tie by taking the third position in any row, column, or diagonal that the first player tries to fill. If the first player goes to a corner or a side position first, the second player can guarantee a tie by going to the middle first and then following the same blocking technique. Note that, if the first player picks a corner, the second player picks the middle, and the first player then picks the corner opposite from her original play, then the second player must not pick one of the remaining corners if she is to ensure at least a tie.

[4] An easily accessible statement of the rules of chess and much more is at Wikipedia, at http://en.wikipedia.org/wiki/Chess.

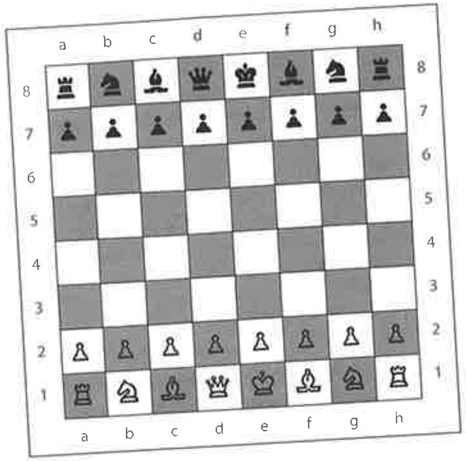

FIGURE 3.8 Chessboard

the pieces illustrated in Figure 3.8, White can open with any one of 20 moves,[5] and Black can respond with any of 20. Therefore 20 branches emerge from the first node of the game tree, each leading to a second node from each of which 20 more branches emerge. After only two moves, there are already 400 branches, each leading to a node from which many more branches emerge. And the total number of possible moves in chess has been estimated to be 10^{120}, or a "one" with 120 zeros after it. A supercomputer a thousand times as fast as your PC, making a trillion calculations a second, would need more than 10^{100} years to check out all these moves.[6] Astronomers offer us less than 10^{10} years before the sun turns into a red giant and swallows the earth.

The general point is that, although a game may be amenable in principle to a complete solution by rollback, its complete tree may be too complex to permit such solution in practice. Faced with such a situation, what is a player to do? We can learn a lot about this by reviewing the history of attempts to program computers to play chess.

[5]He can move one of eight pawns forward either one square or two or he can move one of the two knights in one of two ways (to squares a3, c3, f3, or h3).

[6]This would have to be done only once because, after the game has been solved, anyone can use the solution and no one will actually need to play. Everyone will know whether White has a win or whether Black can force a draw. Players will toss to decide who gets which color. They will then know the outcome, shake hands, and go home.

When computers first started to prove their usefulness for complex calculations in science and business, many mathematicians and computer scientists thought that a chess-playing computer program would soon beat the world champion. It took a lot longer, even though computer technology improved dramatically while human thought progressed much more slowly. Finally, in December 1992, a German chess program called Fritz2 beat world champion Gary Kasparov in some blitz (high-speed) games. Under regular rules, where each player gets $2\frac{1}{2}$ hours to make 40 moves, humans retained greater superiority for longer. A team sponsored by IBM put a lot of effort and resources into the development of a specialized chess-playing computer and its associated software. In February 1996, this package, called Deep Blue, was pitted against Gary Kasparov in a best-of-six series. Deep Blue caused a sensation by winning the first game, but Kasparov quickly figured out its weaknesses, improved his counterstrategies, and won the series handily. In the next 15 months, the IBM team improved Deep Blue's hardware and software, and the resulting Deeper Blue beat Kasparov in another best-of-six series in May 1997.

To sum up, computers have progressed in a combination of slow patches and some rapid spurts, while humans have held some superiority but have not been able to improve sufficiently fast to keep ahead. Closer examination reveals that the two use quite different approaches to thinking through the very complex game tree of chess.

When contemplating a move in chess, looking ahead to the end of the whole game may be too hard (for humans and computers both). How about looking part of the way—say, 5 or 10 moves ahead—and working back from there? The game need not end within this limited horizon; that is, the nodes that you reach after 5 or 10 moves will not generally be terminal nodes. Only terminal nodes have payoffs specified by the rules of the game. Therefore you need some indirect way of assigning plausible payoffs to nonterminal nodes, because you are not able to explicitly roll back from a full look-ahead. A rule that assigns such payoffs is called an **intermediate valuation function.**

In chess, humans and computer programs both use such partial look-ahead in conjunction with an intermediate valuation function. The typical method assigns values to each piece and to positional and combinational advantages that can arise during play. Quantification of values for different positions are made on the basis of the whole chess-playing community's experience of play in past games starting from such positions or patterns; this is called "knowledge." The sum of all the numerical values attached to pieces and their combinations in a position is the intermediate value of that position. A move is judged by the value of the position to which it is expected to lead after an explicit forward-looking calculation for a certain number—say, five or six—of moves.

The evaluation of intermediate positions has progressed furthest with respect to chess openings—that is, the first dozen or so moves of a game. Each opening can lead to any one of a vast multitude of further moves and positions, but experience enables players to sum up certain openings as being more or less likely to favor one player or the other. This knowledge has been written down in massive books of openings, and all top players and computer programs remember and use this information.

At the end stages of a game, when only a few pieces are left on the board, backward reasoning on its own is often simple enough to be doable and complete enough to give the full answer. The midgame, when positions have evolved into a level of complexity that will not simplify within a few moves, is the hardest to analyze. To find a good move from a midgame position, a well-built intermediate valuation function is likely to be more valuable than the ability to calculate another few moves further ahead.

This is where the art of chess playing comes into its own. The best human players develop an intuition or instinct that enables them to sniff out good opportunities and avoid subtle traps in a way that computer programs find hard to match. Computer scientists have found it generally very difficult to teach their machines the skills of pattern recognition that humans acquire and use instinctively—for example, recognizing faces and associating them with names. The art of the midgame in chess also is an exercise in recognizing and evaluating patterns in the same, still mysterious way. This is where Kasparov has his greatest advantage over Fritz2 or Deep Blue. It also explains why computer programs do better against humans at blitz or limited-time games: a human does not have the time to marshal his art of the midgame.

In other words, the best human players have subtle "chess knowledge," based on experience or the ability to recognize patterns, which endows them with a better intermediate valuation function. Computers have the advantage when it comes to raw or brute-force calculation. Thus although both human and computer players now use a mixture of look-ahead and intermediate valuation, they use them in different proportions: humans do not look so many moves ahead but have better intermediate valuations based on knowledge; computers have less sophisticated valuation functions but look ahead further by using their superior computational powers.

Recently, chess computers have begun to acquire more knowledge. When modifying Deep Blue in 1996 and 1997, IBM enlisted the help of human experts to improve the intermediate valuation function in its software. These consultants played repeatedly against the machine, noted its weaknesses, and suggested how the valuation function should be modified to correct the flaws. Deep Blue benefited from the contributions of the experts and their subtle kind of thinking, which results from long experience and an awareness of complex interconnections among the pieces on the board.

If humans can gradually make explicit their subtle knowledge and transmit it to computers, what hope is there for human players who do not get reciprocal

help from computers? At times in their 1997 encounter, Kasparov was amazed by the human or even superhuman quality of Deep Blue's play. He even attributed one of the computer's moves to "the hand of God." And matters can only get worse: the brute-force calculating power of computers is increasing rapidly while they are simultaneously, but more slowly, gaining some of the subtlety that constitutes the advantage of humans.

The abstract theory of chess says that it is a finite game that can be solved by rollback. The practice of chess requires a lot of "art" based on experience, intuition, and subtle judgment. Is this bad news for the use of rollback in sequential-move games? We think not. It is true that theory does not take us all the way to an answer for chess. But it does take us a long way. Looking ahead a few moves constitutes an important part of the approach that mixes brute-force calculation of moves with a knowledge-based assessment of intermediate positions. And, as computational power increases, the role played by brute-force calculation, and therefore the scope of the rollback theory, will also increase.

Evidence from the study of the game of checkers, as we describe below, suggests that a solution to chess may yet be feasible.

C. Checkers

An astonishing number of computer and person hours have been devoted to the search for a solution to chess. Less famously, but just as doggedly, researchers have been working on solving the somewhat less complex game of checkers. Recently, this work was rewarded; the game of checkers was declared "solved" in July 2007.[7]

Checkers is another two-player game played on an eight-by-eight board. Each player has 12 round game pieces of different colors, as shown in Figure 3.9, and players take turns moving their pieces diagonally on the board, jumping (and capturing) the opponent's pieces when possible. As in chess, the game ends and Player A wins when Player B is either out of pieces or unable to move; the game can also end in a draw if both players agree that neither can win.

Although the complexity of checkers pales somewhat in comparison to that of chess—the number of possible positions in checkers is approximately the square root of the number in chess—there are still 5×10^{20} possible positions, so drawing a game tree is out of the question. Conventional wisdom and evidence from world championships for years suggested that good play should lead to a draw, but there was no proof. Now a computer scientist in Canada has the proof—a computer program named Chinook that can play to a guaranteed tie.

[7]Our account is based on two reports in the journal *Science*. See "Program Proves That Checkers, Perfectly Played, Is a No-Win Situation," Adrian Cho, *Science* (317), 20 July 2007, pp. 308–309, and "Checkers Is Solved," Jonathan Schaeffer et al., *Science* (317), 14 September 2007, pp. 1518–1522.

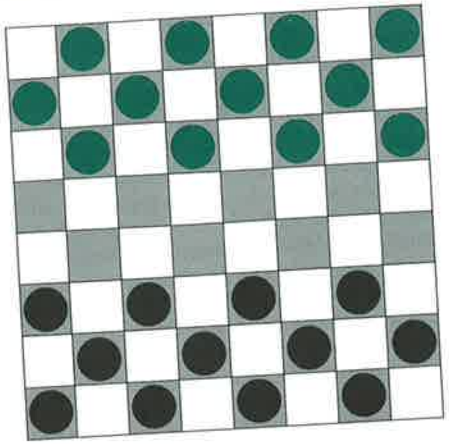

FIGURE 3.9 Checkers

Chinook was first created in 1989; played the world champion, Marion Tinsley, in 1992 (losing four to two with 33 draws) and again in 1994 (when Tinsley's health failed during a series of draws); was put on hold between 1997 and 2001 while its creators waited for computer technology to improve; and finally exhibited a loss-proof algorithm in the spring of 2007. That algorithm uses a combination of endgame rollback analysis and starting position forward analysis along with the equivalent of an intermediate valuation function to trace out the best moves within a database including all possible positions on the board.

The creators of Chinook describe the full game of checkers as "weakly solved"; they know that they can generate a tie and they have a strategy for reaching that tie from the start of the game. For all 39×10^{12} possible positions that include 10 or fewer pieces on the board, they describe checkers as "strongly solved"; not only do they know they can play to a tie, they can reach that tie from any of the possible positions that can arise once only 10 pieces remain. Their algorithm first solved the 10-piece endgames, then went back to the start to search out paths of play in which both players make optimal choices. The search mechanism, involving a complex system of evaluating the value of each intermediate position, invariably led to those 10-piece positions that generate a draw.

Thus our hope for the future of rollback analysis may not be misplaced. We know that for really simple games, we can find the rollback equilibrium by verbal reasoning without having to draw the game tree explicitly. For games having an intermediate range of complexity, verbal reasoning is too hard, but a complete tree can be drawn and used for rollback. Sometimes we may

enlist the aid of a computer to draw and analyze a moderately complicated game tree. For the most complex games, such as checkers and chess, we can draw only a small part of the game tree, and we must use a combination of two methods: (1) calculation based on the logic of rollback; and (2) rules of thumb for valuing intermediate positions on the basis of experience. The computational power of current algorithms has shown that even some games in this category are amenable to solution, provided one has the time and resources to devote to the problem.

Thankfully, most of the strategic games that we encounter in economics, politics, sports, business, and daily life are far less complex than chess or even checkers. The games may have a number of players who move a number of times; they may even have a large number of players or a large number of moves. But we have a chance at being able to draw a reasonable-looking tree for those games that are sequential in nature. The logic of rollback remains valid, and it is also often the case that, once you understand the idea of rollback, you can carry out the necessary logical thinking and solve the game without explicitly drawing a tree. Moreover, it is precisely at this intermediate level of difficulty, between the simple examples that we solved explicitly in this chapter and the insoluble cases such as chess, that computer software such as Gambit is most likely to be useful; this is indeed fortunate for the prospect of applying the theory to solve many games in practice.

6 EVIDENCE CONCERNING ROLLBACK

How well do actual participants in sequential-move games perform the calculations of rollback reasoning? There is very little systematic evidence, but classroom and research experiments with some games have yielded outcomes that appear to counter the predictions of the theory. Some of these experiments and their outcomes have interesting implications for the strategic analysis of sequential-move games.

For instance, many experimenters have had subjects play a single-round bargaining game in which two players, designated A and B, are chosen from a class or a group of volunteers. The experimenter provides a dollar (or some known total), which can be divided between them according to the following procedure. Player A proposes a split—for example, "75 to me, 25 to B." If player B accepts this proposal, the dollar is divided as proposed by A. If B rejects the proposal, neither player gets anything.

Rollback in this case predicts that B should accept any sum, no matter how small, because the alternative is even worse—namely, zero—and, foreseeing this, A should propose "99 to me, 1 to B." This particular outcome almost never happens.

Most players assigned the A role propose a much more equal split. In fact, 50–50 is the single most common proposal. Furthermore, most players assigned the B role turn down proposals that leave them 25% or less of the total and walk away with nothing; some reject proposals that would give them 40% of the pie.[8]

Many game theorists remain unpersuaded that these findings undermine the theory. They counter with some variant of the following argument: "The sums are so small as to make the whole thing trivial in the players' minds. The B players lose 25 or 40 cents, which is almost nothing, and perhaps gain some private satisfaction that they walked away from a humiliatingly small award. If the total were a thousand dollars, so that 25% of it amounted to real money, the B players would accept." But this argument does not seem to be valid. Experiments with much larger stakes show similar results. The findings from experiments conducted in Indonesia, with sums that were small in dollars but amounted to as much as three months' earnings for the participants, showed no clear tendency on the part of the A players to make less equal offers, although the B players tended to accept somewhat smaller shares as the total increased; similar experiments conducted in the Slovak Republic found the behavior of inexperienced players unaffected by large changes in payoffs.[9]

The participants in these experiments typically have no prior knowledge of game theory and no special computational abilities. But the game is extremely simple; surely even the most naive player can see through the reasoning, and answers to direct questions after the experiment generally show that most participants do. The results show not so much the failure of rollback as the theorist's error in supposing that each player cares only about her own money earnings. Most societies instill in their members a strong sense of fairness, which then causes most A players to offer 50–50 or something close and the B players to reject anything that is grossly unfair. This argument is supported by the observation that even in a most drastic variant called the *dictator game*, where the A player decides on the split and the B player has no choice at all, many As give significant shares to the Bs.[10]

[8]Read Richard H. Thaler, "Anomalies: The Ultimate Game," *Journal of Economic Perspectives*, vol. 2, no. 4 (fall 1988), pp. 195–206, and Douglas D. Davis and Charles A. Holt, *Experimental Economics* (Princeton: Princeton University Press, 1993), pp. 263–269, for a detailed account of this game and related ones.

[9]The results of the Indonesian experiment are reported in Lisa Cameron, "Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia," *Economic Inquiry*, vol. 37, no. 1 (January 1999), pp. 47–59. Slonim and Roth report results similar to Cameron's, but also found that offers (in all rounds of play) were rejected less often as the payoffs were raised. See Robert Slonim and Alvin Roth, "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic," *Econometrica*, vol. 66, no. 3 (May 1998), pp. 569–596.

[10]One could argue that this social norm of fairness may actually have value in the ongoing evolutionary game being played by the whole of society. Players who are concerned with fairness reduce transaction costs and the costs of fights that can be beneficial to society in the long run. These matters will be discussed in Chapters 11 and 12.

Some experiments have been conducted to determine the extent to which fairness plays into behavior in ultimatum and dictator games. The findings show that changes in the information available to players and to the experimenter have a significant effect on the final outcome. In particular, when the experimental design is changed so that not even the experimenter can identify who proposed (or accepted) the split, the extent of sharing drops noticeably. Evidence from the new field of "neuroeconomics" includes similar results. Alan Sanfey and colleagues conducted experiments of the ultimatum game in which they took MRI readings of the subjects' brains as they made their choices. They found "that activity in a region [of the brain] well known for its involvement in negative emotion" was stimulated in the responders (B players) when they rejected "unfair" (less than 50:50) offers. Thus deep instincts or emotions of anger and disgust seem to be implicated in these rejections. They also found that "unfair" (less than 50:50) offers were rejected less often when responders knew that the offerer was a computer than when they knew that the offerer was human.[11]

Another experimental game with similarly paradoxical outcomes goes as follows. Two players are chosen and designated A and B. The experimenter puts a dime on the table. Player A can take it or pass. If A takes the dime, the game is over, with A getting the 10 cents and B getting nothing. If A passes, the experimenter adds a dime, and now B has the choice of taking the 20 cents or passing. The turns alternate, and the pile of money grows until reaching some limit—say, a dollar—that is known in advance by both players.

We show the tree for this game in Figure 3.10. Because of the appearance of the tree, this type of game is often called the *centipede game*. You may not even need the tree to use rollback on this game. Player B is sure to take the dollar at the last stage; so A should take the 90 cents at the penultimate stage, and so on. Thus A should take the very first dime and end the game.

However, in most classroom or experimental settings, such games go on for at least a few rounds. Remarkably, by behaving "irrationally," the players as a

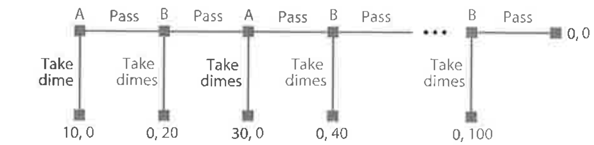

FIGURE 3.10 The Centipede Game

[11]See Alan Sanfey, James Rilling, Jessica Aronson, Leigh Nystrom, and Jonathan Cohen, "The Neural Basis of Economic Decision-Making in the Ultimatum Game," *Science*, vol. 300 (June 13, 2003), pp. 1755–1758.

group make more money than they would if they followed the logic of backward reasoning. Sometimes A does better and sometimes B, but sometimes they even solve this conflict or bargaining problem. In a classroom experiment that one of us (Dixit) conducted, one such game went all the way to the end. Player B collected the dollar, and quite voluntarily gave 50 cents to player A. Dixit asked A, "Did you two conspire? Is B a friend of yours?" and A replied, "No, we didn't even know each other before. But he is a friend now."

Once again, what is revealed is not that players cannot calculate and use game-theoretic logic but that their value systems and payoffs are different from those attributed to them by the theorist who predicted that the game should end with A taking the dime on the first step. We will come across some similar evidence of cooperation that seems to contradict rollback reasoning when we look at finitely repeated prisoners' dilemma games in Chapter 11.[12]

The examples discussed here seem to indicate that apparent violations of strategic logic can be explained by recognizing that people do not care merely about their own money payoffs, but internalize concepts such as fairness. But not all observed plays, contrary to the precepts of rollback, have some such explanation. People do fail to look ahead far enough, and they do fail to draw the appropriate conclusions from attempts to look ahead. For example, when issuers of credit cards offer favorable initial interest rates or no fees for the first year, many people fall for them without realizing that they may have to pay much more later. Therefore the game-theoretic analysis of rollback and rollback equilibria serves an advisory or prescriptive role as much as it does a descriptive role. People equipped with the theory of rollback are in a position to make better strategic decisions and get higher payoffs, no matter what is included in their payoff calculations. And game theorists can use their expertise to give valuable advice to those who are placed in complex strategic situations but lack the skill to determine their own best strategies.

7 STRATEGIES IN THE *SURVIVOR* GAME

The examples in the preceding sections were deliberately constructed to illustrate and elucidate basic concepts such as nodes, branches, moves, and strategies, as well as the technique of rollback. Now we show how all of them can be applied, by considering a real-life (or at least "reality-TV-life") situation.

In the summer of 2000, CBS television broadcast the first of the series of *Survivor* shows, which became an instant hit and helped launch the whole

[12]Once again, one wonders what would happen if the sum added at each step were a thousand dollars instead of a dime.

new genre of "reality TV." Leaving aside many complex details and some earlier stages not relevant for our purpose, the concept was as follows. A group of contestants, called a "tribe," was put on an uninhabited island and left largely to fend for themselves for food and shelter. Every 3 days they had to vote one of themselves out of the tribe. The person who had the most votes cast against him or her at a meeting of the remaining players (called the "tribal council") was the victim of the day. However, before each meeting of the tribal council, the survivors up to that point competed in a game of physical or mental skill that was devised by the producers of the game for that occasion, and the winner of this competition, called a "challenge," was immune to being voted off at the following meeting. Also, one could not vote against oneself. Finally, when two people were left, the seven who had been voted off most recently returned as a "jury" to pick one of the two remaining survivors as the million-dollar winner of the whole game.

The strategic problems facing all contestants were: (1) to be generally regarded as a productive contributor to the tribe's search for food and other tasks of survival, but to do so without being regarded as too strong a competitor and therefore a target for elimination, (2) to form alliances to secure blocks of votes to protect oneself from being voted off, (3) to betray these alliances when the numbers got too small and one had to vote against someone, but (4) to do so without seriously losing popularity with the other players, who would ultimately have the power of the vote on the jury.

We pick up the story when just three contestants were left: Rudy, Kelly, and Rich. Of them, Rudy was the oldest contestant, an honest and blunt person who was very popular with the contestants who had been previously voted off. It was generally agreed that, if he was one of the last two, then he would be voted the million-dollar winner. So it was in the interests of both Kelly and Rich that they should face each other, rather than face Rudy, in the final vote. But neither wanted to be seen as instrumental in voting off Rudy. With just three contestants left, the winner of the immunity challenge is effectively decisive in the cast-off vote, because the other two must vote against each other. Thus the jury would know who was responsible for voting off Rudy and, given his popularity, would regard the act of voting him off with disfavor. The person doing so would harm his or her chances in the final vote. This was especially a problem for Rich, because he was known to have an alliance with Rudy.

The immunity challenge was one of stamina; each contestant had to stand on an awkward support and lean to hold one hand in contact with a totem on a central pole, called the "immunity idol." Anyone whose hand lost contact with the idol, even for an instant, lost the challenge; the one to hold on longest was the winner.

An hour and a half into the challenge, Rich figured out that his best strategy was to deliberately lose this immunity challenge. Then, if Rudy won immunity, he would maintain his alliance and keep Rich—Rudy was known to be a man who always kept his word. Rich would lose the final vote to Rudy in this case, but that would make him no worse off than if he won the challenge and kept Rudy. If Kelly won immunity, the much more likely outcome, then it would be in her interest to vote off Rudy—she would have at least some chance against Rich but zero against Rudy. Then Rich's chances of winning were quite good. Whereas, if Rich himself held on, won immunity, and then voted off Rudy, his chances against Kelly would be decreased by the fact that he voted off Rudy.

So Rich deliberately stepped off, and later explained his reasons quite clearly to the camera. His calculation was borne out. Kelly won that challenge and voted off Rudy. And, in the final jury vote between Rich and Kelly, Rich won by one vote.

Rich's thinking was essentially a rollback analysis along a game tree. He did this analysis instinctively, without drawing the tree, while standing awkwardly and holding on to the immunity idol, but it took him an hour and a half to come to his conclusion. With all due credit to Rich, we show the tree explicitly, and can reach the answer faster.

Figure 3.11 shows the tree. You can see that it is much more complex than the trees encountered in earlier sections. It has more branches and moves; in addition, there are uncertain outcomes, and the chances of winning or losing in various alternative situations have to be estimated instead of being known precisely. But you will see how we can make reasonable assumptions about these chances and proceed with the analysis.

At the initial node, Rich decides whether to continue or to give up in the immunity challenge. In either case, the winner of the challenge cannot be forecast with certainty; this is indicated in the tree by letting "Nature" make the choice, as we did with the coin-toss situation in Figure 3.1. If Rich continues, Nature chooses the winner from the three contestants. We don't know the actual probabilities, but we will assume particular values for exposition and point out the crucial assumptions. The supposition is that Kelly has a lot of stamina and that Rudy, being the oldest, is not likely to win. So we posit the following probabilities of a win when Rich chooses to continue: 0.5 (50%) for Kelly, 0.45 for Rich, and only 0.05 for Rudy. If Rich gives up on the challenge, Nature picks the winner of the immunity challenge randomly between the two who remain; in this case, we assume that Kelly wins with probability 0.9 and Rudy with probability 0.1.

The rest of the tree follows from each of the three possible winners of the challenge. If Rudy wins, he keeps Rich as he promised, and the jury votes Rudy the winner.[13] If Rich wins immunity, he has to decide whether to keep Kelly or

[13]Technically, Rudy faces a choice between keeping Rich or Kelly at the action node after he wins the immunity challenge. Because everyone placed zero probability on his choosing Kelly (owing to the Rich–Rudy alliance), we illustrate only Rudy's choice of Rich. The jury, similarly, has a choice between Rich and Rudy at the last action node along this branch of play. Again, the foregone conclusion is that Rudy wins in this case.

Rudy. If he keeps Rudy, the jury votes for Rudy. If he keeps Kelly, it is not certain whom the jury chooses. We assume that Rich alienates some jurors by turning on Rudy and that, despite being better liked than Kelly, he gets the jury's vote in this situation only with probability 0.4. Similarly, if Kelly wins immunity, she can either keep Rudy and lose the jury's vote, or keep Rich. If she keeps Rich, his probability of winning the jury's vote is higher, at 0.6, because in this case he is both better liked by the jury and hasn't voted off Rudy.

What about the players' actual payoffs? We can safely assume that both Rich and Kelly want to maximize the probability of his or her emerging as the ultimate winner of the $1 million. Rudy similarly wants to get the prize, but keeping

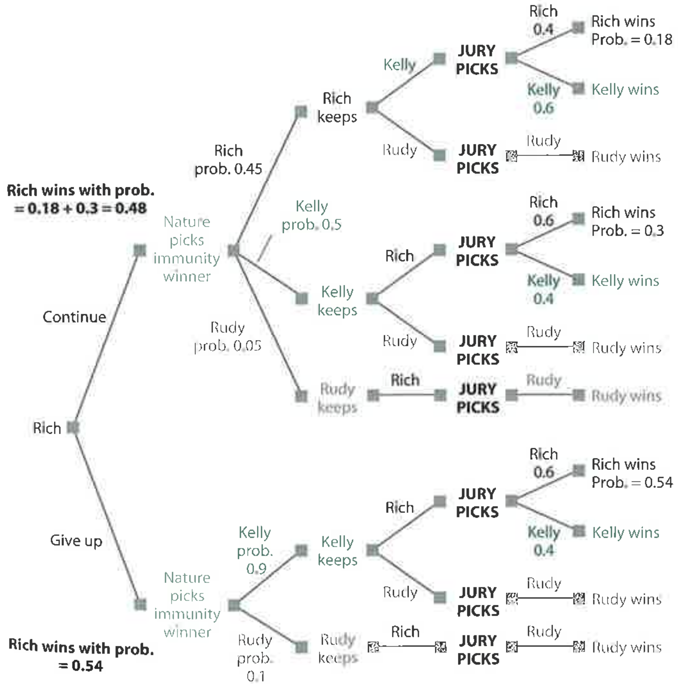

FIGURE 3.11 *Survivor* Immunity Game Tree

his word to Rich is paramount. With these preferences of the various players in mind, Rich can now do rollback analysis along the tree to determine his own initial choice.

Rich knows that, if he wins the immunity challenge (the uppermost path after his own first move and Nature's move), he will have to keep Kelly to have a 40% chance of eventual victory; keeping Rudy at this stage would mean a zero probability of eventual victory. Rich can also calculate that, if Kelly wins the immunity challenge (which occurs once in each of the upper and lower halves of the tree), she will choose to keep him for similar reasons, and then the probability of his eventual victory will be 0.6.

What are Rich's chances as he calculates them at the initial node? If Rich chooses Give Up at the initial node, then there is only one way for him to emerge as the eventual winner—if Kelly wins immunity (probability 0.9), if she then keeps Rich (probability 1), and if the jury votes for Rich (probability 0.6). Because all three things need to happen for Rich to win, his overall probability of victory is the product of the three probabilities—namely, $0.9 \times 1 \times 0.6 = 0.54$.[14] If Rich chooses Continue at the initial node, then there are two ways in which he can win. First, he wins the game if he wins the immunity challenge (probability 0.45), if he then eliminates Rudy (probability 1), and if he still wins the jury's vote against Kelly (probability 0.4); the total probability of winning in this way is $0.45 \times 0.4 = 0.18$. Second, he wins the game if Kelly wins the challenge (probability 0.5), if she eliminates Rudy (probability 1), and if Rich gets the jury's vote (probability 0.6); total probability here is $0.5 \times 0.6 = 0.3$. Rich's overall probability of eventual victory if he chooses Continue is the sum of the probabilities of these two paths to victory—namely, $0.18 + 0.3 = 0.48$.

Rich can now compare his probability of winning the million dollars when he chooses Give Up (0.54) with his probability of winning when he chooses Continue (0.48). Given the assumed values of the various probabilities in the tree, Rich has a better chance of victory if he gives up. Thus, Give Up is his optimal strategy. Although this result is based on assuming specific numbers for the probabilities, Give Up remains Rich's optimal strategy as long as (1) Kelly is very likely to win the immunity challenge once Rich gives up and (2) Rich wins the jury's final vote more often when Kelly has voted out Rudy than when Rich has done so.[15]

This example serves several purposes. Most important, it shows how a complex tree, with much external uncertainty and missing information about precise probabilities, can still be solved by using rollback analysis. We

[14] Readers who need instruction or a refresher course in the rules for combining probabilities will find a quick tutorial in the Appendix to Chapter 7.

[15] Readers who can handle the algebra of probabilities can solve this game by using more general symbols instead of specific numbers for the probabilities, as in Exercise U9 of this chapter.

hope this gives you some confidence in using the method and some training in converting a somewhat loose verbal account into a more precise logical argument. You might counter that Rich did this reasoning without drawing any trees. But knowing the system or general framework greatly simplifies the task even in new and unfamiliar circumstances. Therefore it is definitely worth the effort to acquire the systematic skill.

A second purpose is to illustrate the seemingly paradoxical strategy of "losing to win." Another instance of this strategy can be found in some sporting competitions that are held in two rounds, such as the soccer World Cup. The first round is played on a league basis in several groups of four teams each. The top two teams from each group then go to the second round, where they play others chosen according to a prespecified pattern; for example, the top-ranked team in group A meets the second-ranked team in group B, and so on. In such a situation, it may be good strategy for a team to lose one of its first-round matches if this loss causes it to be ranked second in its group; that ranking might earn it a subsequent match against a team that, for some particular reason, it is more likely to beat than the team that it would meet if it had placed first in its group in the first round.

SUMMARY

Sequential-move games require players to consider the future consequences of their current moves before choosing their actions. Analysis of pure sequential-move games generally requires the creation of a *game tree*. The tree is made up of *nodes* and *branches* that show all of the possible actions available to each player at each of her opportunities to move, as well as the payoffs associated with all possible outcomes of the game. Strategies for each player are complete plans that describe actions at each of the player's decision nodes contingent on all possible combinations of actions made by players who acted at earlier nodes. The equilibrium concept employed in sequential-move games is that of *rollback equilibrium*, in which players' equilibrium strategies are found by looking ahead to subsequent nodes and the actions that would be taken there and by using these forecasts to calculate one's current best action. This process is known as *rollback*, or *backward induction*.

Different types of games entail advantages for different players, such as *first-mover advantages*. The inclusion of many players or many moves enlarges the game tree of a sequential-move game but does not change the solution process. In some cases, drawing the full tree for a particular game may require more space or time than is feasible. Such games can often be solved by identifying strategic similarities between actions that reduces the size of the tree or by simple logical thinking.

When solving larger games, verbal reasoning can lead to the rollback equilibrium if the game is simple enough or a complete tree may be drawn and analyzed. If the game is sufficiently complex that verbal reasoning is too difficult and a complete tree is too large to draw, we may enlist the help of a computer program. Checkers has been "solved" with the use of such a program, although full solution of chess will remain beyond the powers of computers for a long time. In actual play of these truly complex games, elements of both art (identification of patterns and of opportunities versus peril) and science (forward-looking calculations of the possible outcomes arising from certain moves) have a role in determining player moves.

Tests of the theory of sequential-move games seem to suggest that actual play shows the irrationality of the players or the failure of the theory to predict behavior adequately. The counterargument points out the complexity of actual preferences for different possible outcomes and the usefulness of strategic theory for identifying optimal actions when actual preferences are known.

KEY TERMS

action node (48)
backward induction (56)
branch (48)
decision node (48)
decision tree (48)
equilibrium path of play (60)
extensive form (48)
first-mover advantage (62)
game tree (48)
initial node (48)
intermediate valuation function (67)
move (50)
node (48)
path of play (60)
prune (54)
rollback (56)
rollback equilibrium (56)
root (48)
second-mover advantage (62)
terminal node (50)

SOLVED EXERCISES

S1. Suppose two players, Hansel and Gretel, take part in a sequential-move game. Hansel moves first, Gretel moves second, and each player moves only once.

(a) Draw a game tree for a game in which Hansel has two possible actions (Up or Down) at each node and Gretel has three possible actions (Top, Middle, or Bottom) at each node. How many of each node type—decision and terminal—are there?

(b) Draw a game tree for a game in which Hansel and Gretel each have three possible actions (Sit, Stand, or Jump) at each node. How many of the two node types are there?

(c) Draw a game tree for a game in which Hansel has four possible actions (North, South, East, or West) at each node and Gretel has two possible actions (Stay or Go) at each node. How many of the two node types are there?

S2. In each of the following games, how many pure strategies (complete plans of action) are available to each player? List out all of the pure strategies for each player.

(a)

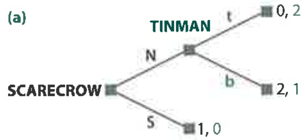

(b)

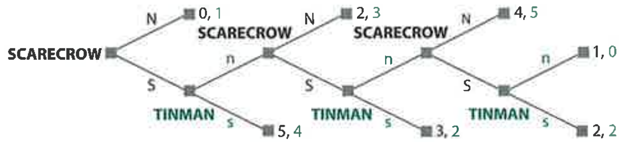

(c)

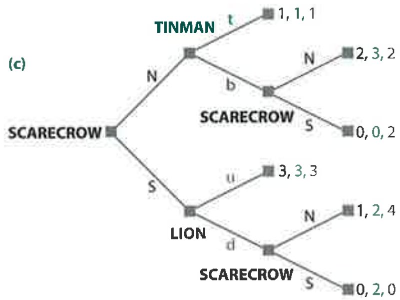

S3. For each of the games illustrated in Exercise S2, identify the rollback equilibrium outcome and the complete equilibrium strategy for each player.

S4. Consider the rivalry between Airbus and Boeing to develop a new commercial jet aircraft. Suppose Boeing is ahead in the development process and

Airbus is considering whether to enter the competition. If Airbus stays out, it earns zero profit, whereas Boeing enjoys a monopoly and earns a profit of $1 billion. If Airbus decides to enter and develop the rival airplane, then Boeing has to decide whether to accommodate Airbus peaceably or to wage a price war. In the event of peaceful competition, each firm will make a profit of $300 million. If there is a price war, each will lose $100 million because the prices of airplanes will fall so low that neither firm will be able to recoup its development costs.

Draw the tree for this game. Find the rollback equilibrium and describe the firms' equilibrium strategies.

S5. Consider a game in which two players, Fred and Barney, take turns removing matchsticks from a pile. They start with 21 matchsticks, and Fred goes first. On each turn, each player may remove either one, two, three, or four matchsticks. The player to remove the last matchstick wins the game.

(a) Suppose there are only six matchsticks left, and it is Barney's turn. What move should Barney make to guarantee himself victory? Explain your reasoning.

(b) Suppose there are 12 matchsticks left, and it is Barney's turn. What move should Barney make to guarantee himself victory? (Hint: Use your answer to part (a) and roll back.)

(c) Now start from the beginning of the game. If both players play optimally, who will win?

(d) What are the optimal strategies (complete plans of action) for each player?

S6. Consider the game in the previous exercise. Suppose the players have reached a point where it is Fred's move and there are just five matchsticks left.

(a) Draw the game tree for the game starting with five matchsticks.

(b) Find the rollback equilibria for this game starting with five matchsticks.

(c) Would you say this five-matchstick game has a first-mover advantage or a second-mover advantage?

(d) Explain why you found more than one rollback equilibrium. How is your answer related to the optimal strategies you found in part (c) of the previous exercise?

S7. A slave has just been thrown to the lions in the Roman Colosseum. Three lions are chained down in a line, with Lion 1 closest to the slave. Each lion's chain is short enough that he can only reach the two players immediately adjacent to him.

The game proceeds as follows. First, Lion 1 decides whether or not to eat the slave.

If Lion 1 has eaten the slave, then Lion 2 decides whether or not to eat Lion 1 (who is then too heavy to defend himself). If Lion 1 has not eaten the slave, then Lion 2 has no choice: he cannot try to eat Lion 1, because a fight would kill both lions.

Similarly, if Lion 2 has eaten Lion 1, then Lion 3 decides whether or not to eat Lion 2.

Each lion's preferences are fairly natural: best (4) is to eat and stay alive, next best (3) is to stay alive but go hungry, next (2) is to eat and be eaten, and worst (1) is to go hungry and be eaten.

(a) Draw the game tree, with payoffs, for this three-player game.

(b) What is the rollback equilibrium to this game? Make sure to describe the strategies, not just the payoffs.

(c) Is there a first-mover advantage to this game? Explain why or why not.

(d) How many complete strategies does each lion have? List them.

S8. Consider three major department stores—Big Giant, Titan, and Frieda's—contemplating opening a branch in one of two new Boston-area shopping malls. Urban Mall is located close to the large and rich population center of the area; it is relatively small and can accommodate at most two department stores as "anchors" for the mall. Rural Mall is farther out in a rural and relatively poorer area; it can accommodate as many as three anchor stores. None of the three stores wants to have branches in both malls because there is sufficient overlap of customers between the malls that locating in both would just mean competing with itself. Each store prefers to be in a mall with one or more other department stores than to be alone in the same mall, because a mall with multiple department stores will attract sufficiently many more total customers that each store's profit will be higher. Further, each store prefers Urban Mall to Rural Mall because of the richer customer base. Each store must choose between trying to get a space in Urban Mall (knowing that if the attempt fails, it will try for a space in Rural Mall) and trying to get a space in Rural Mall right away (without even attempting to get into Urban Mall).

In this case, the stores rank the five possible outcomes as follows: 5 (best), in Urban Mall with one other department store; 4, in Rural Mall with one or two other department stores; 3, alone in Urban Mall; 2, alone in Rural Mall; and 1 (worst), alone in Rural Mall after having attempted to get into Urban Mall and failed, by which time other nondepartment stores have signed up the best anchor locations in Rural Mall.

The three stores are sufficiently different in their managerial structures that they experience different lags in doing the paperwork required to request an expansion space in a new mall. Frieda's moves quickly, followed by Big Giant, and finally by Titan, which is the least efficient in readying a location plan. When all three have made their requests, the malls decide which stores to let in. Because of the name recognition that both Big Giant and Titan have with the potential customers, a mall would take either (or both) of those stores before it took Frieda's. Thus, Frieda's does not get one of the two spaces in Urban Mall if all three stores request those spaces; this is true even though Frieda's moves first.

(a) Draw the game tree for this mall location game.

(b) Illustrate the rollback pruning process on your game tree and use the pruned tree to find the rollback equilibrium. Describe the equilibrium by using the (complete) strategies employed by each department store. What are the payoffs to each store at the rollback equilibrium outcome?

S9. (Optional) Consider the following ultimatum bargaining game, which has been studied in laboratory experiments. The Proposer moves first, and proposes a split of $10 between himself and the Responder. Any whole-dollar split may be proposed. For example, the Proposer may offer to keep the whole $10 for himself, he may propose to keep $9 for himself and give $1 to the Responder, $8 to himself and $2 to the Responder, and so on. (Note that the Proposer therefore has eleven possible choices.) After seeing the split, the Responder can choose to accept the split or reject it. If the Responder accepts, both players get the proposed amounts. If she rejects, both players get $0.

(a) Write out the game tree for this game.

(b) How many complete strategies does each player have?

(c) What is the rollback equilibrium to this game, assuming the players care only about their cash payoffs?

(d) Suppose Rachel the Responder would accept any offer of $3 or more, and reject any offer of $2 or less. Suppose Pete the Proposer knows Rachel's strategy, and he wants to maximize his cash payoff. What strategy should he use?

(e) Rachel's true payoff (her "utility") might not be the same as her cash payoff. What other aspects of the game might she care about? Given your answer, propose a set of payoffs for Rachel that would make her strategy optimal.

(f) In laboratory experiments, players typically do not play the rollback equilibrium. Proposers typically offer an amount between $2 and $5 to the Responder. Responders often reject offers of $3, $2, and especially $1. Explain why you think this might occur.

UNSOLVED EXERCISES

U1. "In a sequential-move game, the player who moves first is sure to win." Is this statement true or false? State the reason for your answer in a few brief sentences, and give an example of a game that illustrates your answer.

U2. In each of the following games, how many pure strategies (complete plans of action) are available to each player? List all of the pure strategies for each player.

(a)

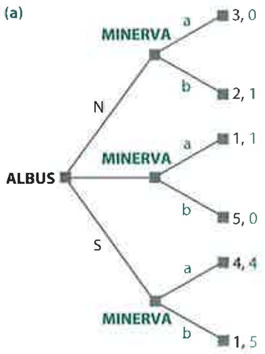

(b)

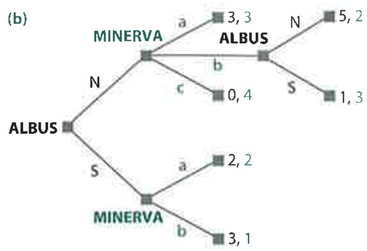

(c)

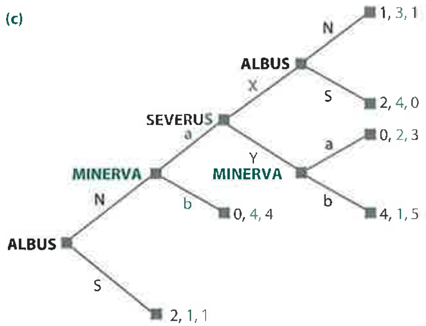

U3. For each of the games illustrated in Exercise U2, identify the rollback equilibrium outcome and the complete equilibrium strategy for each player.

U4. Two distinct proposals, A and B, are being debated in Washington. The Congress likes proposal A, and the president likes proposal B. The proposals are not mutually exclusive; either or both or neither may become law. Thus there are four possible outcomes, and the rankings of the two sides are as follows, where a larger number represents a more favored outcome.

Outcome	Congress	President
A becomes law	4	1
B becomes law	1	4
Both A and B become law	3	3
Neither (status quo prevails)	2	2

(a) The moves in the game are as follows. First, the Congress decides whether to pass a bill and whether it is to contain A or B or both. Then the president decides whether to sign or veto the bill. Congress does not have enough votes to override a veto. Draw a tree for this game and find the rollback equilibrium.

(b) Now suppose the rules of the game are changed in only one respect: the president is given the extra power of a line-item veto. Thus, if the Congress passes a bill containing both A and B, the president may choose not only to sign or veto the bill as a whole, but also to veto just one of the two items. Show the new tree and find the rollback equilibrium.

(c) Explain intuitively why the difference between the two equilibria arises.

U5. Two players, Amy and Beth, play the following game with a jar containing 100 pennies. The players take turns; Amy goes first. Each time it is a player's turn, she takes between one and 10 pennies out of the jar. The player whose move empties the jar wins.

(a) If both players play optimally, who will win the game? Does this game have a first-mover advantage? Explain your reasoning.

(b) What are the optimal strategies (complete plans of action) for each player?

(c) Now suppose we change the rules so that the player whose move empties the jar loses. Does this game have a first-mover advantage? Explain your reasoning.

(d) In this second variant, what are the optimal strategies for each player?

U6. Now Amy and Beth play a game with two jars, each containing 100 pennies. The players take turns; Amy goes first. Each time it is a player's turn, she chooses one of the jars and removes anywhere from one to 10 pennies from it. The player whose move leaves both jars empty wins. (Note that when a player empties the second jar, the first jar must already have been emptied in some previous move by one of the players.)

(a) Does this game have a first-mover advantage or a second-mover advantage? Explain which player can guarantee victory, and how she can do it. (Hint: simplify the game by starting with a smaller number of pennies in each jar, and see if you can generalize your finding to the actual game.)

(b) What are the optimal strategies (complete plans of action) for each player? (Hint: First think of a starting situation in which both jars have equal numbers of pennies. Then consider starting positions in which the two jars differ by one to 10 pennies. Finally, consider starting positions in which the jars differ by more than 10 pennies.)

U7. Modify Exercise S7 so that there are now four lions.

(a) Draw the game tree, with payoffs, for this four-player game.

(b) What is the rollback equilibrium to this game? Make sure to describe the strategies, not just the payoffs.

(c) Is the additional lion good or bad for the slave? Explain.

U8. To give Mom a day of rest, Dad plans to take his two children, Bart and Cassie, on an outing on Sunday. Bart prefers to go to the amusement park (A), whereas Cassie prefers to go to the science museum (S). Each child gets 3 units of utility from his/her more preferred activity, and only 2 units of utility from his/her less preferred activity. Dad gets 2 units of utility for either of the two activities.

To choose their activity, Dad plans first to ask Bart for his preference, then to ask Cassie after she hears Bart's choice. Each child can choose either the amusement park (A) or the science museum (S). If both children choose the same activity, then that is what they will all do. If the children choose different activities, Dad will make a tie-breaking decision. As the parent, Dad has an additional option: he can choose the amusement park, the science museum, or his personal favorite, the mountain hike (M). Bart and Cassie each get 1 unit of utility from the mountain hike, and Dad gets 3 units of utility from the mountain hike.

Because Dad wants his children to cooperate with one another, he gets 2 extra units of utility if the children choose the same activity (no matter which one of the two it is).

(a) Draw the game tree, with payoffs, for this three-person game.

(b) What is the rollback equilibrium to this game? Make sure to describe the strategies, not just the payoffs.

(c) How many different complete strategies does Bart have? Explain.

(d) How many complete strategies does Cassie have? Explain.

U9. **(Optional—more difficult)** Consider the *Survivor* game tree illustrated in Figure 3.11. We might not have guessed exactly the values Rich estimated for the various probabilities, so let's generalize this tree by considering other possible values. In particular, suppose that the probability of winning the immunity challenge when Rich chooses Continue is x for Rich, y for Kelly, and $1 - x - y$ for Rudy; similarly, the probability of winning when Rich gives up is z for Kelly and $1 - z$ for Rudy. Further, suppose that Rich's chance of being picked by the jury is p if he has won immunity and has voted Rudy off the island; his chance of being picked is q if Kelly has won immunity and has voted Rudy off. Continue to assume that if Rudy wins immunity, he keeps Rich with probability 1, and that Rudy wins the game with probability 1 if he ends up in the final two. Note that in the example of Figure 3.11, we had $x = 0.45$, $y = 0.5$, $z = 0.9$, $p = 0.4$, and $q = 0.6$. (In general, the variables p and q need not sum to 1, though this happened to be true in Figure 3.11.)

(a) Find an algebraic formula, in terms of x, y, z, p, and q, for the probability that Rich wins the million dollars if he chooses Continue. (Note: Your formula might not contain all of these variables.)

(b) Find a similar algebraic formula for the probability that Rich wins the million dollars if he chooses Give Up. (Again, your formula might not contain all of the variables.)

(c) Use these results to find an algebraic inequality telling us under what circumstances Rich should choose Give Up.

(d) Suppose all the values are the same as in Figure 3.11 except for z. How high or low could z be so that Rich would still prefer to Give Up? Explain intuitively why there are some values of z for which Rich is better off choosing Continue.

(e) Suppose all the values are the same as in Figure 3.11 except for p and q. Assume that since the jury is more likely to choose a "nice" person who doesn't vote Rudy off, we should have $p > 0.5 > q$. For what values of the ratio (p/q) should Rich choose Give Up? Explain intuitively why there are some values of p and q for which Rich is better off choosing Continue.

4

Simultaneous-Move Games with Pure Strategies I: Discrete Strategies

Recall from Chapter 2 that games are said to have simultaneous moves if players must move without knowledge of what their rivals have chosen to do. It is obviously so if players choose their actions at exactly the same time. A game is also simultaneous when players choose their actions in isolation, with no information about what other players have done or will do, even if the choices are made at different hours of the clock. (For this reason, simultaneous-move games have *imperfect information* in the sense we defined in Chapter 2, Section 2.D.) This chapter focuses on games that have such purely simultaneous interactions among players. We consider a variety of types of simultaneous games, introduce a solution concept called Nash equilibrium for these games, and study games with one equilibrium, many equilibria, or no equilibrium at all.

Many familiar strategic situations can be described as simultaneous-move games. The various producers of television sets, stereos, or automobiles make decisions about product design and features without knowing what rival firms are doing about their own products. Voters in U.S. elections simultaneously cast their individual votes; no voter knows what the others have done when she makes her own decision. The interaction between a soccer goalie and an opposing striker during a penalty kick requires both players to make their decisions simultaneously—the goalie cannot afford to wait until the ball has actually been kicked to decide which way to go, because then it would be far too late.

When a player in a simultaneous-move game chooses her action, she obviously does so without any knowledge of the choices made by other players. She also