A Beautiful Equilibrium

BIG GAME OF COORDINATION

Fred and Barney are Stone Age rabbit hunters. One evening, while carousing, they happen to engage in some shop talk. As they exchange information and ideas, they realize that by cooperating they could hunt much bigger game, such as stag or bison. One person on his own cannot expect any success hunting either stag or bison. But done jointly, each day's stag or bison hunting is expected to yield six times as much meat as a day's rabbit hunting by one person. Cooperation promises great advantage: each hunter's share of meat from a big-game hunt is three times what he can get hunting rabbits on his own.

The two agree to go big-game hunting together the following day and return to their respective caves. Unfortunately, they caroused too well, and both have forgotten whether they decided to go after stag or bison. The hunting grounds for the two species are in opposite directions. There were no cell phones in those days, and this was before the two became neighbors, so one could not quickly visit the other's cave to ascertain where to go. Each would have to make the decision the next morning in isolation.

Therefore the two end up playing a simultaneous-move game of deciding where to go. If we call each hunter's quantity of meat from a day's rabbit hunting 1, then the share of each from successful coordination in hunting either stag or bison is 3. So the payoff table of the game is as shown here:

		Barney's choice		
		Stag	Bison	Rabbit
Fred's choice	Stag	3 3	0 0	1 0
	Bison	0 0	3 3	1 0
	Rabbit	0 1	0 1	1 1

This game differs from the prisoners' dilemma of the previous chapter in many ways. Let us focus on one crucial difference. Fred's best choice depends on what Barney does, and vice versa. For neither player is there a strategy that is best regardless of what the other does; unlike in the prisoners' dilemma, this game has no dominant strategies. So each player has to think about the other's choice and figure out his own best choice in light of that.

Fred's thinking goes as follows: "If Barney goes to the grounds where the stags are, then I will get my share of the large catch if I go there too, but nothing if I go to the bison grounds. If Barney goes to the bison grounds, things are the other way around. Rather than take the risk of going to one of these areas and finding that Barney has gone to the other, should I go by myself after rabbits and make sure of my usual, albeit small, quantity of meat? In other words, should I take 1 for sure instead of risking either 3 or nothing? It depends on what I think Barney is likely to do, so let me put myself in his shoes (bare feet?) and think what he is thinking. Oh, he is wondering what I am likely to do and is trying to put himself in my shoes! Is there any end to this circular thinking about thinking?"

SQUARING THE CIRCLE

John Nash's beautiful equilibrium was designed as a theoretical way to square just such circles of thinking about thinking about other people's choices in games of strategy.* The idea is to look for an outcome where each player in the game chooses the strategy that best serves his or her own interest, in response to the other's strategy. If such a configuration of strategies arises, neither player has any reason to change his choice unilaterally. Therefore, this is a potentially stable outcome of a game where the players make individual and simultaneous choices of strategies. We begin by illustrating the idea with some examples of it in action. Later in this chapter we discuss how well it predicts outcomes in various games; we find reasons for cautious optimism and for making Nash equilibrium a starting point of the analysis of almost all games.

Let us develop the concept by considering a more general version of the pricing game between Rainbow’s End and B. B. Lean. In chapter 3 we allowed each company the choice of just two prices for the shirt, namely $80 and $70. We also recognized the strength of the temptation for each to cut the price. Let us therefore allow more choices in a lower range, going in $1 steps from $42 to $38.† In the earlier example, when both charge $80, each sells 1,200 shirts. If one of them cuts its price by $1 while the other holds its price unchanged, then the price cutter gains 100 customers, 80 of whom shift from the other firm and 20 of whom shift from some other firm that is not a part of this game or decide to buy a shirt when they would otherwise not have done so. If both firms reduce their price by $1, existing customers stay put, but each gains 20 new ones. So when both firms charge $42 instead of $80, each gains 38 × 20 = 760 customers above the original 1,200. Then each sells 1,960 shirts and makes a profit of (42 × 20) × 1,960 = 43,120 dollars. Doing similar calculations for the other price combinations, we have the game table below.

		B. B. Lean's price				
		42	41	40	39	38
Rainbow's End's price	42	43,120 43,120	*43,260* 41,360	43,200 39,600	42,940 37,840	42,480 36,080
	41	41,360 *43,260*	41,580 41,580	*41,600* 39,900	41,420 38,220	41,040 36,540
	40	39,600 43,200	39,900 *41,600*	*40,000* *40,000*	39,900 *38,400*	39,600 36,800
	39	37,840 42,940	38,220 41,420	*38,400* 39,900	38,380 38,380	38,160 *36,860*
	38	36,080 42,480	36,540 41,040	36,800 39,600	*36,860* 38,160	36,700 36,700

The table may seem daunting but is in fact easy to construct using Microsoft Excel or any other spreadsheet program.

TRIP TO THE GYM NO. 2

Try your hand at constructing this table in Excel.

Best Responses

Consider the thinking of RE's executives in charge of setting prices. (From now on, we will simply say "RE's thinking," and similarly for BB.) If RE believes that BB is choosing $42, then RE's profits from choosing various possible prices are given by the numbers in the southwest corners of the first column of profits in the above table. Of those five numbers, the highest is $43,260, corresponding to RE's price $41. Therefore this is RE's "best response" to BB's choice of $42. Similarly, RE's best response is $40 if it believes that BB is choosing $41, $40, or $39, and $39 if it believes BB is choosing $38. We show these best-response profit numbers in bold italics for clarity. We also show BB's best responses to the various possible prices of RE, using bold, italicized numbers in the northeast corners of the appropriate cells.

Before proceeding, we must make two remarks about best responses. First, the term itself requires clarification. The two firms' choices are simultaneous. Therefore, unlike the situation in chapter 2, each firm is not observing the other's choice and then "responding" with its own best choice given the other firm's actual choice. Rather, each firm is formulating a belief (which may be based on thinking or experience or educated guesswork) about what the other firm is choosing, and responding to this belief.

Second, note that it is not always best for one firm to undercut the other's price. If RE believes that BB is choosing $42, RE should choose a lower price, namely $41; but if RE believes that BB is choosing $39, RE's best response is higher, namely $40. In choosing its best price, RE has to balance two opposing considerations: undercutting will increase the quantity it sells, but will leave it a lower profit margin per unit sold. If RE believes that BB is setting a very low price, then the reduction in RE's profit margin from undercutting BB may be too big, and RE's best choice may be to accept a lower sales volume to get a higher profit margin on each shirt. In the extreme case where RE thinks BB is pricing at cost, namely $20, matching this price will yield RE zero profit. RE does better to choose a higher price, keeping some loyal customers and extracting some profit from them.

Nash Equilibrium

Now return to the table and inspect the best responses. One fact immediately stands out: one cell, namely the one where each firm charges $40, has both of its numbers in bold italics, yielding a profit of $40,000 to each firm. If RE believes that BB is choosing the price of $40, then its own best price is $40, and vice versa. If the two firms choose to price their shirts at $40 each, the beliefs of each about the other's price are confirmed by the actual outcome. Then there would be no reason for one firm to change its price if the truth about the other firm's choice were somehow revealed. Therefore these choices constitute a stable configuration in the game.

Such an outcome in a game, where the action of each player is best for him given his beliefs about the other's action, and the action of each is consistent with the other's beliefs about it, neatly squares the circle of thinking about thinking. Therefore it has a good claim to be called a resting point of the players' thought processes, or an equilibrium of the game. Indeed, this is just a definition of Nash equilibrium.

To highlight the Nash equilibrium, we shade its cell in gray and will do the same in all the game tables that follow.

The price-setting game in chapter 3, with just two price choices of $80 and $70, was a prisoners' dilemma. The more general game with several price choices shares this feature. If both firms could make a credible, enforceable agreement to collude, they could both charge prices considerably higher than the Nash equilibrium price of $40, and this would yield larger profits to both. As we saw in chapter 3, a common price of $80 gives each of them $72,000, as opposed to only $40,000 in the Nash equilibrium. The result should impress upon you how consumers can suffer if an industry is a monopoly or a producers' cartel.

In the above example, the two firms were symmetrically situated in all relevant matters of costs and in the quantity sold for each combination of own and rival prices. In general this need not be so, and in the resulting Nash equilibrium the two firms' prices can be different. For those of you who want to acquire a better grasp of the methods and the concepts, we offer this as an "exercise"; casual readers should feel free to peek at the answer in the workouts.

TRIP TO THE GYM NO. 3

Suppose Rainbow's End locates a cheaper source for its shirts, so its cost per shirt goes down from $20 to $11.60, while B. B. Lean's cost remains at $20. Recalculate the payoff table and find the new Nash equilibrium.

The pricing game has many other features, but they are more complex than the material so far. Therefore we postpone them to a position later in this chapter. To conclude this section, we make a few general remarks about Nash equilibria.

Does every game have a Nash equilibrium? The answer is essentially yes, provided we generalize the concept of actions or strategies to allow mixing of moves. This was Nash's famous theorem. We will develop the idea of mixing moves in the next chapter. Games that have no Nash equilibrium, even when mixing is allowed, are so complex or esoteric that we can safely leave them to very advanced treatments of game theory.

Is Nash equilibrium a good solution for simultaneous-move games? We will discuss some arguments and evidence bearing on this issue later in this chapter, and our answer will be a guarded yes.

Does every game have a unique Nash equilibrium? No. In the rest of this chapter we will look at some important examples of games with multiple Nash equilibria and discuss the new issues they raise.

Which Equilibrium?

Let us try Nash's theory on the hunting game. Finding best responses in the hunting game is easy. Fred should simply make the same choice that he believes Barney is choosing. Here is the result.

		Barney's choice	
Fred's choice	Stag	Bison	Rabbit
Stag	*3*, *3*	0, 0	0, 1
Bison	0, 0	*3*, *3*	0, 1
Rabbit	1, 0	1, 0	*1*, *1*

So the game has three Nash equilibria.* Which of these will emerge as the outcome? Or will the two fail to reach any of the equilibria at all? The idea of Nash equilibrium does not by itself give the answers. Some additional and different consideration is needed.

If Fred and Barney had met at the stag party of a mutual friend, that might make the choice of Stag more prominent in their minds. If the ritual in their society is that as the head of the family sets out for the day's hunting he calls out in farewell, “Bye, son,” the choice of Bison might be prominent. But if the ritual is for the family to call out in farewell “Be safe,” the prominence might attach to the safer choice that guarantees some meat regardless of what the other chooses, namely rabbit hunting.

But what, precisely, constitutes “prominence”? One strategy, say Stag, may be prominent in Fred’s mind, but that is not enough for him to make that choice. He must ask himself whether the same strategy is also prominent for Barney. And that in turn involves asking whether Barney will think it prominent to Fred. Selecting among multiple Nash equilibria requires resolution of a similar problem of thinking about thinking as does the concept of Nash equilibrium itself.

To square the circle, the “prominence” must be a multilevel back-and-forth concept. For the equilibrium to be selected successfully when the two are thinking and acting in isolation, it must be obvious to Fred that it is obvious to Barney that it is obvious to Fred…that is the right choice. If an equilibrium is obvious ad infinitum in this way, that is, if the players’ expectations converge upon it, we call it a *focal point*. The development of this concept was just one of Thomas Schelling’s many pioneering contributions to game theory.

Whether a game has a focal point can depend on many circumstances, including most notably the players’ common experiences, which may be historical, cultural, linguistic, or purely accidental. Here are some examples.

We begin with one of Schelling’s classics. Suppose you are told to meet someone in New York City on a specific day but without being told where or at what time. You don’t even know who the other person is, so you cannot contact him/her in advance (but you are told how you would identify each other if and when you do meet). You are also told that the other person has been given identical instructions.

Your chances of success might seem slim; New York City is huge, and the day is long. But in fact people in this situation succeed surprisingly often. The time is simple: noon is the obvious focal point; expectations converge on it almost instinctively. The location is harder, but there are just a few landmark locations on which expectations can converge. This at least narrows down the choices considerably and improves the chances of a successful meeting.

Schelling conducted experiments in which the subjects were from the Boston or New Haven areas. In those days they traveled to New York by train and arrived at Grand Central Station; for them the clock in that station was focal. Nowadays, many people would think the Empire State Building is a focal point because of the movie *Sleepless in Seattle* (or *An Affair to Remember*); others would think Times Square the obvious “crossroads of the world.”

One of us (Nalebuff) performed this experiment in an ABC *Primetime* program titled *Life: The Game*.[1] Six pairs of mutual strangers were taken to different parts of New York and told to find others about whom they had no information except that the other pair would be looking for them under similar conditions. The discussions within each pair followed Schelling's reasoning remarkably well. Each thought about what they believed would be the obvious places to meet and about what others would think they were thinking: each team, say team A, in its thinking recognized the fact that another team, say B, was simultaneously thinking about what was obvious to A. Eventually, three of the pairs went to the Empire State Building and the other three to Times Square. All chose noon for the time. There remained some further issues to be sorted out: the Empire State Building has observation decks on two different levels, and Times Square is a big place. But with a little ingenuity, including a display of signs, all six pairs were successful in meeting.[*]

What is essential for success is not that the place is obvious to you, or obvious to the other team, but that it is obvious to each that it is obvious to the others that…And, if the Empire State Building has this property, then each team has to go there even though it may be inconvenient for them to get there, because it is the only place each can expect the other team to be. If there were just two teams, one of them might think the Empire State Building the obvious focal point and the other might think Times Square equally obvious; then the two would fail to meet.

Professor David Kreps of Stanford Business School conducted the following experiment in his class. Two students were chosen to play the game, and each had to make his/her choice without any possibility of communication with the other. Their job was to divide up a list of cities between them. One student was assigned Boston, and the other was assigned San Francisco (and these assignments were public so that each knew the other's city). Each was then given a list of nine other U.S. cities—Atlanta, Chicago, Dallas, Denver, Houston, Los Angeles, New York, Philadelphia, and Seattle—and asked to choose a subset of these cities. If their choices resulted in a complete and nonoverlapping division, both got a prize. But if their combined list missed a city or had any duplicates, then they both got nothing.

How many Nash equilibria does this game have? If the student assigned Boston chooses, say, Atlanta and Chicago, while the student assigned San Francisco chooses the rest (Dallas, Denver, Houston, Los Angeles, New York, Philadelphia, and Seattle), that is a Nash equilibrium: given the choice of one, any change in the choice of the other will create either an omission or an overlap and would lower the payoff to the deviator. The same argument applies if, say, one chooses Dallas, Los Angeles, and Seattle while the other chooses the other six. In other words, there are as many Nash equilibria as there are ways of dividing up the list of nine numbers into two distinct subsets. There are 2^9, or 512, such ways; therefore the game has a huge number of Nash equilibria.

Can the players' expectations converge to create a focal point? When both players were Americans or long-time U.S. residents, over 80 percent of the time they chose the division geographically; the student assigned Boston chose all the cities east of the Mississippi and the student assigned San Francisco chose those west of the Mississippi.[*] Such coordination was much less likely when one or both students were non-U.S. residents. Thus nationality or culture can help create a focal point. When Kreps's pairs lacked such common experience, choices were sometimes made alphabetically, but even then there was no clear dividing point. If the total number of cities was even, an equal split might be focal, but with nine cities, that is not possible. Thus one should not assume that players will always find a way to select one of multiple Nash equilibria by a convergence of expectations; failure to find a focal point is a distinct possibility.[†]

Next, suppose each of two players is asked to choose a positive integer. If both choose the same number, both get a prize. If the two choose different numbers, neither gets anything. The overwhelmingly frequent choice is 1: it is the first among the whole numbers (positive integers), it is the smallest, and so on; therefore, it is focal. Here the reason for its salience is basically mathematical.

Schelling gives the example of two or more people who have gone to a crowded place together and get separated. Where should each go in the expectation of finding the other? If the place, say a department store or a railway station, has a Lost and Found window, it has a good claim to be focal. Here the reason for its salience is linguistic. Sometimes meeting places are deliberately created to guarantee a convergence of expectations; for example, many railway stations in Germany and Switzerland have a well-signposted *Treffpunkt* (meeting point).

What is neat about the game of meeting is not just that the two players find each other but that the focal point ends up being relevant to so many strategic interactions. Probably the most important is the stock market. John Maynard Keynes, arguably the twentieth century's most famous economist, explained its behavior by analogy with a newspaper contest that was common in his time, where a number of photographs of faces were presented, and readers had to guess which face the majority of other voters would judge the most beautiful.[2] When everyone thinks along these lines, the question becomes which face most people think that most others will think that most others will think…is the most beautiful. If one contestant was significantly more beautiful than all the others, this could provide the necessary focal point. But the reader's job was rarely that easy. Imagine instead that the hundred finalists were practically indistinguishable except for the color of their hair. Of the hundred, only one is a redhead. Would you pick the redhead?

The aim becomes not to make any absolute judgment of beauty but to find a focal point of this process of thinking. How do we agree on that? The reader must figure out the realized convention without the benefit of communication. "Pick the most beautiful" might be the stated rule, but that could be significantly more difficult than picking the redhead, or the one with an interesting gap between her two front teeth (Lauren Hutton) or the mole (Cindy Crawford). Anything that distinguishes becomes a focal point and allows people's expectations to converge. For this reason, we should not be surprised that many of the world's top models do not have perfect features; rather, they are almost perfect but have some interesting flaw that gives their look a personality and a focal point.

Keynes used the beauty contest as a metaphor for the stock market, where each investor wants to buy the stocks that will rise in price, which means the stocks that investors, in general, think will appreciate. The hot stock is the one that everyone thinks that everyone else thinks…is the hot stock. There can be different reasons why different sectors or stocks become hot at different times—a well-publicized initial public offering, a famous analyst's recommendation, and so on. The focal point concept also explains the attention paid to round numbers: 10,000 for the Dow, or 2,500 for the Nasdaq. These indexes are just values of a specified portfolio of stocks. A number like 10,000 does not have any intrinsic meaning; it serves as a focal point only because expectations can converge more easily on round numbers.

The point of all this is that equilibrium can easily be determined by whim or fad. There is nothing fundamental that guarantees the most beautiful contestant will be chosen or the best stock will appreciate the fastest. There are some forces that work in the right direction. High forecast earnings are similar to the beauty contestant's complexion—one of the many necessary but by no means sufficient conditions needed to anchor otherwise arbitrary whims and fads.

Many mathematical game theorists dislike the dependence of an outcome on historical, cultural, or linguistic aspects of the game or on purely arbitrary devices like round numbers; they would prefer the solution be determined purely by the abstract mathematical facts about the game—the number of players, the strategies available to each, and the payoffs to each in relation to the strategy choices of all. We disagree. We think it entirely appropriate that the outcome of a game played by humans interacting in a society should depend on the social and psychological aspects of the game.

Think of the example of bargaining. Here the players' interests seem to be totally conflicting; a larger share for one means a smaller share for the other. But in many negotiations, if the two parties fail to agree, neither will get anything and both may suffer serious damage, as happens when wage bargaining breaks down and a strike or a lockout ensues. The two parties' interests are aligned to the extent that both want to avoid such disagreement. They can do so if they can find a focal point, with the common expectation that neither will concede anything beyond that point. That is why a 50:50 split is so often observed. It is simple and clear, it has the advantage of appearing fair, and, once such considerations get a foothold, it serves for the convergence of expectations.

Consider the problem of excessive compensation of CEOs. Often a CEO really cares about prestige. Whether the person gets paid $5 million or $10 million won't really have a big impact on the person's life. (That's easy for us to say from where we sit, where both numbers are quite abstract.) What's the meeting place that the CEOs care about? It is being better than average. Everyone wants to be in the top half. They all want to meet there. The problem is that this meeting spot only allows in half of the folks. But the way they get around this is via escalating pay. Every firm pays its CEO above last year's average, so everyone can think they have an above-average CEO. The end result is wildly escalating CEO salaries. To solve the problem, we need to find some other focal meeting point. For example, historically CEOs got prestige in their community via public service. Competing in that dimension was good all around. The current focal point on pay was created by *Business Week* surveys and compensation consultants. Changing it won't be easy.

The issue of fairness is also one of choosing a focal point. The Millennium Development Goals and Jeff Sachs's book *The End of Poverty* emphasize that contributing 1 percent of gross domestic product (GDP) to development will end poverty by 2025. The key point here is that the focal point of contributions is based on a percentage of income, not an absolute amount. Thus rich countries have a bigger obligation to contribute than the less rich. The apparent fairness of this can contribute to the convergence of expectations. Whether the promised funds will actually materialize remains to be seen.

BATTLES AND CHICKENS

In the hunting game, the two players' interests are perfectly aligned; both prefer one of the big-game equilibria, and the only question is how they can coordinate their beliefs on a focal point. We now turn to two other games, which also have non-unique Nash equilibria, but have an element of conflicting interests. Each leads to different ideas about strategy.

Both of these games date from the 1950s and have stories that fit those times. We will illustrate them using variants of the game between our Stone Age hunters, Fred and Barney. But we will relate the original sexist stories too, partly because they explain the names that have come to be attached to these games and partly for the amusement value of looking back on the quaint thoughts and norms of old times.

The first game is generically called battle of the sexes. The idea is that a husband and wife have different preferences in movies, and the two available choices are very different. The husband likes lots of action and fighting; he wants to see *300*. The wife likes three-handkerchief weepies; her choice is *Pride & Prejudice* (or *A Beautiful Mind*). But both prefer watching either movie in the other's company to watching any movie on their own.

In the hunting version, remove the Rabbit choice and keep only Stag and Bison. But suppose Fred prefers stag meat and rates the outcome of a jointly conducted stag hunt 4 instead of 3, while Barney has the opposite preference. The revised game payoff table is as shown below.

Fred's choice \ Barney's choice	Stag	Bison
Stag	***4***, ***3***	0, 0
Bison	0, 0	***3***, ***4***

As usual, best responses are shown in bold italics. We see at once that the game has two Nash equilibria, one where both choose Stag, and the other where both choose Bison. Both players prefer to have either equilibrium outcome than to hunt alone in one of the two nonequilibrium outcomes. But they have conflicting preferences over the two equilibria: Fred would rather be in the Stag equilibrium and Barney in the Bison equilibrium.

How might one or the other outcome be sustained? If Fred can somehow convey to Barney that he, Fred, is credibly and unyieldingly determined to choose Stag, then Barney must make the best of the situation by complying.

However, Fred faces two problems in using such a strategy.

First, it requires some method of communication before the actual choices are made. Of course, communication is usually a two-way process, so Barney might try the same strategy. Fred would ideally like to have a device that will let him send messages but not receive them. But that is not without its own problems; how can Fred be sure that Barney has received and understood the message?

Second, and more important, is the problem of credibly conveying an unyielding determination. This can be faked, and Barney might put it to the test by defying Fred and choosing Bison, which would leave Fred with a pair of bad choices: give in and choose Bison, which leads to humiliation and destruction of reputation, or go ahead with the original choice of Stag, which means missing the opportunity of the joint hunt, getting zero meat, and ending up with a hungry family.

In chapter 7 we will examine some ways that Fred could make his determination credible and achieve his preferred outcome. But we will also examine some ways that Barney could undermine Fred’s commitment.

If they have two-way communication before the game is played, this is essentially a game of negotiation. The two prefer different outcomes, but both prefer some agreement to complete disagreement. If the game is repeated, they may be able to agree to a compromise—for example, alternate between the two grounds on alternate days. Even in a single play, they may be able to achieve a compromise in the sense of a statistical average by tossing a coin and choosing one equilibrium if it comes up heads and the other equilibrium if it comes up tails. We will devote an entire chapter to the important subject of negotiation.

The second classic game is called chicken. In the standard telling of this story, two teenagers drive toward each other on a straight road, and the first one to swerve to avoid a collision is the loser, or chicken. If both keep straight, however, they crash, and that is the worst outcome for both. To create a game of chicken out of the hunting situation, remove the Stag and Bison choices, but suppose there are two areas for rabbit hunting. One, located to the south, is large but sparse; both can go there and each will get 1 of meat. The other, located to the north, is plentiful but small. If just one hunter goes there, he can get 2 of meat. If both go there, they will merely interfere and start fighting with each other and get nothing. If one goes north and the other goes south, the one who goes north will enjoy his 2 of meat. The one going south will get his 1. But his and his family’s feeling of envy for the other who comes back at the end of the day with 2 will reduce his enjoyment, so we will give him a payoff of only 1/2 instead of 1. This yields the game payoff table shown below.

Fred's choice	Barney's choice: North	Barney's choice: South
North	0 0	***1/2*** ***2***
South	***2*** ***1/2***	1 1

As usual, best responses are shown in bold italics. We see at once that the game has two Nash equilibria, with one player going north and the other going south. The latter is then the chicken; he has made the best of a bad situation in responding to the other's choice of North.

Both games, the battle of the sexes and chicken, have a mixture of common and conflicting interests: in both, the two players agree in preferring an equilibrium outcome to a nonequilibrium outcome, but they disagree as to which equilibrium is better. This conflict is sharper in chicken, in the sense that if each player tries to achieve his preferred equilibrium, both end up in their worst outcome.

Methods for selecting one of the equilibria in chicken are similar to those in the battle of the sexes. One of the players, say Fred, may make a commitment to choosing his preferred strategy, namely going north. Once again, it is important to make this commitment credible and to ensure that the other player knows it. We will consider commitments and their credibility more fully in chapters 6 and 7.

There is also the possibility of compromise in chicken. In a repeated interaction Fred and Barney may agree to alternate between North and South; in a single play, they may use a coin toss or other randomizing method to decide who gets North.

Finally, chicken shows a general point about games: even though the players are perfectly symmetric as regards their strategies and payoffs, the Nash equilibria of the game can be asymmetric, that is, the players choose different actions.

A LITTLE HISTORY

In the course of developing examples in this chapter and the one before it, we have introduced several games that have become classics. The prisoners' dilemma, of course, everyone knows. But the game of the two Stone Age hunters trying to meet is almost equally well known. Jean-Jacques Rousseau introduced it in an almost identical setting—of course he did not have *Flintstones* characters to add color to the story.

The hunters' meeting game differs from the prisoners' dilemma because Fred's best response is to take the same action as Barney does (and vice versa), whereas in a prisoners' dilemma game Fred would have a dominant strategy (just one action—for example, Rabbit—would be his best choice regardless of what Barney does) and so would Barney. Another way to express the difference is to say that in the meeting game, Fred would go stag hunting if he had the assurance, whether by direct communication or because of the existence of a focal point, that Barney would also go stag hunting, and vice versa. For this reason, the game is often called the *assurance game.*

Rousseau did not put his idea in precise gametheoretic language, and his phrasing leaves his meaning open to different interpretations. In Maurice Cranston's translation, the large animal is a deer, and the statement of the problem is as follows: "If it was a matter of hunting a deer, everyone well realized that he must remain faithfully at his post; but if a hare happened to pass within the reach of one of them, we cannot doubt that he would have gone off in pursuit of it without scruple and, having caught his own prey, he would have cared very little about having caused his companions to lose theirs."[3] Of course if the others were going for the hare, then there would be no point in any one hunter's attempting the deer. So the statement seems to imply that each hunter's dominant strategy is to go after a hare, which makes the game a prisoners' dilemma. However, the game is more commonly interpreted as an assurance game, where each hunter prefers to join the stag hunt if all the others are doing likewise.

In the version of chicken made famous by the movie *Rebel Without a Cause*, two teenagers drive their cars in parallel toward a cliff; the one who first jumps out of his car is the chicken. The metaphor of this game was used for nuclear brinkmanship by Bertrand Russell and others. The game was discussed in detail by Thomas Schelling in his pioneering gametheoretic analysis of strategic moves, and we will pick this back up in chapter 6.

To the best of our knowledge, the battle of the sexes game does not have such roots in philosophy or popular culture. It appears in the book *Games and Decisions* by R. Duncan Luce and Howard Raiffa, an early classic on formal game theory.[4]

FINDING NASH EQUILIBRIA

How can we find Nash equilibrium for a game? In a table, the worst-case method is cell-by-cell inspection. If both of the pay-off entries in a cell are best responses, the strategies and payoffs for that cell constitute a Nash equilibrium. If the table is large, this procedure can get tedious. But God made computers precisely to rescue humans from the tedium of inspection and calculation. Software packages to find Nash equilibria are readily available.[5]

Sometimes there are shortcuts; we now describe one that is often useful.

Successive Elimination

Return to the pricing game between Rainbow's End and B. B. Lean. Here again is the table of payoffs:

Rainbow's End's price	B. B. Lean's price: 42	41	40	39	38
42	43,120 43,120	*43,260* 41,360	43,200 39,600	42,940 37,840	42,480 36,080
41	41,360 *43,260*	41,580 41,580	*41,600* 39,900	41,420 38,220	41,040 36,540
40	39,600 43,200	39,900 *41,600*	*40,000* *40,000*	39,900 *38,400*	39,600 36,800
39	37,840 42,940	38,220 41,420	*38,400* 39,900	38,380 38,380	38,160 *36,860*
38	36,080 42,480	36,540 41,040	36,800 39,600	*36,860* 38,160	36,700 36,700

RE does not know what price BB is choosing. But it can figure out what price or prices BB is not choosing: BB will never set its price at $42 or $38. There are two reasons (both of which apply in our example, but in other situations only one may apply).[6]

First, each of these strategies is uniformly worse for BB than another available strategy. No matter what it thinks RE is choosing, $41 is better for BB than $42, and $39 is better than $38. To see this, consider the $41 versus $42 comparison; the other is similar. Look at the five numbers for BB's profits from choosing $41 (shaded in dark gray) versus those from $42 (shaded in light gray). For each of RE's five possible choices, BB's profit from choosing $42 is smaller than that from choosing $41:

43,120 < 43,260,
41,360 < 41,580,
39,600 < 39,900,
37,840 < 38,220,
36,080 < 36,540.

So no matter what BB expects RE to do, BB will never choose $42, and RE can confidently expect BB to rule out the $42 strategy, and, likewise, $38.

When one strategy, say A, is uniformly worse for a player than another, say B, we say that A is *dominated* by B. If such is the case, that player will never use A, although whether he uses B remains to be seen. The other player can confidently proceed in thinking on this basis; in particular, he need not consider playing a strategy that is the best response only to A. When solving the game, we can remove dominated strategies from consideration. This reduces the size of the game table and simplifies the analysis.*

The second avenue for elimination and simplification is to look for strategies that are *never best responses* to anything the other player might be choosing. In this example, $42 is never BB's best response to anything RE might be choosing within the range we are considering. So, RE can confidently think, "No matter what BB is thinking about my choice, it will never choose $42."

Of course, anything that is dominated is a never best response. It is more instructive to look at BB's option to price at $39. This can *almost* be eliminated for being a never best response. A price of $39 is only a best response to an RE price of $38. Once we know that $38 is dominated, then we can conclude that a BB price of $39 will never be a best response to anything RE will ever play. The advantage, then, of looking for never best responses is that you are able to eliminate strategies that are not dominated but would still never be chosen.

We can perform a similar analysis for the other player. RE's $42 and $38 strategies are eliminated, leaving us with a 3-by-3 game table:

		B. B. Lean's price		
		41	40	39
Rainbow's End's price	41	41,580 41,580	***41,600*** 39,900	41,420 38,220
	40	39,900 ***41,600***	***40,000*** ***40,000***	39,900 ***38,400***
	39	38,220 41,420	***38,400*** 39,900	38,380 38,380

In this simplified game, each firm has a dominant strategy, namely $40. Therefore our Rule 2 (from chapter 3) indicates that as a solution for the game.

The $40 strategy was not dominant in the original larger game; for example, if RE thought that BB would charge $42, then its profits from setting its own price at $41, namely $43,260, would be more than its profits from choosing $40, namely $43,200. The elimination of some strategies can open up the way to eliminate more in a second round. Here just two rounds sufficed to pin down the outcome. In other examples it may take more rounds, and even then the range of outcomes may be narrowed somewhat but not all the way to uniqueness.

If successive elimination of dominated strategies (or never-best-response strategies) and choice of dominant strategies does lead to a unique outcome, that is a Nash equilibrium. When this works, it is an easy way to find Nash equilibria. Therefore we summarize our discussion of finding Nash equilibria into two rules:

RULE 3: Eliminate from consideration any dominated strategies and strategies that are never best responses, and go on doing so successively.

RULE 4: Having exhausted the simple avenues of looking for dominant strategies or ruling out dominated ones, next search all the cells of the game table for a pair of mutual best responses in the same cell, which is a Nash equilibrium of the game.

GAMES WITH INFINITELY MANY STRATEGIES

In each of the versions of the pricing game we discussed so far, we allowed each firm only a small number of price points: only $80 and $70 in chapter 3, and only between $42 and $38 in $1 steps in this chapter. Our purpose was only to convey the concepts of the prisoners' dilemma and Nash equilibrium in the simplest possible context. In reality, prices can be any number of dollars and cents, and for all intents and purposes it is as if they can be chosen over a continuous range of numbers.

Our theory can cope with this further extension quite easily, using nothing more than basic high-school algebra and geometry. We can show the prices of the two firms in a two-dimensional graph, measuring RE's price along the horizontal or X axis and BB's price along the vertical or Y axis. We can show the best responses in this graph instead of showing bold italic profit numbers in a game table of discrete price points.

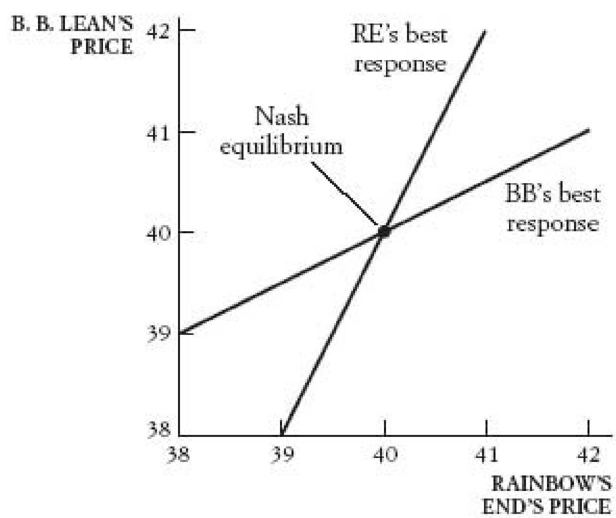

We do this for the original example where the cost of each shirt to each store was $20. We omit the details of the mathematics and merely tell you the result.[7] The formula for BB's best response in terms of RE's price (or BB's belief about the price RE is setting) is

BB's best response price = 24 + 0.4 × RE's price (or BB's belief about it).

This is shown as the flatter of the two lines in the above graph. We see that for each $1 cut in RE's price, BB's best response should be to cut its own price but by less, namely 40 cents. This is the result of BB's calculation, striking the best balance between losing customers to RE and accepting a lower profit margin.

The steeper of the two curves in the figure is RE's best response to its belief about BB's price. Where the two curves intersect, the best response of each is consistent with the other's beliefs; we have a Nash equilibrium. The figure shows that this occurs when each firm charges $40. Moreover, it shows that this particular game has exactly one Nash equilibrium. Our finding a unique Nash equilibrium in the table where prices had to be multiples of $1 was not an artificial consequence of that restriction.

Such graphs or tables that allow much more detail than we could in the simple examples are a standard method for computing Nash equilibria. The calculation or graphing can quickly get too complicated for paper-and-pencil methods, and too boring besides, but that's what computers are for. The simple examples give us a basic understanding of the concept, and we should reserve our human thinking skills for the higher-level activity of assessing its usefulness. Indeed, that is our very next topic.

A BEAUTIFUL EQUILIBRIUM?

John Nash's equilibrium has a lot of conceptual claim to be the solution of a game where each player has the freedom of choice. Perhaps the strongest argument in its favor takes the form of a counterargument to any other proposed solution. A Nash equilibrium is a configuration of strategies where each player's choice is his best response to the other player's choice (or the other players' choices when there are more than two players in the game). If some outcome is not a Nash equilibrium, at least one player must be choosing an action that is not his best response. Such a player has a clear incentive to deviate from that action, which would destroy the proposed solution.

If there are multiple Nash equilibria, we do need some additional method for figuring out which one will emerge as the outcome. But that just says we need Nash plus something else; it does not contradict Nash.

So we have a beautiful theory. But does it work in practice? One answers this question by looking for instances where such games are played in the real world, or by creating them in a laboratory setting and then comparing the actual outcomes against the predictions of the theory. If the agreement is sufficiently good, that supports the theory; if not, the theory should be rejected. Simple, right? In fact the process turns complicated very quickly, both in implementation and in interpretation. The results are mixed, with some reasons for optimism for the theory but also some ways in which the theory must be augmented or altered.

The two methods—observation and experiment—have different merits and flaws. Laboratory experiments allow proper scientific “control.” The experimenters can specify the rules of the game and the objectives of the participants quite precisely. For example, in pricing games where the subjects play the roles of the managers of the firms, we can specify the costs of the two firms and the equations for the quantities each would sell in relation to the prices both charge, and give the players the appropriate motivation by paying them in proportion to the profits they achieve for their firm in the game. We can study the effects of a particular factor, keeping all other things constant. By contrast, games that occur in real life have too many other things going on that we cannot control and too many things about the players—their true motivations, the firms’ costs of production, and so on—that we do not know. That makes it hard to make inferences about the underlying conditions and causes by observing the outcomes.

On the other hand, real-world observations do have some advantages. They lack the artificiality of laboratory experiments, in which the subjects are usually students, who have no previous experience in business or the similar applications that motivate the games. Many are novices even to the setting of the laboratory where the games are staged. They have to understand the rules of the game and then play it, all in a matter of an hour or two. Think how long it took you to figure out how to play even simple board games or computer games; that will tell you how naïve the play in such settings can be. We already discussed some examples of this problem in chapter 2. A second issue concerns incentives. Although the experimenter can give the students the correct incentives by designing the structure of their monetary payments to fit their performance in the game, the sizes of the payments are usually small, and even college students may not take them sufficiently seriously. By contrast, business games and even professional sports in the real world are played by experienced players for large stakes.

For these reasons, one should not rely solely on any one form of evidence, whether it supports or rejects a theory, but should use both kinds and learn from each. With these cautions in mind, let us see how the two types of empirical approaches do.

The field of industrial organization in economics provides the largest body of empirical testing of gametheoretic competition among firms. Industries like auto manufacturing have been studied in depth. These empirical investigators start with several handicaps. They do not know the firms' costs and demands from any independent source, and must estimate these things from the same data that they want to use for testing the pricing equilibrium. They do not know precisely how the quantities sold by each firm depend on the prices charged by all. In the examples in this chapter, we simply assumed a linear relationship, but the real-world counterparts (demand functions, in the jargon of economics) can be nonlinear in quite complicated ways. The investigator must assume some specific form of the nonlinearity. Real-life competition among firms is not just about prices; it has many other dimensions—advertising, investment, research and development. Real-life managers may not have the pure and simple aims of profit (or shareholder value) maximization that economic theory usually assumes. And competition among firms in real life extends over several years, so an appropriate combination of backward reasoning and Nash equilibrium concepts must be specified. And many other conditions, such as income and costs, change from one year to the next, and firms enter or exit the industry. The investigator must think about what all these other things might be and make proper allowance for (control for, in statistical jargon) their effects on quantities and prices. Real-world outcomes are also affected by many random factors and so, uncertainty must be allowed for.

A researcher must make a choice in each of these matters and then derive equations that capture and quantify all the relevant effects. These equations are then fitted to the data, and statistical tests performed to see how well they do. Then comes an equally difficult problem: What does one conclude from the findings? For example, suppose the data do not fit your equations very well. Something in your specification that led to the equations was not correct, but what was it? It could be the nonlinear form of the equations you chose; it could be the exclusion of some relevant variable, like income, or of some relevant dimension of competition, like advertising; or it could be that the Nash equilibrium concept used in your derivations is invalid. Or, it could be a combination of all these things. You cannot conclude that Nash equilibrium is incorrect when something else might be wrong. (But you would be right to raise your level of doubt about the equilibrium concept.)

Different researchers have made different choices in all these matters and, predictably, have found different results. After a thorough survey of this research, Peter Reiss and Frank Wolak of Stanford University give a mixed verdict: "The bad news is that the underlying economics can make the empirical models extremely complex. The good news is that the attempts so far have begun to define the issues that need to be addressed."[8] In other words, more research is needed.

Another active area for empirical estimation concerns auctions where a small number of strategically aware firms interact in bidding for things like bandwidths in the airwave spectrum. In these auctions, asymmetry of information is a key issue for the bidders and also for the auctioneer. Therefore we postpone the discussion of auctions to chapter 10, after we have examined the general issues of information in games in chapter 8. Here we merely mention that empirical estimation of auction games is already having considerable success.[9]

What do laboratory experiments have to say about the predictive power of game theory? Here the record is also mixed. Among the earliest experiments were the markets set up by Vernon Smith. He found surprisingly good results for game theory as well as for economic theory: small numbers of traders, each with no direct knowledge of the others' costs or values, could achieve equilibrium exchanges very quickly.

Other experiments with different kinds of games yielded outcomes that seemed contradictory to theoretical predictions. For example, in the ultimatum game, where one player makes a take-it-or-leave-it offer to the other for dividing a given sum between the two, the offers were surprisingly generous. And in prisoners' dilemmas, good behavior occurred far more frequently than theory might lead people to believe. We discussed some of these findings in chapters 2 and 3. Our general conclusion was that the participants in these games had different preferences or valuations than the purely selfish ones that used to be the natural assumption in economics. This is an interesting and important finding on its own; however, once the realistic "social" or "other-regarding" preferences are allowed for, the theoretical concepts of equilibrium—backward reasoning in sequential-move games and Nash in simultaneous-move games—yield generally good explanations of the observed outcomes.

When a game does not have a unique Nash equilibrium, the players have the additional problem of locating a focal point or some other method of selection among the possible equilibria. How well they succeed depends on the context, in just the way that theory suggests. If the players have sufficiently common understanding for their expectations to converge, they will succeed in settling on a good outcome; otherwise disequilibrium may persist.

Most experiments work with subjects who have no prior experience playing the particular game. The behavior of these novices does not initially conform to equilibrium theory, but it often converges to equilibrium as they gain experience. But some uncertainty about what the other player will do persists, and a good concept of equilibrium should allow players to recognize such uncertainty and respond to it. One such extension of the Nash equilibrium concept has become increasingly popular; this is the *quantal response equilibrium*, developed by professors Richard McKelvey and Thomas Palfrey of Caltech. This is too technical for a book like ours, but some readers may be inspired to read and study it.[10]

After a detailed review of the relevant work, two of the top researchers in the field of experimental economics, Charles Holt of the University of Virginia and Alvin Roth of Harvard University, offer a guardedly optimistic prognosis: “In the last 20 years, the notion of Nash equilibrium has become a required part of the tool kit for economists and other social and behavioral scientists…. There have been modifications, generalizations, and refinements, but the basic equilibrium analysis is the place to begin (and sometimes end) the analysis of strategic interactions.”[11] We think that to be exactly the right attitude and recommend this approach to our readers. When studying or playing a game, begin with the Nash equilibrium, and then think of reasons why, and the manner in which, the outcome may differ from the Nash predictions. This dual approach is more likely to give you a good understanding or success in actual play than either a totally nihilistic—anything goes—attitude or a slavishly naïve adherence to the Nash equilibrium with additional assumptions, such as selfishness.

CASE STUDY: HALF WAY

A Nash equilibrium is a combination of two conditions:

- i. Each player is choosing a best response to what he believes the other players will do in the game.
- ii. Each player’s beliefs are correct. The other players are doing just what everyone else thinks they are doing.

It is easier to describe this outcome in a two-player game. Our two players, Abe and Bea, each have beliefs about what the other will do. Based on those beliefs, Abe and Bea each choose to take an action that maximizes their payoffs. The beliefs prove right: Abe’s best response to what he thinks Bea is doing is just what Bea thought Abe would do, and Bea’s best response to what she thought Abe would do is indeed just what Abe expected her to do.

Let’s look at these two conditions separately. The first condition is quite natural. If otherwise, then you’d have to argue that someone is not taking the best action given what he or she believes. If he or she had something better, why not do it?

Mostly, the rub comes in the second condition—that everyone is correct in what they believe. For Sherlock Holmes and Professor Moriarty this was not a problem:

> “‘All that I have to say has already crossed your mind,’ said he.
> ‘Then possibly my answer has crossed yours,’ I replied.
> ‘You stand fast?’
> ‘Absolutely.’”

For the rest of us, correctly anticipating what the other side will do is often a challenge.

The following simple game will help illustrate the interplay between these two conditions and why you might or might not want to accept them.

Abe and Bea are playing a game with the following rules: Each player is to pick a number between 0 and 100, inclusive. There is a $100 prize to the player whose number is closest to half the other person’s number.

We’ll be Abe and you can play Bea. Any questions?

> What if there’s a tie?

Okay, in that case we split the prize. Any other questions?

> No.

Great, then let’s play. We’ve picked our number. Time for you to pick yours. What is your number? To help keep yourself honest, write it down.

Case Discussion

We picked 50. No, we didn’t. To see what we actually picked, you’ll have to read on.

Let’s start by taking a step back and use the two-step approach to finding a Nash equilibrium. In step 1, we believe that your strategy had to be an optimal response to something we might have done. Since our number has to be something between 0 and 100, we figure that you couldn’t have picked any number bigger than 50. For example, the number 60 is only an optimal response if you thought we would pick 120, something we couldn’t do under the rules.

What that tells us is that if your choice was truly a best response to something we might have done, you had to pick a number between 0 and 50. By the same token, if we picked a number based on something that you might have done, we would have picked something between 0 and 50.

Believe it or not, many folks stop right there. When this game is played among people who haven’t read this book, the most common response is 50. Frankly, we think that is a pretty lame answer (with apologies if that’s what you picked). Remember that 50 is only the best choice if you think that the other side was going to pick 100. But, in order for the other side to pick 100, they would have to have misunderstood the game. They would have had to pick a number that had (almost) no chance of winning. Any number less than 100 will beat 100.

We will assume that your strategy was a best response to something we might have done and so it is between 0 and 50. That means our best choice should be something between 0 and 25.

Note that at this juncture, we have taken a critical step. It may seem so natural that you didn’t even notice. We are no longer relying on our first condition that our strategy is a best response. We have taken the next step and proposed that our strategy should be a best response to something that is a best response from you.

If you are going to do something that is a best response, we should be doing something that is a best response to a best response.

At this point, we are beginning to form some beliefs about your actions. Instead of imagining that you can do anything allowed by the rules, we are going to assume that you will actually have picked a move that is a best response. Given the quite sensible belief that you are not going to do something that doesn’t make sense, it then follows that we should only pick a number between 0 and 25.

Of course, by the same token, you should be realizing that we won’t be picking a number bigger than 50. If you think that way, then you won’t pick a number bigger than 25.

As you might have guessed, the experimental evidence shows that after 50, 25 is the most common guess in this game. Frankly, 25 is a much better guess than 50. At least it has a chance of winning if the other player was foolish enough to pick 50.

If we take the view that you are only going to pick a number between 0 and 25, then our best response is now limited to numbers between 0 and 12.5. In fact, 12.5 is our guess. We'll win if our guess is closer to half your number than your number is to half ours. That means we win if you picked anything higher than 12.5.

Did we win?

Why did we pick 12.5? We thought you would pick a number between 0 and 25, and that's because we thought you'd think we'd pick a number between 0 and 50. We could of course go on with our reasoning and conclude that you'd figure we'd pick a number between 0 and 25, leading you to choose something between 0 and 12.5. If you had thought that, then you'd be one step ahead of us and would have won. Our experience suggests that most people don't think more than two or three levels, at least on their first go-around.

Now that you've had some practice and better understand the game, you might want a rematch. That's fair. So write down your number again—we promise not to peek.

We are pretty confident that you expect us to pick something less than 12.5. That means you'll pick something less than 6.25. And if we think you'll pick something less than 6.25, we should pick a number less than 3.125.

Now if this were the first go-around, we might stop there. But we just explained that most folks stop after two levels of reasoning, and this time we expect that you are determined to beat us, so you'll engage in at least one more level of thinking ahead. If you expect us to pick 3.125, then you'll pick 1.5625, which leads us to think of 0.78125.

At this point, we are guessing that you can see where this is all heading. If you think we are going to pick a number between 0 and X, then you should pick something between 0 and X/2. And if we think you are going to pick something between 0 and X/2, then we should pick something between 0 and X/4.

The only way that we can both be right is if we both pick 0. That's what we've done. This is the Nash equilibrium. If you pick 0, we want to pick 0; if we pick 0, you want to pick 0. Thus if we both correctly anticipate what the other will do, we both do best picking 0, just what we expected the other to do.

We should have picked 0 the first time around as well. If you pick X and we pick 0, then we win. That is because 0 is closer to X/2 than X is to 0/2 = 0. We knew this all along but didn't want to give it away the first time we played.

As it turned out, we didn't actually need to know anything about what you might be doing to pick 0. But that is a highly unusual case and an artifact of having only two players in the game.

Let's modify the game to add more players. Now the person whose number is closest to half the average number wins. Under these rules, it is no longer the case that 0 always wins.* But it is still the case that the best responses converge to zero. In the first round of reasoning, all players will pick something between 0 and 50. (The average number picked can't be above 100, so half the average is bounded by 50.) In the second iteration of logic, if everyone thinks others will play a best response, then in response everyone should pick something between 0 and 25. In the third iteration of logic, they'll all pick something between 0 and 12.5.

How far people are able to go in this reasoning is a judgment call. Again, our experience suggests that most people stop at two or three levels of reasoning. The case of a Nash equilibrium requires that the players follow the logic all the way. Each player picks a best response to what he or she believes that the other players are doing. The logic of Nash equilibrium leads us to the conclusion that all players will pick 0. Everyone picking 0 is the only strategy where each of the players is choosing a best response to what they believe other players are doing and each is right about what they believe the others will be doing.

When people play this game, they rarely pick zero on the first go-around. This is convincing evidence against the predictive power of Nash equilibrium. On the other hand, when they play the game even two or three times, they get very close to the Nash result. That is convincing evidence in favor of Nash.

Our view is that both perspectives are correct. To get to a Nash equilibrium, all players have to choose best responses—which is relatively straightforward. They also all have to have correct beliefs about what the other players will be doing in the game. This is much harder. It is theoretically possible to develop a set of internally consistent beliefs without playing the game, but it is often easier to play the game. To the extent that players learn that their beliefs were wrong by playing the game and then learn how to do a better job predicting what others will do, they will converge to a Nash equilibrium.

While experience is helpful, it is no guarantee of success. One problem arises when there are multiple Nash equilibria. Consider the annoying problem of what to do when a mobile phone call gets dropped. Should you wait for the other person to call you, or should you call? Waiting is a best response if you think the other person will call, and calling is a best response if you think the other person will wait. The problem here is that there are two equally attractive Nash equilibria: You call and the other person waits; or you wait and the other person calls.

Experience doesn't always help get you there. If you both wait, then you might decide to call, but if you both happen to call at the same time, then you get busy signals (or at least you did in the days before call waiting). To resolve this dilemma, we often turn to social conventions, such as having the person who first made the call do the callback. At least that way you know the person has the number.

EPILOGUE TO PART I

In the previous four chapters, we introduced several concepts and methods, using examples from business, sports, politics, and so forth as vehicles. In the chapters to follow, we will put the ideas and techniques to work. Here we recapitulate and summarize them for ready reference.

A *game* is a situation of strategic interdependence: the outcome of your choices (strategies) depends upon the choices of one or more other persons acting purposely. The decision makers involved in a game are called players, and their choices are called *moves*. The interests of the players in a game may be in strict conflict; one person's gain is always another's loss. Such games are called *zero-sum*. More typically, there are zones of commonality of interests as well as of conflict and so, there can be combinations of mutually gainful or mutually harmful strategies. Nevertheless, we usually refer to the other players in a game as one's rivals.

The moves in a game may be *sequential* or *simultaneous*. In a game of sequential moves, there is a linear chain of thinking: If I do this, my rival can do that, and in turn I can respond in the following way. Such a game is studied by drawing a *game tree*. The best choices of moves can be found by applying *Rule 1: Look forward and reason backward*.

In a game with simultaneous moves, there is a logical circle of reasoning: I think that he thinks that I think that…and so on. This circle must be squared; one must see through the rival's action even though one cannot see it when making one's own move. To tackle such a game, construct a *table* that shows the outcomes corresponding to all conceivable combinations of choices. Then proceed in the following steps.

Begin by seeing if either side has a *dominant strategy*—one that outperforms all of that side's other strategies, irrespective of the rival's choice. This leads to *Rule 2: If you have a dominant strategy, use it*. If you don't have a dominant strategy, but your rival does, then count on his using it, and choose your best response accordingly.

Next, if neither side has a dominant strategy, see if either has a *dominated strategy*—one that is uniformly worse for the side playing it than all the rest of its strategies. If so, apply *Rule 3: Eliminate dominated strategies from* *consideration*. Go on doing so successively. If during the process any dominant strategies emerge in the smaller games, they should be chosen. If this procedure ends in a unique solution, you have found the prescriptions of action for the players and the outcome of the game. Even if the procedure does not lead to a unique outcome, it can reduce the size of the game to a more manageable level. Finally, if there are neither dominant nor dominated strategies, or after the game has been simplified as far as possible using the second step, apply *Rule 4: Look for an equilibrium, a pair of strategies in which each player's action is the best response to the other's*. If there is a unique equilibrium of this kind, there are good arguments why all players should choose it. If there are many such equilibria, one needs a commonly understood rule or convention for choosing one over the others. If there is no such equilibrium, that usually means that any systematic behavior can be exploited by one's rivals, which indicates the need for *mixing one's plays*, the subject of the next chapter.

In practice, games can have some sequential moves and some simultaneous moves; in that case a combination of these techniques must be employed to think about and determine one's best choice of actions.