

How often have I said to you that when you have eliminated the impossible, whatever remains, however improbable, must be the truth?

—SHERLOCK HOLMES IN *A STUDY IN SCARLET*, BY SIR ARTHUR CONAN DOYLE

3.1 Introduction

IN CHAPTER 2, WE LEARNED how to construct a game, in both the extensive and strategic forms. But having built a game, what do we do with it? I say, let's play amateur sleuth and investigate people's behavior in strategic scenarios. To do so, we need to know how to solve a game, which is what this chapter is all about.

Figure 3.3 presents the strategic form of the kidnapping situation whose extensive form was illustrated in Figure 2.1. Recall that Guy, the kidnapper, has four strategies. He can either not kidnap Orlando, but in the event that he does, then kill him (*Do not kidnap/Kill*); not kidnap Orlando, but in the event that he does, then release him (*Do not kidnap/Release*); kidnap Orlando and kill him (*Kidnap/Kill*); or kidnap Orlando and release him (*Kidnap/Release*). Vivica, who is Orlando's kin, can either pay the ransom or not. The first number in a cell is the row player's payoff (Guy's payoff), and the second number is the column player's payoff (Vivica's payoff). Recall that payoff numbers are rank ordered from the least to the most preferred outcome.

Strategic Form of the Kidnapping Game		
	Vivica (kin of victim)	
	<i>Pay ransom</i>	<i>Do not pay ransom</i>
<i>Do not kidnap</i>	5, 5	5, 5
<i>Kidnap</i>	1, 2	3, 4

Will there be a kidnapping? If so, will ransom be paid? Will Orlando survive? Solving this game means answering these questions by selecting among the eight possible pairs of strategies. We need to weed out unreasonable and implausible strategy profiles and, ideally, identify a unique compelling one. The fewer the solutions, the more precise is our prediction about behavior. To derive a solution, we'll need to assume something about how a player selects among his or her

various strategies. Of course, what makes this task challenging is the fact that how a player selects a strategy may well depend on how she thinks *other* players are selecting. This is a complex undertaking not quickly dispensed with, and it is best that we start at the beginning and start simple.

The plan is to progressively make more assumptions about players and explore what we can say about how they'll behave. We begin with assuming that players are rational (Section 3.2), then further assume that each player believes that all players are rational (Section 3.3), and then assume on top of that that each player believes that all players believe that all players are rational (Section 3.3). We then generalize this sequence of solution techniques (Section 3.4) and conclude with some experimental evidence regarding the strategic sophistication of college undergraduates (Section 3.5).

3.2 Solving a Game When Players Are Rational

Stupid is a condition. Ignorance is a choice. —WILEY MILLER

IN MODELING A PLAYER'S SELECTION of a strategy, we'll begin by assuming that *players are rational*. A player is **rational** when she acts in her own best interests. More specifically, given a player's beliefs as to how other players will behave, the player selects a strategy in order to maximize her payoff. Note that rationality has nothing to say about what are reasonable beliefs to hold regarding what others will do; rationality just says that a player chooses the strategy that maximizes her payoff, *given* her beliefs as to the strategies of other players.

Is it reasonable to assume that people act only in their self-interest? Does this assumption mean that people are selfish? Although rationality does mean pursuing your own interests, it places few restrictions on what those interests might be. It can encompass Ebenezer Scrooge either before Christmas Eve—when all he cares about is money—or after he is visited by a series of ghosts—when he cares about his fellow human beings. Rationality is a church that welcomes all people, from the narcissistic to the altruistic. To be rational means only to pursue your interests, however they are defined.

Initially, we will assume a particular implication of rationality: A player will not use a strategy (call it s') when another strategy exists (call it s'') that always produces a strictly higher payoff, regardless of what strategies are used by the other players. Thus, strategy s' is *never* the right thing to do, as s'' will always outperform s' . It would then be rather stupid to play s' (whether you're Mother Teresa or Kim Kardashian). We'll go even further and assume that each player believes that other players avoid stupid strategies, that each player believes that other players believe that players avoid such strategies, and so forth. In other words, it is *common knowledge* among players that a player will not use a particular strategy when another strategy that is always strictly better is available.

In Appendix 3.7, we consider the concept of **rationalizability**, which is a stronger implication of rationality being common knowledge. Because it is a rather subtle and complex concept, we make it optional here and leave it for the more adventuresome student (or the more exacting instructor).

3.2.1 Strict Dominance

Let's revisit Puccini's opera *Tosca*, the strategic form game of which is reproduced in **FIGURE 3.2**. Recall from Section 2.5 that Baron Scarpia, the chief of police, has condemned Floria Tosca's lover Mario Cavaradossi to death. Scarpia

tells Tosca that he'll have the firing squad use blank cartridges in exchange for sexual favors from Tosca. Scarpia first tells the firing squad to use real or blank cartridges, and then he meets Tosca, at which point she must decide whether to consent to Scarpia's demands or stab him. Soon thereafter, Cavaradossi is brought before the firing squad. Both Tosca and Scarpia move without knowing what the other has chosen. The payoffs reflect that Tosca cares foremost that her lover survives and secondarily that she not consent to Scarpia, and that Scarpia longs to have relations with Tosca and only secondarily desires to execute Cavaradossi.

FIGURE 3.2 The Tosca Game

		Scarpia	
		<i>Real</i>	<i>Blanks</i>
Tosca	<i>Stab</i>	2,2	4,1
	<i>Consent</i>	1,4	3,3

Can we say what these two operatic characters will do if all we assume is that they are rational? Well, consider Tosca. If Scarpia chooses *real*, then Tosca's payoff from choosing *stab* is 2 and from choosing *consent* is 1. (See FIGURE 3.3.) Thus, she clearly prefers to stab Scarpia if she expects him to have ordered the firing squad to use real cartridges. What about if she expects him to have chosen blank cartridges? Her payoff is 4 from *stab* and 3 from *consent*, so, once again, she prefers *stab*. Thus, regardless of what Tosca believes Scarpia will do, *stab* gives

FIGURE 3.3 Tosca's Payoffs in the Event That Scarpia Chooses Real Bullets

		Scarpia	
		<i>Real</i>	<i>Blanks</i>
Tosca	<i>Stab</i>	2,2	4,1
	<i>Consent</i>	1,4	3,3

Tosca a strictly higher payoff than *consent*. Therefore, Tosca should most definitely *not* choose *consent*. Since she has to do something, she has nothing left to do but to stab Scarpia. To see this another way, *stab* is superior to *consent* regardless of what Scarpia will do, so a rational Tosca should most surely choose *stab*.

By a similar argument, the rationality of Scarpia implies that he will choose *real*. If Tosca chooses *stab*, he earns a payoff of 2 from killing Cavaradossi and only 1 from not doing so. If Tosca chooses *consent*, then the payoff from *real* is 4, which once again exceeds that from *blanks* (which is 3). Hence, regardless of what he thinks Tosca will do, Scarpia should have the firing squad use real bullets. In conclusion, game theory makes a very clear (and bloody) prediction that Tosca will stab Scarpia and Scarpia will see that Cavaradossi dies at the hands of the firing squad.

The strategy *consent* is said to be *strictly dominated* by the strategy *stab* for Tosca, which just means that *stab* delivers a higher payoff than *consent* for any strategy of Scarpia.

† **DEFINITION 3.1** A strategy s' **strictly dominates** a strategy s'' if the payoff from s' is strictly higher than that from s'' for any strategies chosen by the other players.*

A strategy that strictly dominates every other strategy for a player is said to be a **dominant strategy**. Obviously, with only two strategies, if *consent* is strictly dominated for Tosca, then *stab* must be the dominant strategy.

† **DEFINITION 3.2** A strategy is the **dominant strategy** if it strictly dominates every other strategy.

*A more formal mathematical presentation of Definitions 3.1 and 3.2 is provided in Section 3.6, which is an appendix to this chapter.

If a strategy is strictly dominated, then it is not optimal for any beliefs regarding what other players will do; thus, a rational player will avoid using such a strategy. Furthermore, if a player has a dominant strategy, then, if he is rational, he will use it. When each player has a dominant strategy, the unique reasonable solution is that each player uses his or her dominant strategy.

INSIGHT A rational player never uses a strictly dominated strategy. A rational player always uses a dominant strategy.

Before we apply these new tools to a few other games, take note of an interesting property of the outcome of *Tosca*. When Tosca stabs Scarpia and Scarpia has Cavaradossi killed, they each receive a payoff of 2. Now consider the alternative strategy pair in which Tosca consents and Scarpia has the firing squad use blanks, so that Cavaradossi survives. Now Tosca and Scarpia each earn a payoff of 3; they are both better off!

We see in the *Tosca* situation an important distinction between individual rationality and collective rationality: It is individually rational for Tosca to stab Scarpia (because it is her dominant strategy), and it is individually rational for Scarpia to use real bullets (because it is his dominant strategy); however, it is collectively rational—in the sense that everyone would be better off—if Tosca and Scarpia were to commit, respectively, to consenting and using blanks. Thus, what may be in an individual's best interests need not be in the best interests of the group. (We'll have more to say on this matter later in the book.)

So, what happens in the opera? Tosca stabs Scarpia and Scarpia uses real bullets, so that both Cavaradossi and Scarpia die. When she learns that Cavaradossi is dead, Tosca jumps to her death from the castle's ramparts. In spite of the carnage, love wins out over lechery. Or, if this was a tennis match, the score would be love: 15, lechery: love.

► SITUATION: **WHITE FLIGHT AND RACIAL SEGREGATION IN HOUSING**

What do you think you are going to gain by moving into a neighborhood where you just aren't wanted and where some elements—well—people can get awful worked up when they feel that their whole way of life and everything they've ever worked for is threatened. [The (white) Karl Linder speaking to the (black) younger family in response to Lena Younger's purchase of a home in his neighborhood.] —A RAISIN IN THE SUN, BY LORAIN HANSBERRY

The renting or sale of homes and apartments in an all-white neighborhood to African-Americans was a contentious racial issue in the 1960s. The term "white flight" refers to the exodus of white families upon the arrival of a few black families, making the neighborhood "tip" from being all white to all black. Could white flight occur even if both blacks and whites prefer to have a racially integrated neighborhood?¹

Suppose that a black family is willing to pay higher rent than a white family to live in what is currently an all-white neighborhood. A willingness to pay higher rent could reflect fewer options for blacks when it comes to attractive homes and good schools. For a scenario in which eight identical homes are in a certain neighborhood, TABLE 3.1 lists the monthly rent that a black family and a white family are hypothetically willing to pay, depending on how many black families

TABLE 3.1 Rent in the Housing Market

No. of Black Families	Rent Paid by a Black Family in Dollars	Rent Paid by a White Family in Dollars	Total Rent in Dollars
0	—	100	800
1	110	105	845
2	115	110	890
3	120	100	860
4	110	90	800
5	100	75	725
6	90	75	690
7	85	70	665
8	80	—	640

are in the neighborhood. Notice that both blacks and whites are willing to pay the highest rent for a (somewhat) racially integrated neighborhood.

The column titled "Total rent" is the sum of the monthly rents collected by the landlords of the eight homes. For example, if there are three black families, then each black family pays \$120/month and each of the five white families pays \$100/month, for a total collected monthly rent of \$860. Note that the landlords' rent is maximized at \$890, when two of the homes are rented to black families.

The game of interest is not between black and white families, but rather between the landlords who own the eight homes. Suppose that each home is owned by a different person and the payoff of a landlord equals the rent collected. Then, in deciding to whom to rent their property, will the landlords select strategies that result in a racially integrated neighborhood? Or will their decisions cause the neighborhood to "tip" from being all white to all black?

The concept of strict dominance allows us to say what will happen. Consider the decision faced by an individual landlord, and suppose the other seven landlords are currently renting to only white families. Then the first landlord will earn \$10 more a month (\$110 versus \$100) by renting to a black family. Hence, he will prefer to rent to a black family when the other seven homes are rented to white families. Now suppose instead that, of the other seven homes, six are rented to white families and one to a black family. Then if the landlord in question rents to a black family, he'll earn \$115 (since two black families are now in the neighborhood), an amount of money that exceeds what he can earn by renting to a white family, which is only \$105 (when there is only one black family). It is thus to the landlord's advantage to rent to a black family when six of the other seven houses are rented to white families. Continuing in this manner, you can show that, regardless of what the other seven landlords are doing (in terms of the race of their tenants), an individual landlord makes more money by renting to a black family. In other words, renting to a black family *strictly dominates* renting to a white family.

The monetary benefit applies to each of the landlords, so if each uses his or her dominant strategy, then each will rent to a black family. As a result, the neighborhood shifts from being all white to all black. Notice, however, that if

all the houses are rented to black families, then the landlords end up with a *lower* total rent of \$640, compared with the \$800 they got when they were all renting to white families. Sadly, then, landlords are poorer and racial integration is not achieved. ◀◀◀

3.1 CHECK YOUR UNDERSTANDING

Player 1 chooses a number from $\{5, 6, \dots, 20\}$, player 2 chooses a number from $\{10, 11, \dots, 30\}$, and player 3 chooses a number from $\{30, 31, \dots, 40\}$. Letting x_i denote the number (or strategy) chosen by player i , the payoff to player i is $100 - x_i$ when $x_1 + x_2 + x_3 \geq 50$, and is $-x_i$ when $x_1 + x_2 + x_3 < 50$. For each player, find the strategies that are strictly dominated.*

*Answers to Check Your Understanding are in the back of the book.

► SITUATION: BANNING CIGARETTE ADVERTISING ON TELEVISION

*Winston tastes good like a cigarette should!**

If you grew up in the 1960s, you would know that popular jingle from television commercials for Winston cigarettes. In fact, it was so common that many people (like me!) remember it even though it hasn't been broadcast for more than 40 years. Since a federally mandated ban went into effect in 1971, no advertisements for tobacco products have either been seen or heard on television or radio in the United States.

Typically, a government-imposed restriction is something that a company disdains. Taking away options usually means limiting avenues for achieving higher profit. In a game-theoretic setting, however, losing options is not always bad. Is it possible that the TV and radio ban might have *increased* the profits of the tobacco manufacturers?

Consider Philip Morris and R. J. Reynolds, which were (and still are) the largest cigarette manufacturers. Most people know their Marlboro and Winston brands of cigarettes. In considering how much advertising is needed, it is critical to understand how advertising affects the number of packs of cigarettes sold. Some studies by economists show that advertising doesn't have much of an impact on the number of smokers and instead just shifts the existing set of smokers among the different brands; however, other evidence indicates that advertising dissuades smokers from stopping and lures nonsmokers (in particular, youth) into trying smoking. To keep our model simple, let us assume that advertising doesn't affect the total number of packs sold and just shifts smokers among the different brands.

Suppose the annual demand for cigarettes is 1 billion packs, and that the market share of a company depends on how much it spends on advertising relative to what its rival spends. Let ADV_{PM} denote the advertising expenditures of Philip Morris (PM) and ADV_{RJR} denote the advertising expenditures of R. J. Reynolds (RJR). Assume that the market share of PM equals

$$\frac{ADV_{PM}}{ADV_{PM} + ADV_{RJR}}$$

*In response to complaints about grammar from language mavens, Winston responded with a new slogan: "What do you want, good grammar or good taste?"

This quotient says that PM's share of all packs sold equals its share of advertising. Such a model is overly simplistic, but all that we really need to assume is that sales are higher when more is spent on advertising. The total number of packs sold by PM is then

$$1,000,000,000 \times \left(\frac{ADV_{PM}}{ADV_{PM} + ADV_{RJR}} \right),$$

and by a similar argument, the corresponding number for RJR is

$$1,000,000,000 \times \left(\frac{ADV_{RJR}}{ADV_{PM} + ADV_{RJR}} \right).$$

If each pack sold generates a profit of 10 cents (remember, we're back in 1971), then the profit that PM gets from spending ADV_{PM} dollars is

$$0.1 \times 1,000,000,000 \times \left(\frac{ADV_{PM}}{ADV_{PM} + ADV_{RJR}} \right) - ADV_{PM}$$

or

$$100,000,000 \times \left(\frac{ADV_{PM}}{ADV_{PM} + ADV_{RJR}} \right) - ADV_{PM}$$

Analogously, RJR's profit is

$$100,000,000 \times \left(\frac{ADV_{RJR}}{ADV_{PM} + ADV_{RJR}} \right) - ADV_{RJR}$$

Our objective is to say something about how much these companies advertise. To keep things simple, assume just three levels of advertising exist: \$5 million, \$10 million, and \$15 million. In that case, the payoff matrix is as shown in FIGURE 3.4 (where strategies and payoffs are in millions of dollars). For example, if PM spends \$5 million and RJR spends \$15 million, then PM's market share is $5/(5 + 15)$, or 25%. PM then sells 250 million packs (0.25 multiplied by 1 billion) and, at 10 cents per pack, makes a gross profit of \$25 million. Once we net out the cost of advertising, PM's profit (or payoff) is \$20 million.

FIGURE 3.4 The Cigarette Advertising Game

		R. J. Reynolds		
		Spend 5	Spend 10	Spend 15
Philip Morris	Spend 5	45,45	28,57	20,60
	Spend 10	57,28	40,40	30,45
	Spend 15	60,20	45,30	35,35

TABLE 3.2 shows that a strategy of spending \$15 million strictly dominates spending either \$5 million or \$10 million. For example, if RJR spends \$5 million, then PM earns \$60 million from spending \$15 million (and gets 75% of the market), while PM earns \$57 million with an advertising budget of \$10 million and only \$45 million by matching RJR's paltry expenditure of \$5 million. Similarly, a budget of \$15 million for PM outperforms the other two options when RJR spends \$10 million and when it spends \$15 million. Thus, PM prefers the heavy

TABLE 3.2 Spending 15 Is a Dominant Strategy for PM					
RJR Strategy	PM Payoff from 15		PM Payoff from 10		PM Payoff from 5
Spend 5	60	>	57	>	45
Spend 10	45	>	40	>	28
Spend 15	35	>	30	>	20

advertising campaign regardless of what RJR chooses, so heavy advertising is a dominant strategy for PM. Because the same can be shown for RJR, the prediction is that both cigarette companies inundate our television sets with attractive men and women spewing forth smoke.

FIGURE 3.5 Cigarette Advertising Game When TV and Radio Commercials Are Excluded			
		R. J. Reynolds	
		Spend 5	Spend 10
Philip Morris	Spend 5	45,45	28,57
	Spend 10	57,28	40,40

Now suppose the ban on TV and radio advertising is put into effect and it has the impact of making it infeasible for the cigarette companies to spend \$15 million on advertising. That is, the most that a company can spend using the remaining advertising venues is \$10 million. In the context of this simple game, each company's strategy set is then constrained to comprise the choices of spending \$5 million and spending \$10 million, as shown in FIGURE 3.5.

The solution to this game is that both companies spend moderately on advertising, since spending \$10 million strictly dominates spending 5 million. And what has this intrusive government policy done to their profits? They have increased! Each company's profit rises from \$35 million to \$40 million.

In the original game, each company had a dominant strategy of spending \$15 million. This heavy advertising tended to cancel out, so each ended up with 50% of the market. If they both could have restrained their spending to \$10 million, they would each still have had half of the market—thus leaving them with the same gross profits—and would have spent less on advertising, which translates into higher net profit.

By reducing the options for advertising, the TV and radio ban served to restrain competition, reduce advertising expenditures, and raise company profits. Of course, there's nothing wrong with that if that is indeed what happened, since the objective of the ban was to reduce smoking, not lower companies' profits. ◀◀◀

◆ **INSIGHT** A rational player never smokes cigarettes. (Okay, I'm making that one up.)

3.2.2 Weak Dominance

Not to be absolutely certain is, I think, one of the essential things in rationality. —BERTRAND RUSSELL

Returning to the world of Italian opera yet again, suppose we now assume that Scarpia, upon being stabbed by Tosca, does not care whether Cavaradossi is killed.

The resulting payoff matrix is shown in FIGURE 3.6. Although *stab* continues to be the dominant strategy for Tosca (indeed, we haven't changed her payoffs), using real bullets no longer strictly dominates using blanks for Scarpia. Nevertheless, *real* would seem the reasonable course of action for Scarpia. If Tosca consents, then Scarpia strictly prefers to have used real bullets; he receives a payoff of 4 as opposed to 3. If Tosca stabs him, then he doesn't care, as his payoff is 2 regardless of what he chooses. Thus, he can't be any worse off by using real bullets, and he might just be better off. We say that the strategy *real* weakly dominates the strategy *blanks*.

FIGURE 3.6 Revised Tosca Game

		Scarpia	
		<i>Real</i>	<i>Blanks</i>
Tosca	<i>Stab</i>	2,2	4,2
	<i>Consent</i>	1,4	3,3

† **DEFINITION 3.3** A strategy s' **weakly dominates** a strategy s'' if (1) the payoff from s' is at least great as that from s'' for any strategies chosen by the other players; and (2) the payoff from s' is strictly greater than that from s'' for some strategies of the other players.*

Because most people are cautious and lack absolute confidence as to what other players will do, it seems prudent to avoid weakly dominated strategies. Doing so means that you can never be any worse off and you just might end up being better off. There'll be no regrets by avoiding weakly dominated strategies. If a **weakly dominant** strategy exists—which means that it weakly dominates all other strategies—it would be wise to use that strategy.

INSIGHT A rational and cautious player never uses a weakly dominated strategy. A rational and cautious player always uses a weakly dominant strategy.

3.2.3 Bidding at an Auction

It's a very sobering feeling to be up in space and realize that one's safety factor was determined by the lowest bidder on a government contract.

—ASTRONAUT ALAN SHEPHERD

It's been 20 years since you've graduated from college, and you've just sold your Internet company for a cool \$50 million. With all this cash on hand, you decide to indulge your passion for modern art. An auction house is selling an Andy Warhol piece that you've been coveting for some time. The rules are that all interested parties must submit a written bid by this Friday at 5 P.M. Whoever submits the highest bid wins the Warhol piece and pays a price equal to the bid—a format known as the **first-price auction**.

The Warhol piece is worth \$400,000 to you. If you win the item, your payoff equals \$400,000 less the price you paid, while if you don't win, your payoff is zero. Hence, if you end up paying \$400,000, you're no better off, while you're better (worse) off if you get it for less (more) than \$400,000.

You've just learned that there is only one other bidder: your old college girl friend, who has recently cashed in stock options after being CEO for a biomedical company. You know that she values the piece at \$300,000, and furthermore,

*A more formal mathematical definition is provided in Section 3.6, which is an appendix to this chapter.

is the second-highest bid. In that case, your payoff is 2 ($= 4 - 2$). In fact, your payoff is 2 whether you bid 3, 4, or 5, because the price you pay is not your bid, but the other bidder's bid. Your bid only influences whether or not you win. If you were to bid 1, then, since she is bidding 2, your payoff would be affected—it is now zero—since your low bid causes you to lose the auction.

Inspection of Figure 3.8 reveals that a bid of 4 weakly dominates every other bid for you. It would then make sense for you to bid 4, regardless of how you think your former girl friend will bid. As for her, a bid of 3 weakly dominates every other one of her bids. Note that for each of you, the weakly dominant bid equals your valuation. This is not coincidental: in every second-price auction, bidding your valuation weakly dominates every other bid!

In the first-price auction, the motivation for shading your bid below your valuation is to lower the price you pay in the event that you win. That strategy doesn't work in the second-price auction, since the price you pay is not what *you* bid, but what *someone else* bid. Bidding below your valuation only reduces your chances of winning at a price below your valuation, and that's a bad deal.

Figuring out your optimal bid at a second-price auction is a piece of cake. A bidder just needs to determine what the item is worth and bid that value, without the need for a certified psychologist to help you evaluate the psyche of other bidders or for the services of a well-trained game theorist to tell you how to bid! You just need to know yourself.

3.2 CHECK YOUR UNDERSTANDING

For the game in FIGURE 3.9, find the strategies that are strictly dominated and those that are weakly dominated.*

FIGURE 3.9

		Player 2			
		w	x	y	z
Player 1	a	3,2	1,1	4,3	3,5
	b	1,3	3,0	2,4	4,2
	c	2,1	0,1	1,2	1,0
	d	1,0	2,0	2,1	4,0

*Answers to Check Your Understanding are in the back of the book.

► SITUATION: THE PROXY BID PARADOX AT eBAY

Have you ever bid for an item at eBay? Since an auction typically takes days, eBay was smart enough to provide a mechanism that doesn't require a bidder to hang out online 24/7; instead, you can enter a proxy bid, which works as follows: As long as the highest bid of the other bidders is below your proxy bid, you'll be the top bidder, with a bid equal to the highest bid of the other bidders plus the minimum bid increment. As soon as the highest bid of the other bidders exceeds your proxy bid, you drop out of the bidding, although you can always return with a higher proxy bid.

So what should your proxy bid be? Note that if, at the end of the auction, you submitted the highest proxy bid, then you win the item and pay a price equal to the *second-highest* proxy bid (plus the minimum bid increment). In this way, the

eBay auction has the property of a second-price auction, and accordingly, you should submit a proxy bid equal to your valuation. Furthermore, once you've submitted such a bid, you can just return to the auction site at its completion to find out whether you've won. How simple!

This argument for setting your proxy bid equal to your valuation is hit with a full body slam when it gets in the ring with reality. Contrary to what the theory prescribes, people frequently *change* their proxy bid over the course of an eBay auction. For example, say Dave enters a proxy bid of \$150 for a ticket to a Rolling Stones concert and, coming back a day later, sees that the highest bid has reached \$180, so that it exceeds his proxy bid. Dave then changes his proxy bid to \$200. But if it was originally worth \$200 to Dave to see Mick Jagger and his buddies, why didn't he just submit a proxy bid of \$200 at the start? Why mess around with this lower proxy bid?

Because the phenomenon of bidders changing their proxy bids happens fairly often, we cannot summarily dismiss it as "stupid bidding." The phenomenon represents systematic behavior, and as social scientists, our objective is to understand it, not judge it. If we accept the idea that the bidders are doing exactly what they intend to do, it's the theory that's stupid—or, to say it more eloquently, our auction model is missing some relevant factors.

What could be missing? Several possibilities have been identified, but we have space to discuss only one. A potentially significant departure between the model and reality is that eBay runs multiple auctions for the same item. Think about how this could alter someone's bidding strategy. Perhaps the Stones concert is worth \$200 to you, which means that you would prefer to pay anything less than \$200 than not get a ticket. But if you're bidding for a ticket at an eBay auction, losing the auction doesn't necessarily mean not getting a ticket; you might instead participate in another auction for a Stones ticket.

To see what difference this new information makes, imagine that you're participating at an auction that ends two weeks prior to the concert. If you win the auction at a price of \$199, your payoff is then \$1, which is your valuation less the price. Although \$1 is higher than your payoff from not having a ticket, which is zero, it may not be higher than your expected payoff from participating in another auction. You might prefer not to win at \$199 in order to have the option of winning at a lower price in a later auction.

It isn't hard to see how this scenario could cause you to change your proxy bid over time. Suppose that you are currently watching two auctions and auction I ends tomorrow and auction II ends in two days. Suppose also that you have a proxy bid in auction II, but you're keeping track of the price in auction I. As just argued, your proxy bid is not your valuation, but instead something that depends on what kind of price you think you would need to pay to win at another auction. If auction I closes at a higher price than you expected, you may conclude that the remaining tickets will go for higher prices, and this conclusion could cause you to raise your proxy bid at auction II. Your optimal proxy bid changes over time as you learn what these tickets are selling for at other eBay auctions.

The gap between the theory's prediction and actual behavior at eBay auctions indicates a problem, not with game theory, but rather with the particular game-theoretic model. Game theory is immensely flexible and, when combined with an observant and clever mind, can offer cogent explanations of many social phenomena. ◀◀◀

3.3 CHECK YOUR UNDERSTANDING

Consider a two-player game in which each player's strategy set is $\{0, 1, \dots, 5\}$. If q_1 and q_2 denote the strategies of players 1 and 2, respectively, then player 1's payoff is $(10 - 2q_1 - q_2)q_1$ and player 2's payoff is $(10 - 2q_2 - q_1)q_2$. Note that this game is symmetric. Find all strictly dominated strategies and all weakly dominated strategies.*

*Answers to Check Your Understanding are in the back of the book.

3.3 Solving a Game When Players Are Rational and Players Know That Players Are Rational

[Saddam Hussein] starts out with a very menacing image. It sets you back a bit. I remember looking at my hands, and I was sweating. I was conscious that he knew what his reputation was. And he knew that I knew his reputation.

—BILL RICHARDSON

IN A SECOND-PRICE AUCTION, a player's optimal bid could be determined without figuring out what bids others would submit. However, in a first-price auction, how much a bid should be shaded below your valuation depends on how aggressively you think other bidders will bid. Games commonly require the kind of thinking that goes into bidding in a first-price auction, in that a player must prognosticate what others will do.

In this section, we begin our journey into solving that problem by considering some games for which it's not enough to assume players are rational. However, if we assume just a bit more—such as the assumption that each player knows that the other players are rational—then reasonable conclusions can be drawn about how players will behave, at least in some situations.

► SITUATION: TEAM-PROJECT GAME

Stanford is sort of a big, incredibly smart high school, the high school that we never had. We've got the jocks, the nerds, the sorority girls, the frat boys, the indie kids, the preps, the 'whatever' college kids. . . . —TAM VO IN *THE STANFORD DAILY*

Consider a college class with a diverse array of students, and let's indulge ourselves with a few stereotypes. Some of the students are underachieving jocks who, as long as it means minimal studying, are content to get a grade of C (fondly known as the "hook"—for looking like one—at my alma mater, the University of Virginia). Then there are the frat boys and sorority girls who are satisfied with a B, but are willing to work hard to avoid a lower grade and the dissatisfaction of their parents. And let us not forget the overachieving nerds who work hard to get an A and find that the best place for their noses is buried in books. (Is there anyone I have not offended?)

Determining how much effort a student will exert is fairly straightforward when it comes to an individual assignment such as an exam. The nerd will study hard; the frat boy will study moderately, and the jock will study just enough to pass. But what happens when they are thrown together in a team project? The quality of the project, and thereby the grade, depends on what all of the team members do. How much effort a student should exert may well

FIGURE 3.10 Team-Project Game with a Nerd and a Jock

		Nerd		
		Low	Moderate	High
Jock	Low	3,1	4,2	5,3
	Moderate	2,2	3,3	4,4
	High	1,3	2,4	3,5

depend on how hard other team members are expected to work.

To keep things simple, let's consider two-person team projects and initially examine a team made up of a nerd and a jock. The associated payoff matrix is shown in FIGURE 3.10. Each student has three levels of effort: *low*, *moderate*, and *high*. The grade on the project is presumed to increase as a function of the effort of both students. Hence, a student's payoff is always increasing with the effort of the other student, as an increasing effort by the other student means a better grade without having to work harder.

Jocks strongly dislike academic work, so their payoffs are ordered to reflect a distaste for effort. Regardless of the effort exerted by her nerdy partner (yes, there are female jocks!), the jock's payoff is lower when she works harder. For example, if the nerd exerts a moderate effort, then the jock's payoff falls from 4 to 3 to 2 as her effort goes from *low* to *moderate* to *high*. You can confirm that *low* is the jock's dominant strategy, since exerting a low effort yields a higher payoff than any other strategy, regardless of the effort chosen by her partner.

What about the nerd? The nerd's payoff increases with effort. Regardless of the effort of his partner, a nerd prefers to work harder in order to improve the project's grade. Thus, a high effort is the dominant strategy for the nerd. The outcome of the game in Figure 3.10 is then clear: If students are rational (and sober), then the jock will exert a low effort and the nerd will exert a high effort. The jock gets a payoff of 5—she does great because she's matched up with someone who is willing to work hard—and the nerd gets a payoff of 3 (while muttering "stupid lazy jock" under his breath).

Next, consider a frat boy and a nerd being matched up. The payoff matrix is presented in FIGURE 3.11. As before, the nerd's payoffs increase with effort. The frat boy is a bit more complicated than the nerd and the jock. He wants a reasonably good grade and is willing to work hard to get it if that is what is required, but he isn't willing to work hard just to go from a B to an A. The frat boy then lacks a dominant strategy. If his partner is lazy, then the frat boy is willing to work hard in order to get that B. If his partner "busts his buns," then the frat boy is content to do squat, as he'll still get the B. And if the partner exerts a moderate effort then the frat boy wants to do the same.

Simply knowing that the frat boy is rational doesn't tell us how he'll behave. Can we solve this game if we assume more than just that players are rational? Remember that the game is characterized by common knowledge: the frat boy knows that he's matched up with a nerd. (It's pretty apparent from the tape around the bridge of his glasses.) Suppose the frat boy not only is rational, but knows that his partner is rational. Since a rational player uses a dominant strategy when he has one, the frat boy can infer from his partner's being rational (and a nerd) that he will exert a high effort. Then, given that his partner exerts a high

FIGURE 3.11 Team-Project Game with a Nerd and a Frat Boy

		Nerd		
		Low	Moderate	High
Frat boy	Low	0,1	2,2	6,3
	Moderate	1,2	4,3	5,4
	High	2,3	3,4	3,5

effort, the effort. Thus, are matched and the fra order to deri to assume th are rational that the nerd

Finally, su up with his f girl. The payc Assuming th: that each play The trick that as neither pla have to wait

► SITUATION: EX

It is as im for Harry, A QUOTE

Philosophers long time. Or century by B tician who, a to the Lord. Pascal's wag of God. Ther penalty will Brady Bunch the cost to y safe and beli In other wor loses nothing

Pascal's w beliefs? Does based, not on is not a text criticism of P maker. Shoul God to decid and what will God then Man has the in God. In de (or disbelief) God's existen dence of God

effort, the frat boy should exert a low effort. Thus, when a nerd and a frat boy are matched, the nerd will hunker down and the frat boy will lounge about. In order to derive this conclusion, we needed to assume that the nerd and the frat boy are rational *and* that the frat boy knows that the nerd is rational.

Finally, suppose the frat boy is matched up with his female counterpart, the sorority girl. The payoff matrix is given in FIGURE 3.12. Assuming that the players are rational and that each player knows that the other is rational is not enough to solve this game. The trick that solves the game between the frat boy and the nerd won't work here, as neither player has a dominant strategy. Learning how to solve this situation will have to wait until Chapter 4. ◀◀◀

FIGURE 3.12 Team-Project Game with a Frat Boy and a Sorority Girl

		Sorority girl		
		Low	Moderate	High
Frat boy	Low	0,0	2,1	6,2
	Moderate	1,2	4,4	5,3
	High	2,6	3,5	3,3

► SITUATION: **EXISTENCE-OF-GOD GAME**

It is as impossible for man to demonstrate the existence of God as it would be for Harry Potter to demonstrate the existence of J. K. Rowling. —AN UPDATING OF A QUOTE BY FREDERICK BUECHNER

Philosophers have wrestled with the issue of whether God exists for a very long time. One of the most famous approaches was developed in the mid-17th century by Blaise Pascal (1623–1662). Pascal was a highly talented mathematician who, at a point in his life, threw aside mathematics to dedicate his life to the Lord. His take on the issue of belief in God has come to be known as Pascal's wager. It goes like this: Suppose you're not sure about the existence of God. Then if you fail to believe in God and it turns out God does exist, the penalty will be mighty severe. (Think of white-hot pitchforks and endless Brady Bunch reruns.) However, if you believe in God and God does not exist, the cost to you is rather minimal. Pascal then argues that one should play it safe and believe in God in order to avoid the excruciatingly horrible outcome. In other words, the atheist gains nothing by being right, and the Christian loses nothing by being wrong.

Pascal's wager has been critiqued many times. Can one really "choose" one's beliefs? Does God reward beliefs as opposed to actions? Should belief in God be based, not on faith or love, but on the cold, calculating logic of wagers? But this is not a text on philosophy or theology; rather, it is about game theory. So my criticism of Pascal's wager is that the problem, as cast, involves only one decision maker. Shouldn't we allow God to be a player? In particular, suppose we allow God to decide whether or not to reveal Her existence to Man. What will God do and what will Man do in that instance?²

God then has two strategies: *reveal* Herself to Man and *hide* Her existence. Man has the two strategies laid out by Pascal: *believe* in God and *do not believe* in God. In describing payoffs, suppose Man cares most about having his belief (or disbelief) confirmed. If he believes in God, he wants to see evidence of God's existence. If he doesn't believe in God, he surely doesn't want to see evidence of God. Secondly, Man prefers to believe in God's existence. As for

FIGURE 3.13 Existence-of-God Game

		Man	
		Believe	Do not believe
God	Reveal	3 4	1 1
	Hide	4 2	2 3

God, She cares most about Man believing in God and secondarily prefers not revealing Herself. The strategic form of the game is revealed (yes, pun intended) in **FIGURE 3.13**.

No dominant strategy exists for Man. If God intends to reveal Her existence, then Man wants to believe in God. If God does not intend to reveal Her existence, then Man doesn't want to believe in God. Knowing that Man is rational isn't enough to tell us what Man will do. In contrast, God does have a dominant strategy: regardless of Man's belief or disbelief in God, God prefers to hide Her existence. Doing so yields a payoff of 4 versus 3 for when Man believes in God and a payoff of 2 versus 1 when Man does not. A rational God will then hide Her existence.

If Man believes that God is rational, then Man knows that God will hide Her existence, since that is a dominant strategy for God. Given that God hides Her existence, Man's optimal strategy is not to believe in God. We conclude that the answer to the riddle is that Man should not believe in God and God should hide Her existence from Man.* ◀◀◀

3.4 CHECK YOUR UNDERSTANDING

For the game shown in 3.3 Check Your Understanding, find the strategies that are consistent with the players being rational and each player believing that the other player is rational.*

*Answers to Check Your Understanding are in the back of the book.

► SITUATION: BOXED-PIGS GAME

While this book aims to show how game theory can be used to understand human behavior, it can explain the behavior of lesser animals as well (which we explore more fully in Chapters 16 and 17). Let's consider an experiment in which two pigs—one large and one small—are placed in a cage. At one end of the cage is a lever and at the other end is a food dispenser. When the lever is pressed, 10 units of food are dispensed at the other end. Suppose either pig incurs a utility cost of 2 units (measured in food) from pressing the lever. How the 10 units of dispensed food is divided up depends on both who gets to the food dispenser first and a pig's size. If the large pig is there first, then it gets 9 units and the small pig gets only 1 unit. The large pig not only has heft, but also positioning. (Imagine LeBron James posting up against Justin Bieber on the basketball court.) If, instead, the small pig is there first, it gets 4 of the 10 units, as it consumes some before the large pig arrives to shove it out of the way. If both pigs get there at the same time, the small pig is presumed to get 3 of the 10 units (perhaps mostly from eating the food that falls out of the large pig's mouth).

*Don't get sidetracked on such matters as whether rationality—a concept intended for Man—is applicable to God, or how Man can play a game against someone he's not sure exists, or whether the result is blasphemous because it says that Man should not or will not believe in God. This example is just intended to be a thought-provoking application of game theory.

Each pig decides whether to press the lever or wait at the dispenser. Those are the two strategies in their strategy sets. Assuming that a pig's payoff is the number of units of food consumed less any disutility from pressing the lever, the strategic form of the game is shown in FIGURE 3.14.

FIGURE 3.14 The Boxed-Pigs Game

		Large pig	
		<i>Press lever</i>	<i>Wait at dispenser</i>
Small pig	<i>Press lever</i>	1,5	-1,9
	<i>Wait at dispenser</i>	4,4	0,0

Does the large pig rule the room by being the one that gets to wait at the dispenser? Actually, no. In this setting, "weakness is strength," as it is the large pig that presses the lever while the small pig waits at the dispenser to start consuming the food. How does that outcome emerge?

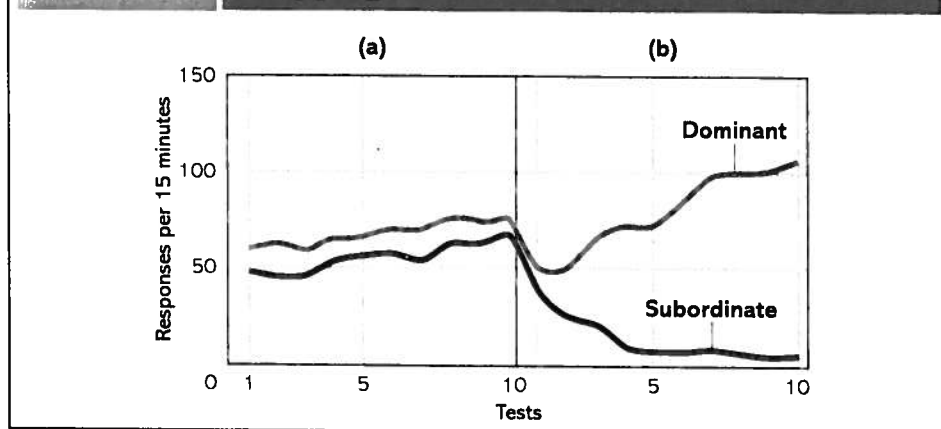
Key to the outcome is that the small pig has a dominant strategy. If the large pig presses the lever, it is preferable for the small pig to wait at the dispenser, since it gets more food then and avoids the disutility from pressing the lever; its payoff is 4 from waiting at the dispenser, compared with 1 from also pressing the lever. If, instead, the large pig waits at the dispenser, then the small pig doesn't get enough food to justify the bother of pressing the lever. It gets only 1 unit of food, and the cost of pressing the lever is 2 units, so its payoff is -1. It would prefer not to press the lever and get a zero payoff. Thus, if the small pig is rational (that might sound a bit odd) then it waits at the dispenser regardless of what the large pig does.

The large pig does not have a dominant strategy. It prefers to wait at the dispenser if the small pig is going to press the lever, but it prefers to press the lever if the small pig is going to wait at the dispenser. If the large pig believes that the small pig is rational (now, that most definitely sounds odd!), then the large pig knows that the small pig will wait at the dispenser. That the small pig won't get enough food to make it worth its while to press the lever serves to make it credible that it won't press the lever. The large pig then has no choice but to press the lever, even though this means that the small pig has the advantage of being at the food first. Of course, once the large pig gets there, it can be assured of getting enough food to justify having pressed the lever. Thus, if pigs are rational and each pig believes that the other pig is rational, then the solution is for the large pig to press the lever and the small pig to wait at the dispenser.

Is saying "if pigs are rational" like saying "if pigs could fly"? It actually is perfectly reasonable to assume that pigs are rational, since this just means that pigs act in their own best interests; don't all species? More problematic is assuming that pigs believe that other pigs are rational; that's dicey. But before you dismiss this solution, let's see if it works by comparing the solution with how pigs actually behave.

The experiment was conducted in a cage measuring about 2.8 meters by 1.8 meters.³ To ensure a strong desire to eat, the pigs were not fed for 24 hours. Each pig was initially put in a cage by itself in order for it to learn that pressing the lever resulted in food being dispensed. Rather than focus on size as the determining factor, the experimenters determined dominance by putting the two pigs

FIGURE 3.15 Experimental Results in the Boxed-Pigs Game. The Graphs Record the Average Number of Presses of the Lever (After 10 Trials) During a 15-Minute Period by Dominant and Subordinate Animals After 24 Hours Without Food When Tested (a) Separately and (b) Together.



together in a room with a bowl of food. A pig was classified as dominant if it spent a higher fraction of the time feeding singly from the bowl.

The results are shown in FIGURE 3.15. On the vertical axis is the number of times the lever was pressed per 15 minutes. On the horizontal axis is the trial number. Up through trial 10, the pigs were in separate cages and the dominant pig pressed the lever slightly more. Starting with trial 10, they were placed in the same cage—and the results are striking: the dominant pig increasingly was the one to press the lever.

I am not claiming that the pigs achieved this outcome by each pig thinking about what the other pig was thinking. A more likely explanation is that they got to it through trial and error. Indeed, note that their behavior gets closer and closer to the predicted outcome over time. Perhaps a few times in which the submissive pig presses the lever and ends up with nothing but crumbs could well induce it to stop pressing the lever, and at that point, the dominant pig learns that the only way it'll eat anything is if it presses the lever. Experience can be a substitute for clever reasoning.

3.4 Solving a Game When Rationality Is Common Knowledge

Man in Black: *All right: where is the poison? The battle of wits has begun. It ends when you decide and we both drink and find out who is right and who is dead.*

Vizzini: *But it's so simple. All I have to do is divine from what I know of you. Are you the sort of man who would put the poison into his own goblet, or his enemy's? Now, a clever man would put the poison into his own goblet, because he would know that only a great fool would reach for what he was given. I'm not a great fool, so I can clearly not choose the wine in front of you. But you must have known I was not a great fool; you would have counted on it, so I can clearly not choose the wine in front of me. —FROM THE MOVIE THE PRINCESS BRIDE*

3.4.1 The Doping Game: Is It Rational for Athletes to Use Steroids?

On August 7, 2007, baseball player Barry Bonds hit his 756th career home run, surpassing the career record of 755 home runs by Henry Aaron. Although this should have been a time of awe and praise for Bonds, the achievement was tainted by allegations that his stellar performance was partially due to neither skill nor hard work, but instead to performance-enhancing steroids. A book titled *Game of Shadows* claims that Bonds engaged in significant steroid use beginning in 1998.

It is well recognized that doping is a serious problem in not only professional, but also amateur sports. The societal challenge is to design a system that deters athletes from using steroids. Such a system would be good not only for fans, but, more importantly, the athletes themselves. Taking steroids is intended to give an athlete a relative advantage over other athletes, but if all athletes use them, then the advantage is lost. Although the benefit evaporates, the cost remains, because athletes still suffer the health consequences. Game theory can be useful for investigating how to structure a system of monitoring and punishments to provide the right incentives.⁴

Although we'll not take on that challenging task here, we can at least identify the temptations faced by athletes and how they can affect their behavior. Consider a randomly selected sport—oh say, such as cycling, and three athletes named Bernhard, Floyd, and Lance. They are assumed to differ in both innate skill and their propensity to take steroids, which could be determined by their desire to win. On the raw-skill dimension, suppose Bernhard is better than Floyd and Floyd is better than Lance. As to the propensity to take steroids, Lance is more inclined to use them than Floyd, and Floyd is more inclined than Bernhard. More specifically, Lance will take steroids regardless of whether Floyd and Bernhard do. Floyd will not take steroids if no one else does, but in order to remain competitive, he'll take them if either Lance or Bernhard (or both) does so. Bernhard, who is the most talented without performance-enhancing drugs, won't take steroids unless both Lance and Floyd do so.

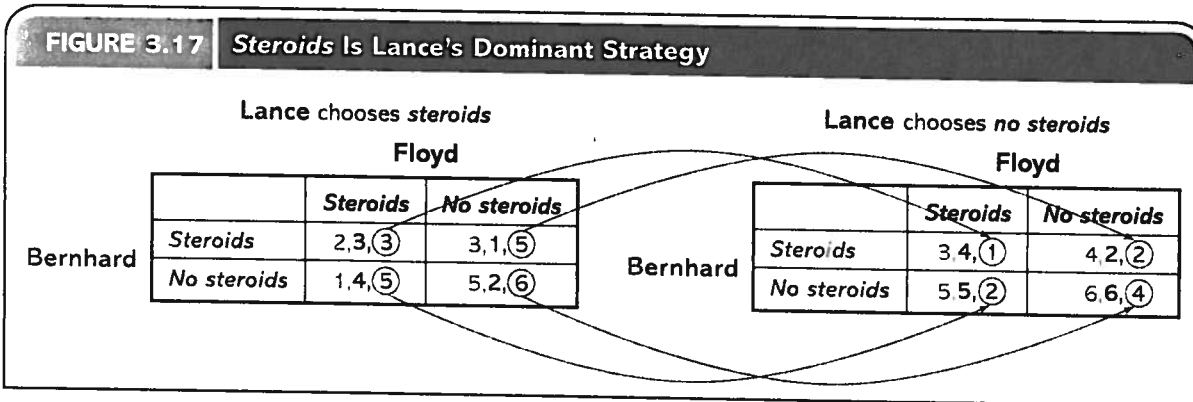
These preferences are embodied in the strategic form game illustrated in FIGURE 3.16, where the first number in a cell is Bernhard's payoff and the second number is Floyd's payoff. Bernhard chooses a row, Floyd chooses a column, and Lance chooses a matrix.

What will these athletes do? Rationality doesn't shed any light on what Floyd and Bernhard will do, as their usage depends on what the other athletes are expected to do. However, Lance has a dominant strategy of taking steroids, as

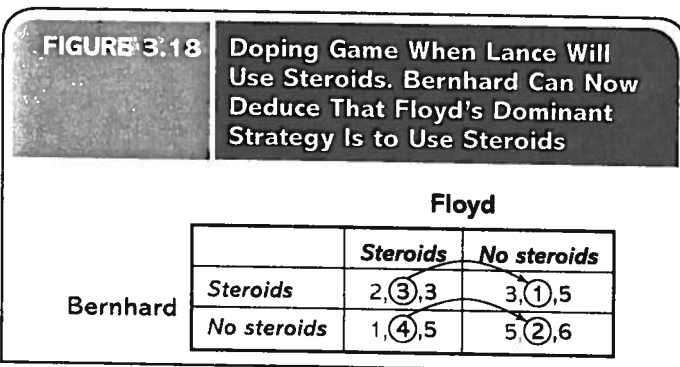
FIGURE 3.16 The Doping Game

		Lance chooses <i>steroids</i>		Lance chooses <i>no steroids</i>	
		Floyd		Floyd	
Bernhard	<i>Steroids</i>	2, 3	3, 1	3, 4	4, 2
	<i>No steroids</i>	1, 4	5, 2	5, 5	6, 6


shown in FIGURE 3.17. If neither Floyd nor Bernhard uses steroids, then Lance's payoff from steroid use is 6, which exceeds his payoff of 4 from abstaining. If one of the other athletes uses steroids (either Floyd or Bernhard), then steroid use for Lance means a payoff of 5, versus a payoff of 2 from staying off of them. Finally, if both Floyd and Bernhard take steroids, then Lance's payoff from using steroids is 3, versus 1 from abstention. Thus, rationality implies that Lance will use steroids.



Let us assume not only that these athletes are rational, but also that each believes that the other two athletes are rational. This assumption implies that both Floyd and Bernhard believe that Lance will use steroids, since the rationality of Lance implies steroid use. From the perspective of Floyd and Bernhard, the game then looks like that shown in FIGURE 3.18, where we've eliminated the *no steroids* strategy for Lance. Floyd now has a dominant strategy of taking steroids. Given that he knows that Lance is going to use them (because Floyd knows that Lance is rational and that *steroids* is the dominant strategy for Lance), it follows that Floyd should do so as well, because it is the best strategy for him, regardless of whether Bernhard uses steroids. Bernhard still lacks a dominant strategy.



Thus far, we know that Lance will use steroids—because he is rational—and that Floyd will use steroids—because Floyd is rational and Floyd knows that Lance is rational. So, what will Bernhard do? Let us make the assumption that each athlete knows that athletes know that athletes are rational. What this assumption buys us is that Bernhard knows that Floyd knows that Lance is rational. Hence, Bernhard knows that Floyd knows that Lance will use steroids, and, therefore, Bernhard knows that Floyd will use steroids. Thus, Bernhard eliminates *no steroids* for Floyd so the situation Bernhard faces is as shown in FIGURE 3.19. Given that



as shown in FIGURE 3.19. Given that

Bernhard then expects both Lance and Floyd to resort to taking steroids, Bernhard finds it optimal to use steroids as well, since it gives a payoff of 2 as opposed to 1.

We conclude that if (1) all athletes are rational, (2) each athlete believes that the other athletes are rational, and (3) each athlete believes that the other athletes believe that the other athletes are rational, then all three of the athletes use steroids. What is depressing about this conclusion is that two of the three athletes don't even want to take steroids and do so only because others are taking them. Lance's strong temptation to enhance his performance through chemicals results in the other two athletes succumbing as well. This is the challenge that sports faces today.

This solution has a ring of truth to it. In *Game of Shadows*, the authors contend that Bonds turned to taking steroids only after the 1998 season, when Mark McGwire and Sammy Sosa were center stage, battling to break Roger Maris's single-season home-run record of 61. Both McGwire and Sosa did in fact surpass 61 home runs; McGwire, who has since admitted to being "juiced" with steroids, set the new record of 70 home runs. Three years later, Bonds broke that record with 73 dingers. If Bonds did take steroids, was it a reaction to remaining competitive with the other top home-run hitters in baseball?

Drug tests are "intelligence tests"—if you can't get around them, you don't deserve to play.— BRUCE SCHNEIER, *WIRED*

3.5 CHECK YOUR UNDERSTANDING

For the game shown in FIGURE 3.20, find the strategies that are consistent with the players being rational, each player believing the other player is rational, and each player believing the other player believes the player is rational.*

FIGURE 3.20

		Player 2			
		w	x	y	z
Player 1	a	3,2	1,1	4,3	3,5
	b	1,3	3,0	2,4	4,2
	c	2,1	0,1	1,2	1,0
	d	1,0	2,0	2,1	4,0

*Answers to Check Your Understanding are in the back of the book.

3.4.2 Iterative Deletion of Strictly Dominated Strategies

Up to now, we have progressively used additional levels of knowledge about rationality in order to solve games. The *Tosca* game was solved with only the assumption that players are rational. The Existence-of-God game required not only that the players were rational, but also that each player believed that all players were rational. (Specifically, we needed God and Man to be rational and Man to believe that God is rational). And with the Doping game, the athletes' decisions regarding steroid use could be derived only when all the players were assumed to be rational, each player believed that all players were

rational, and each player believed that all players believed that all players were rational. These are all examples of a more general procedure for solving a game—a procedure known as the *iterative deletion of strictly dominated strategies* (IDSDS).

The IDSDS algorithm is defined by the following series of steps:

- Step 1** Delete all strictly dominated strategies from the original game. (This step is predicated on the assumption that players are rational.)
- Step 2** Delete all strictly dominated strategies from the game derived after performing step 1. (This step is predicated on the assumption that each player believes that all players are rational.)
- Step 3** Delete all strictly dominated strategies from the game derived after performing step 2. (This step is predicated on the assumption that each player believes that all players believe that all players are rational.)
- Step 4** Delete all strictly dominated strategies from the game derived after performing step 3. (This step is predicated on the assumption that each player believes that all players believe that all players believe that all players are rational.)
- ⋮
- Step t** Delete all strictly dominated strategies from the game derived after performing step $t - 1$.

The procedure continues until no more strategies can be eliminated. In a game with an infinite number of strategies for each player, the procedure could go on forever, but that is not typical. Usually, after a finite number of steps, no more strategies can be eliminated. What remains are the strategies that are said to survive the IDSDS.

Returning to the chapter-opening quote of Sherlock Holmes, we see that IDSDS eliminates the “impossible,” and then whatever remains is what is possible. If only one strategy remains for each player (note that at least one strategy must survive), then the game is **dominance solvable** and the IDSDS delivers a unique prediction regarding behavior.

Let's go through an example to make sure that we understand the procedure. Consider the two-player game illustrated in FIGURE 3.21.

For step 1, consider first player 1. Does she have any strictly dominated strategies? To answer this question, you could consider each strategy and determine whether another strategy is available to produce a strictly higher payoff for every strategy of player 2. A shortcut is to first determine which strategies are optimal for player 1 for *some* strategy of

player 2. Those strategies cannot be strictly dominated, since they are best in some circumstances.

In deploying the tactic of first identifying optimal strategies for player 1, note that if player 2 uses strategy w , then strategy d is optimal for player 1 (giving her

FIGURE 3.21 Applying the IDSDS

		Player 2			
		w	x	y	z
Player 1	a	3,2	4,1	2,3	0,4
	b	4,4	2,5	1,2	0,4
	c	1,3	3,1	3,1	4,2
	d	5,1	3,1	2,3	1,4

a payoff of 5, which exceeds the payoff from any other strategy). Thus, d cannot be strictly dominated. If player 2 uses x , then a is best for player 1, so a cannot be strictly dominated. When player 2 uses y , then c is best, so it is not strictly dominated either. Finally, if player 2 uses z , then c is best, but we already know that c is not strictly dominated. Thus far, we've learned that a , c , and d are not strictly dominated for player 1. This leaves only one remaining strategy to consider, which is b . Though b is not optimal for any strategy of player 2, that property does not imply that b is strictly dominated, so we must check whether or not it is. In fact, b is strictly dominated by d . So, since player 1 is rational, player 1 will avoid using b . Thus, as depicted in FIGURE 3.22, we can delete strategy b from the game in Figure 3.21.

We're not finished with step 1, as the same exercise has to be performed on player 2. Working again with the game in Figure 3.21, we see that if player 1 uses strategy a , then z is best for player 2, in which case z is not strictly dominated. If player 1 uses b , then x is best for player 2, so x is not strictly dominated. If player 1 uses c , then w is optimal for player 2, so w is not strictly dominated. And if player 1 uses d , then z is again optimal for player 2. Hence, strategies w , x , and z are not strictly dominated. The remaining strategy, y , is, however, strictly dominated by z . Since player 2 is rational, we conclude that he will not use y . We can then scratch out strategy y . (See FIGURE 3.23.)

Turning to step 2, we show the reduced game in FIGURE 3.24, where strategy b has been eliminated for player 1 and strategy y has been eliminated for player 2. Are there any strictly dominated strategies that we can eliminate from this game? None are strictly dominated for player 1. (Convince yourself.) For player 2, z strictly dominates x . Note that x was not strictly dominated by z in the original game, because it produced a higher payoff than z (and any other strategy) when player 1 used b . However, since b is strictly dominated for player 1, player 2 doesn't think that player 1 will use it, because player 2 believes that player 1 is rational. Hence, b has been eliminated and, along with it, the reason for keeping x around as a possibly useful strategy. Player 2's other strategies remain undominated.

Since a strategy was eliminated in step 2, the procedure is not over, and we move to step 3. With the elimination of strategy x for player 2, the game is as shown in FIGURE 3.25. Recall that no strategies were

FIGURE 3.22 Eliminating Player 1's Strictly Dominated Strategy

		Player 2			
		w	x	y	z
Player 1	a	3,2	4,1	2,3	0,4
	b	1,1	2,5	1,2	0,4
	c	1,3	3,1	3,1	4,2
	d	5,1	3,1	2,3	1,4

FIGURE 3.23 Eliminating Player 2's Strictly Dominated Strategy


		Player 2			
		w	x	y	z
Player 1	a	3,2	4,1	2,3	0,4
	b	1,1	2,5	1,2	0,4
	c	1,3	3,1	3,1	4,2
	d	5,1	3,1	2,3	1,4

FIGURE 3.24 Reduced Game After One Round of IDSDS

		Player 2		
		w	x	z
Player 1	a	3,2	4,1	0,4
	c	1,3	3,1	4,2
	d	5,1	3,1	1,4

FIGURE 3.25 Reduced Game After Two Rounds of IDSDS

		Player 2	
		w	z
Player 1	a	3,2	0,4
	c	1,3	4,2
	d	5,1	1,4



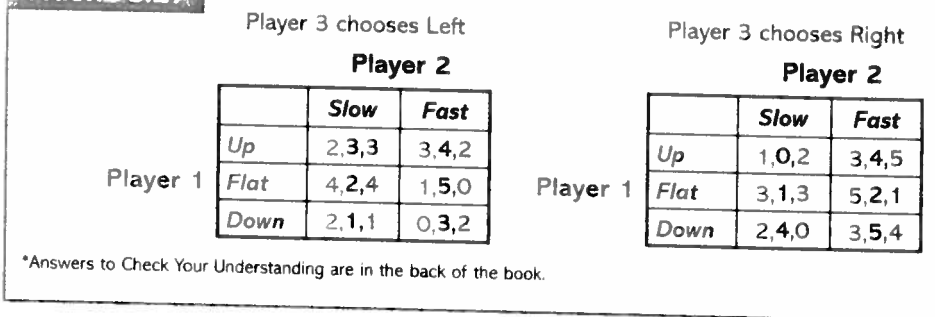
eliminated for player 1. Examining Figure 3.25, note that *d* strictly dominates *a* for player 1, while *c* and *d* remain undominated. Of course, *w* and *z* are still not strictly dominated: They were not strictly dominated in step 2, and the strategies for player 1 in step 3 are the same as those in step 2.

After deleting strategy *a*, we find that the reduced game is as in Figure 3.26, which brings us to step 4. At this point, no strategy is strictly dominated. Since we can't delete any more strategies, the procedure is completed. Our conclusion is that strategies *c* and *d* for player 1 and strategies *w* and *z* for player 2 survive the IDSDS. Thus, assuming that rationality is common knowledge is insufficient to deliver a unique prediction, but it does allow us to eliminate 12 of the 16 possible strategy pairs. All we can say right now is that player 1 will use *c* or *d* and player 2 will use *w* or *z*.

3.6 CHECK YOUR UNDERSTANDING

For the three-player game in Figure 3.27, find the strategies that survive the IDSDS.*

FIGURE 3.27



PLAYING THE GAME

“Guess the Average Number” and Investing in the Stock Market

Should you buy stocks? It depends on whether you think stock prices will rise, which depends on whether other people will find it attractive to own stocks in the future. Should you buy a house or rent? It depends on future housing prices, which depends on whether other people will find it attractive to be a homeowner in the future. Should you buy the painting of a contemporary artist? It depends on the future prices of her paintings, which depends on whether other people will find it attractive to own one of her paintings in the

future. In all of these cases, what is best for you to do depends on what you think will be popular, but what is popular depends on the choices of others. They are like you in that their choices depend on what they think is popular. In other words, popularity involves infinite regress: What is popular depends on what each of us thinks is popular, and that depends on what each of us thinks everyone else thinks is popular, and so on.

In his 1936 landmark book *The General Theory of Employment Interest and Money*, the economist John Maynard Keynes deployed such an argument when comparing the stock market to a beauty contest that was common in English newspapers at the time. The newspaper would print 100 photographs, and people would submit an entry with six faces. Those who identified the faces that were

most popular among the submissions were put into a raffle to win a prize. As Keynes noted: "It is not a case of choosing those [faces] which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees."

These "popularity contests" are ubiquitous in life and are distilled to their essence in a simple game called "Guess the Average Number," which was recently run by National Public Radio's Planet Money: "This is a guessing game. To play, pick a number between 0 and 100. The goal is to pick the number that's closest to half the average of all guesses. So, for example, if the average of all guesses were 80, the winning number would be 40. The game will close at 11:59 P.M. Eastern time on Monday, October 10."


What number should you pick? It doesn't make sense to submit a number above 50 because, even if everyone submitted 100, half of that is 50; thus, the average of all guesses cannot exceed 50. (Formally, submitting 50 weakly dominates submitting any number above 50.) Of course, if everyone submits a number no higher than 50 then, by the same reason-

ing, the average cannot exceed 25. This would seem to argue for not submitting above 25. (Formally, after eliminating weakly dominated numbers, submitting 25 weakly dominates the numbers that remain; that is, it survives two rounds of the iterative deletion of weakly dominated strategies.) But if everyone is following that logic, then the highest submitted number will be 25, in which case the average cannot exceed 12.5. Might that argue for submitting no higher than 12.5? And the argument keeps going all the way down to zero (assuming any number can be submitted and not just integers). That's a big "if," and, more generally, the right answer depends on how many levels you think other people will go.

In the case of NPR's Planet Money, a total of 15,322 listeners submitted numbers; the average was 11.53. FIGURE 3.28 reports the entire distribution of submissions. Some entries are above 50, which makes little sense. More interesting is the spike of submissions around 25, as well as those around 12–14, which indicates that people were following some of the logic expressed above. Of course, this was just a game and the bigger money challenge is figuring out how many levels are used by investors, homebuyers, and art collectors.

www.npr.org/blogs/money/2011/10/03/133654225/please-help-us-pick-a-number www.npr.org/blogs/money/2011/10/11/141249864/heres-the-winner-in-our-pick-a-number-game

FIGURE 3.28



3.5 Do People Believe That People Believe That People Are Rational?

WE'VE BEEN EXPLORING the behavioral implications when players are rational, which we'll refer to as level 1 beliefs), when they further believe other players are rational (level 2), when they yet further believe that other players believe players are rational (level 3), and so forth. It is natural to wonder how many levels people actually use when deciding what to do. Recently, some rather clever experiments have been conducted to shed light on exactly this issue.

FIGURE 3.29

		Other Player		
		x	y	z
You	a	8,8	0,2	0,0
	b	2,2	8,0	2,0
	c	0,2	0,8	0,8

Consider the game in FIGURE 3.29 and suppose we put subjects in an experimental lab to play this game for real money. The experimenter's objective is to determine whether: (1) you are rational; (2) you are rational and you believe the other player is rational; or (3) you are rational, you believe the other player is rational, and you believe the other player believes you are rational. To figure out how the experimenter might do that, let's begin by applying the IDSDS to this game. Strategy *c* is strictly dominated for you, while no strategies are strictly dominated for the other player. If you are rational, then you will not choose *c*. Hence, if we observe you choosing *c* then we can conclude you are not rational.*

In contrast, both strategies *a* and *b* are consistent with rationality; if you believe the other player will choose strategy *x*, then this belief rationalizes the choice of *a*. If you believe the other player will choose *y* or *z*, then that belief would explain why you, as a rational player, would choose *b*. Note that no strategies of the other player are strictly dominated; all of them are consistent with player 2 being rational. Eliminating strategy *c* for you, the game is now as shown in FIGURE 3.30, and we can see that *x* strictly dominates both *y* and *z* for the other player. Finally, in round 3, *a* strictly dominates *b* for you, now that the only strategy that remains for the other player is *x*.

FIGURE 3.30

		Other Player		
		x	y	z
You	a	8,8	0,2	0,0
	b	2,2	8,0	2,0

By this analysis, if you (as player 1) believe the other player believes you are rational then you believe the other player believes you will not choose *c* (recall that it is strictly dominated for you). Given you believe the other player believes you will not choose *c*, then if you believe the other player is rational, you also believe the other player will choose *x*. Given the other

player is expected to choose *x*, then you will choose *a*. Hence, if you have three levels of beliefs—"I am rational," "I believe the other player is rational," and "I believe the other player believes I am rational"—then you will choose strategy *a*. If the experimenter observes you instead playing *b*, she can infer that you do not have three levels of beliefs; you have only one or two levels. However, suppose we observe you playing *a*. That is consistent with three levels, but it is also consistent with just one level. In choosing *a*, you might not have gone through the analysis of trying to figure out what the other player would do if he was rational and believed you were rational, and instead just simply believed (for whatever reasons) that the other player will choose *x*. Observing you choose strategy *a* does not allow us to

*As with all experiments, that inference relies on the validity of the assumption that the payoffs in the game properly represent the payoffs of the subject.

FIGURE 3.31A Player 1's Payoffs

		Player 2		
		<i>p</i>	<i>q</i>	<i>r</i>
Player 1	<i>a</i>	20	8	12
	<i>b</i>	2	18	8
	<i>c</i>	0	12	16

(a)

FIGURE 3.31B Player 2's Payoffs

		Player 3		
		<i>x</i>	<i>y</i>	<i>z</i>
Player 2	<i>p</i>	20	14	8
	<i>q</i>	16	18	2
	<i>r</i>	0	16	16

(b)

FIGURE 3.31C Player 3's Payoffs

		Player 1		
		<i>a</i>	<i>b</i>	<i>c</i>
Player 3	<i>x</i>	12	16	14
	<i>y</i>	8	12	10
	<i>z</i>	6	10	8

(c)

distinguish between the hypothesis that you have level 1 beliefs and the hypothesis that you have level 3 beliefs. Furthermore, this inability to distinguish is not specific to this game but is quite general because if a strategy survives, say, three rounds of IDSDS (and thus is consistent with three levels of beliefs), then it must have satisfied two levels (and thus is consistent with two levels of beliefs) and survived one level (and thus is consistent with one level of beliefs).

Recently, an economist came up with a clever way in which to disentangle low and high levels of beliefs from observed play.⁵ The trick is to have subjects play a particular type of game known as a ring game. A *ring game* is a series of two-player games in which player *i*'s choice of strategy affects player *j*'s payoff, but player *j*'s choice of strategy does not affect player *i*'s payoff. FIGURE 3.31 is an example of a ring game; note that the tables only show the payoffs for a single player. Player 1's payoff depends on the strategies selected by players 1 and 2 (and not 3), player 2's payoff depends on the strategies selected by players 2 and 3 (and not 1), and player 3's payoff depends on the strategies selected by players 3 and 1 (and not 2). For example, consider the strategy profile (*b*, *p*, *z*). Player 1's payoff from 1 choosing *b* and 2 choosing *p* is 2, player 2's payoff from 2 choosing *p* and 3 choosing *z* is 8, and player 3's payoff from 3 choosing *z* and 1 choosing *b* is 10. They are called ring games because of the circular route of interactions as depicted in FIGURE 3.32 where the choice of player 1 impacts the payoff of player 3, the choice of player 3 impacts the payoff of player 2, and the choice of player 2 impacts the payoff of player 1.

If player 1 has three levels of beliefs, let us argue that he will choose strategy *a*. Even though player 3's choice does not affect player 1's payoff, a sophisticated player 1 forms beliefs as to what player 2 believes player 3 will do—as that influences what player 2 does—and player 1 does care about player 2's choice. Given that player 3's strategy *x* strictly dominates *y* and *z*, then if player 1 believes player 2 believes player 3 is rational, then player 1 believes

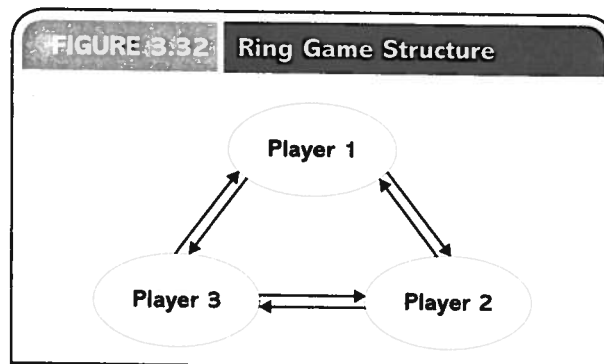


FIGURE 3.33 **Player 3's Payoffs**

		Player 1		
		a	b	c
Player 3	x	8	12	10
	y	12	16	14
	z	6	10	8

player 2 believes player 3 will play *x*, in which case player 2 will optimally play *p* and, therefore, player 1 should play *a*.

Of course, the problem described for the game in Figure 3.29 applies here as well. Even if player 1 does not engage in this highfalutin thinking associated with level 3 beliefs and instead just happens to believe, for whatever reason, that player 2 will choose *p*, then, by virtue of 1 being rational, 1 will choose *a*. In short, choosing *a* is consistent with both level 1 beliefs and level 3 beliefs.

Now comes the critical step for determining whether or not player 1 has level 3 beliefs. After having observed the choice of player 1 in the game of Figure 3.31, the same subject will play the game again, except where the payoffs for player 3 are changed to those in FIGURE 3.33; the payoffs for players 1 and 2 are unchanged. If player 1 has level 1 beliefs, then her play should not change. If player 1 was endowed believing that player 2 would play *p*, there is no reason to think that will change by altering the payoffs of some third player. However, if instead player 1 derived her beliefs about what player 2 would do on the basis that player 2 is rational and that player 2 believes player 3 is rational, then what player 1 thinks player 2 will do could change. With the new payoffs for player 3, strategy *y* strictly dominates *x* and *z* for player 3. Hence, if player 1 believes player 2 believes player 3 is rational, then player 1 believes player 2 believes player 3 will play *y*, in which case player 2 will optimally play *q*, which then means player 1 should play *b*.

If experimental subjects who are given the role of player 1 choose strategy *a* both for the ring game in Figure 3.31 and for the ring game when the payoffs in Figure 3.31c replace those in Figure 3.33, then this subject does not have level 3 beliefs. She may be rational and she may contemplate the implications of player 2's being rational, but she does not contemplate the implications of player 2 believing player 3 is rational. If instead the subject changes her play from strategy *a* to *b*, then this is only consistent with her having level 3 beliefs. She would had to have realized that changing the payoffs of player 3 alters what player 2 will play.

These ring games were part of an experimental design used to empirically uncover the level of beliefs of University of British Columbia undergraduates. The findings are quite intriguing. To begin, 6% of subjects behaved in a manner inconsistent with rationality in that they used strictly dominated strategies. Thus, 94% of subjects were found to be rational. (Does that make you feel better or worse about your fellow man?) Turning to level 2 beliefs, 28% of subjects did not behave in a manner consistent with believing that other players are rational. In other words, the evidence supports 72% of subjects both being rational and believing others are rational. Moving to level 3 beliefs, 44% of subjects were found to be rational, believe others are rational, and believe others believe others are rational. Keep in mind that they may hold even higher level of beliefs, but the experiment was not designed to test for anything above level 3. While it is clear that rationality is not common knowledge, it is also the case that many individuals do engage in higher-level reasoning. Another interesting experiment would be to assess the impact of taking a game theory course on the levels of beliefs!

S
T
e
a
p
w
d
st
d
st

w
(i.
Yc
re
pr

st
str
mc
pla
eli
ing
do
str
in
aft
cor
san
stri
pos

assu
mus
the
but
the
all p
(tha
som
strat
the i
any
I
offer
ests
havin
game
provi

Summary

This chapter outlines methods for solving a game when players are rational, players know that players are rational, players know that players know that players are rational, and so on, and so on. We emphasize the implication that a rational player will not use a **strictly dominated** strategy. A strategy is strictly dominated when another strategy yields a higher payoff, regardless of what the other players do. Thus, no matter what beliefs a player holds concerning the other players' strategies, it is never optimal for him to use a strictly dominated strategy. A **dominant strategy** is surely the unique compelling way to play a game, as it strictly dominates every other strategy.

It is also prudent to avoid playing a weakly dominated strategy. A strategy is **weakly dominated** when the use of another strategy is better some of the time (i.e., for some strategies of the other players) and is at least as good all of the time. You may end up regretting using a weakly dominated strategy, and you'll never regret having avoided it. In some games—such as the second-price auction—the presence of a **weakly dominant strategy** is a good choice for the cautious player.

A procedure known as the **iterative deletion of strictly dominated strategies** (IDSDS) builds on the idea that a rational player does not use a strictly dominated strategy. The key assumption of IDSDS is that it is common knowledge that players are rational and thus common knowledge that players avoid using strictly dominated strategies. The IDSDS procedure eliminates each player's strictly dominated strategies from the game, resulting in a game with fewer strategies (if some strategies are indeed strictly dominated and thus can be deleted). In this smaller game, strictly dominated strategies are again eliminated. A strategy that may not be strictly dominated in the original game may become strictly dominated in the smaller game after some of the other players' strategies are eliminated. This procedure continues—eliminating strictly dominated strategies and then doing the same for the game that remains—until none of the remaining strategies are strictly dominated. Strategies that survive the IDSDS represent the set of possible solutions.

For some games, such as *Tosca*, a unique strategy profile can be derived assuming only that the players are rational. For that to be the case, each player must have a dominant strategy. In other games, like the Existence-of-God game, the derivation of a unique solution requires not only that all players be rational, but also that all players know that all players are rational. Other games, such as the Doping game, require yet more: all players know that all players know that all players are rational. Of course, this procedure of iteratively eliminating stupid (that is, strictly dominated) strategies has traction only if, in the original game, some strategies are indeed stupid. In fact, there are many games in which no strategies are stupid (we'll start seeing them in the next chapter), in which case the iterative deletion of strictly dominated strategies is incapable of eliminating any of the possible strategy profiles.

This chapter has delivered some of the subtle insights that game theory offers. In the *Tosca* game, we showed how players acting in their own best interests can make everyone worse off. In the Cigarette Advertising game, players having fewer options can make themselves better off. And in the Boxed Pigs game, a weaker player can outperform a stronger one. The ensuing chapters will provide many more insightful lessons.

The observant reader will have noticed that we did not solve the game in Figure 3.1 that led off the chapter. Figuring out what happens in the kidnapping scenario is not feasible with the methods of this chapter, because no strategies are strictly dominated. Fortunately, game theory has many more tools up its sleeve, and the next chapter will pull another one out. Like the show business adage says, "Always leave them wanting more."

EXERCISES

1. Derive the strategic form of the Mugging game in Figure 2.9 of Chapter 2 (page 30), and determine whether any strategies are either strictly dominated or weakly dominated.
2. In the Dr. Seuss story "The Zax," a North-Going Zax and a South-Going Zax on their treks soon find themselves facing each other. Each Zax must decide whether to continue in their current direction or move to the side so that the other may pass. As the story reveals, neither of them moves and that stalemate perpetuates for many years. Write down a strategic form game of this situation.
3. For the Team-project game, suppose a jock is matched up with a sorority girl, as shown.

		Sorority girl		
		Low	Moderate	High
Jock	Low	3,0	4,1	5,2
	Moderate	2,2	3,4	4,3
	High	1,6	2,5	3,4

- a. Assume that both are rational and that the jock knows that the sorority girl is rational. What happens?
 - b. Assume that both are rational and that the sorority girl knows that the jock is rational. What happens?
4. Consider the strategic form game shown.

		Player 2		
		x	y	z
Player 1	a	1,3	1,1	0,2
	b	3,1	2,2	1,0
	c	0,2	1,2	3,0

- a. Assume that both players are rational. What happens?
- b. Assume that both players are rational and that each believes that the other is rational. What happens?
- c. Find the strategies that survive the ISDS.

5. For the strategic form game shown, derive the strategies that survive the IDSDS.

		Player 2		
		<i>x</i>	<i>y</i>	<i>z</i>
Player 1	<i>a</i>	5,2	3,4	2,1
	<i>b</i>	4,4	3,2	3,3
	<i>c</i>	3,5	4,4	0,4
	<i>d</i>	2,3	1,5	3,0

6. Two Celtic clans—the Garbh Clan and the Conchubhair Clan—are set to battle. (Pronounce them as you'd like; I don't speak Gaelic.) According to tradition, the leader of each clan selects one warrior and the two warriors chosen engage in a fight to the death, the winner determining which will be the dominant clan. The three top warriors for Garbh are Bevan (which is Gaelic for "youthful warrior"), Cathal (strong in battle), and Duer (heroic). For Conchubhair, it is Fagan (fiery one), Guy (sensible), and Neal (champion). The leaders of the two clans know the following information about their warriors, and each knows that the other leader knows it, and furthermore, each leader knows that the other leader knows that the other leader knows it, and so forth (in other words, the game is common knowledge): Bevan is superior to Cathal against Guy and Neal, but Cathal is superior to Bevan against Fagan. Cathal is superior to Duer against Fagan, Guy, and Neale. Against Bevan, Guy is best. Against Cathal, Neal is best. Against Duer, Fagan is best. Against Bevan, Fagan is better than Neal. Against Cathal, Guy is better than Fagan. Against Duer, Guy and Neal are comparable. Assuming that each leader cares only about winning the battle, what can you say about who will be chosen to fight?

		Player 2			
		<i>w</i>	<i>x</i>	<i>y</i>	<i>z</i>
Player 1	<i>a</i>	1,2	0,5	2,2	4,0
	<i>b</i>	1,3	5,2	5,3	2,0
	<i>c</i>	2,3	4,0	3,3	6,2
	<i>d</i>	3,4	2,1	4,0	7,5


7. Consider the two-player strategic form game depicted.
- Derive the strategies that survive the IDSDS.
 - Derive the strategies that survive the iterative deletion of weakly dominated strategies. (The procedure works the same as the IDSDS, except that you eliminate all *weakly* dominated strategies at each stage.)

8. Consider the three-player game shown. Player 1 selects a row, either a_1 , b_1 or c_1 . Player 2 selects a column, either a_2 or b_2 . Player 3 selects a matrix, either a_3 or b_3 . The first number in a cell is player 1's payoff, the second number is player 2's payoff, and the last number is player 3's payoff. Derive the strategies that survive the IDSDS.

		a_3	
		a_2	b_2
a_1	3, 1, 0	2, 3, 1	
b_1	0, 3, 1	1, 1, 0	
c_1	1, 0, 2	1, 2, 1	

		b_3	
		a_2	b_2
a_1	3, 1, 1	1, 3, 2	
b_1	2, 0, 2	2, 2, 1	
c_1	1, 1, 1	0, 2, 0	

9. A gang controls the drug trade along North Avenue between Maryland Avenue and Barclay Street. The city grid is shown below.



The gang leader sets the price of the drug being sold and assigns two gang members to place themselves along North Avenue. He tells each of them that they'll be paid 20% of the money they collect. The only decision that each of the drug dealers has is whether to locate at the corner of North Avenue and either Maryland Avenue, Charles Street, St. Paul Street, Calvert Street, or Barclay Street. The strategy set of each drug dealer is then composed of the latter five streets. Since the price is fixed by the leader and the gang members care only about money, each member wants to locate so as to maximize the number of units he sells.

For simplicity, assume that the five streets are equidistant from each other. Drug customers live only along North Avenue and are evenly distributed between Maryland Avenue and Barclay Street (so no customers live to the left of Maryland Avenue or to the right of Barclay Street). Customers know that the two dealers set the same price, so they buy from the dealer that is closest to them. The total number of units sold on North Avenue is fixed. The only issue is whether a customer buys from drug dealer 1 or drug dealer 2. This means that a drug dealer will want to locate so as to maximize his share of customers. We can then think about a drug dealer's payoff as being his customer share. The figure below shows the customer shares or payoffs.

Drug Dealers' Payoffs Based on Location

		Dealer 2's location				
		Maryland	Charles	St. Paul	Calvert	Barclay
Dealer 1's location	Maryland	$\frac{1}{2}$	$\frac{7}{8}$	$\frac{3}{4}$	$\frac{5}{8}$	$\frac{1}{2}$
	Charles	$\frac{1}{8}$	$\frac{1}{2}$	$\frac{5}{8}$	$\frac{1}{2}$	$\frac{3}{8}$
	St. Paul	$\frac{1}{4}$	$\frac{3}{8}$	$\frac{1}{2}$	$\frac{3}{8}$	$\frac{1}{4}$
	Calvert	$\frac{3}{8}$	$\frac{1}{2}$	$\frac{5}{8}$	$\frac{1}{2}$	$\frac{1}{8}$
	Barclay	$\frac{1}{2}$	$\frac{5}{8}$	$\frac{3}{4}$	$\frac{7}{8}$	$\frac{1}{2}$

Let us go through a few so that you understand how they were derived. For example, suppose dealer 1 locates at the corner of Maryland and North and dealer 2 parks his wares at the corner of St. Paul and North. All customers who live between Maryland and Charles buy from dealer 1, as he is the closest to them, while the customers who live to the right of Charles buy from dealer 2. Hence, dealer 1 gets 25% of the market and dealer 2 gets 75%. Thus, we see that $(\frac{1}{4}, \frac{3}{4})$ are the payoffs for strategy pair (Maryland, St. Paul). Now, suppose instead that dealer 2 locates at Charles and dealer 1 at Maryland. The customer who lies exactly between Maryland and Charles will be indifferent as to whom to buy from. All those customers to his left will prefer the dealer at Maryland, and they make up one-eighth of the street. Thus, the payoffs are $(\frac{1}{8}, \frac{7}{8})$ for the strategy pair (Maryland, Charles). If two dealers locate at the same street corner, we'll suppose that customers divide themselves equally between the two dealers, so the payoffs are $(\frac{1}{2}, \frac{1}{2})$. Using the IDSDS, find where the drug dealers locate.

- Two students are to take an exam, and the professor has instructed them that the student with the higher score will receive a grade of A and the one with the lower score will receive a B. Student 1's score equals $x_1 + 1.5$, where x_1 is the amount of effort she invests in studying. (That is, I assume that the greater the effort, the higher is the score.) Student 2's score equals x_2 , where x_2 is the amount of effort she exerts. It is implicitly assumed that student 1 is the smarter of the two, in that, if the amount of effort is held fixed, student 1 has a higher score by an amount of 1.5. Assume that x_1 and x_2 can take any value in $\{0, 1, 2, 3, 4, 5\}$. The payoff to student i is $10 - x_i$ if she gets an A and $8 - x_i$ if she gets a B, $i = 1, 2$.
 - Derive the strategies that survive the IDSDS.
 - Derive the strategies that survive the iterative deletion of weakly dominated strategies. (The procedure works the same as the iterative deletion of strictly dominated strategies, except that you eliminate all weakly dominated strategies at each stage.)
- Groucho Marx once said, "I'll never join any club that would have me for a member." Well, Groucho is not interested in joining your investment club, but Julie is. Your club has 10 members, and the procedure for admitting a new member is simple: Each person receives a ballot that has two options: (1) admit Julie and (2) do not admit Julie. Each person can check one of

those two options or abstain by not submitting a ballot. For Julie to be admitted, she must receive at least six votes in favor of admittance. Letting m be the number of ballots submitted with option 1 checked, assume that your payoff function is

$$v(m) = \begin{cases} 1 & \text{if } m = 6, 7, 8, 9, 10 \\ 0 & \text{if } m = 0, 1, 2, 3, 4, 5 \end{cases}$$

- Prove that checking option 1 (admit Julie) is not a dominant strategy.
- Prove that abstaining is a weakly dominated strategy.
- Now suppose you're tired at the end of the day, so that it is costly for you to attend the evening's meeting to vote. By not showing up, you abstain from the vote. This is reflected in your payoff function having the form

$$v(m, \text{action}) = \begin{cases} 1 & \text{if } m = 6, 7, 8, 9, 10 \text{ and you abstained} \\ \frac{1}{2} & \text{if } m = 6, 7, 8, 9, 10 \text{ and you voted} \\ 0 & \text{if } m = 0, 1, 2, 3, 4, 5 \text{ and you abstained} \\ -\frac{1}{2} & \text{if } m = 0, 1, 2, 3, 4, 5 \text{ and you voted} \end{cases}$$

Prove that abstaining is not a weakly dominated strategy.

- Derive all of the rationalizable strategies for the game shown.

		Player 2		
		x	y	z
Player 1	a	0,4	1,1	2,3
	b	1,1	2,2	0,0
	c	3,2	0,0	1,4

- Consider the two-player game:

		Player 2		
		x	y	z
Player 1	a	5,1	4,2	0,1
	b	1,2	0,4	6,3
	c	2,3	1,2	2,1

- Find the strategies that are consistent with both players being rational and each player believing the other player is rational.
 - In addition to that assumed in part (a), assume that player 2 knows player 1 knows player 2 is rational. Find strategies consistent with these beliefs.
- Len and Melanie are deciding what to do Saturday night. The options are to see Mozart's opera *Don Giovanni* or go to the local arena to watch Ultimate Fighter. Len prefers Ultimate Fighter, while Melanie prefers *Don Giovanni*. As a possible compromise, a friend suggests that they attend "*Rocky: The Ballet*," which is a newly produced ballet about Rocky Balboa.

the down-and-out boxer from the streets of Philadelphia who gets a shot at the title. Each would like to go to their most preferred performance, but each also cares about attending with the other person. Also, Len may feel guilty about spending a lot of money for a ticket to Ultimate Fighter when Melanie is not with him; Rocky: The Ballet is cheaper. Don Giovanni and Ultimate Fighter are both expensive tickets, but Melanie would not feel guilty about attending her first choice alone and spending a lot of money. Both Len and Melanie are flying back into town Saturday afternoon and each must independently decide which to attend. The strategic form of the game is shown below. Using the IDSDS, what will they do?

		Melanie		
		Don Giovanni	Ultimate Fighter	Rocky: The Ballet
Len	Don Giovanni	1,5	0,0	0,2
	Ultimate Fighter	3,3	6,1	3,2
	Rocky: The Ballet	4,3	2,0	5,4

15. A total of 10 players are each choosing a number from $\{0,1,2,3,4,5,6,7,8\}$. If a player's number equals exactly half of the average of the numbers submitted by the other nine players, then she is paid \$100; otherwise, she is paid 0. Solve for the strategies that survive the IDSDS.
16. Monica and Isabel are roommates who, on this particular Saturday morning, are trying to decide what scarf to wear. Each has a Burberry scarf (which we'll denote B), a tan scarf (denoted T), and a mauve scarf (denoted M). They care about the scarf but also about whether they end up wearing the same or different scarves. The preference ordering (from best to least preferred outcome) for Monica is: (1) she wears B and Isabel wears T or M; (2) she wears T and Isabel wears B or M; (3) she wears B and Isabel wears B; (4) she wears T and Isabel wears T; (5) she wears M and Isabel wears M; and (6) she wears M and Isabel wears B or T. Isabel's preference ordering is: (1) she wears T and Monica wears B or M; (2) she wears M and Monica wears B or T; (3) she wears T and Monica wears T; (4) she wears M and Monica wears M; (5) she wears B and Monica wears B; and (6) she wears B and Monica wears T or M. Applying the IDSDS, which scarves will be worn?
17. Consider the following game.

		Player 2			
		w	x	y	z
Player 1	a	1,3	4,4	2,2	6,1
	b	0,4	3,2	0,0	5,5
	c	1,2	5,3	2,2	1,6
	d	2,3	2,4	4,2	6,2

- a. Find the strategies that survive the IDSDS.
b. Find the rationalizable strategies.

18. Consider the three-player game below. Player 1 selects a row, either a_1 , b_1 , or c_1 . Player 2 selects a column, either a_2 , b_2 , or c_2 . Player 3 selects a matrix, either a_3 or b_3 or c_3 . The first number in a cell is player 1's payoff, the second number is player 2's payoff, and the last number is player 3's payoff. Derive the strategies that survive the IDSDS.

a_3			
	a_2	b_2	c_2
a_1	3,1,4	2,2,2	3,1,4
b_1	2,4,1	5,3,3	1,2,2
c_1	5,4,5	4,1,6	5,0,1

b_3			
	a_2	b_2	c_2
a_1	1,1,2	3,3,1	2,2,2
b_1	2,2,0	1,1,0	3,0,3
c_1	1,3,3	0,4,1	3,2,2

c_3			
	a_2	b_2	c_2
a_1	4,0,1	3,1,1	3,5,2
b_1	2,5,0	2,4,2	3,2,1
c_1	2,6,3	6,1,3	0,0,0

19. Consider a four-player game in which each player chooses between two strategies: a and b . Their payoffs are shown in the accompanying table for the 16 possible strategy profiles. Find the strategies that survive the IDSDS.

Strategy Profiles				Payoffs			
Player 1	Player 2	Player 3	Player 4	Player 1	Player 2	Player 3	Player 4
a	a	a	a	3	1	2	1
a	a	a	b	2	5	3	3
a	a	b	a	4	2	4	4
a	a	b	b	3	2	5	2
a	b	a	a	2	3	1	0
a	b	a	b	4	4	0	3
a	b	b	a	3	5	2	6
a	b	b	b	2	0	3	5
b	a	a	a	1	5	3	3
b	a	a	b	5	2	1	2
b	a	b	a	1	6	4	5
b	a	b	b	1	3	5	1
b	b	a	a	2	3	2	4
b	b	a	b	2	3	1	0
b	b	b	a	2	7	4	3
b	b	b	b	4	5	3	1

20.

3.6

CONSID
and re
($n - 1$
 $V_i(s'_i, s$
use s'_i

1.

2.

3.

3.7

WE BE
strictl
expec
domi
high
low
tion,

20. For the game below, find the strategies that survive the IDSDS when mixed strategies can be used to eliminate a pure strategy as being strictly dominated.

		Player 2		
		x	y	z
Player 1	a	2.4	3.0	0.1
	b	0.0	1.5	4.2

3.6 Appendix: Strict and Weak Dominance

CONSIDER A GAME WITH n players: $1, 2, \dots, n$. Let S_i denote player i 's strategy set, and read $s_i' \in S_i$ as "strategy s_i' is a member of S_i ." Let S_{-i} be composed of $(n - 1)$ -tuples of strategies for the $n - 1$ players other than player i . Finally, let $V_i(s_i', s_{-i})$ be the payoff of player i when his strategy is s_i' , and the other players use $s_{-i} = (s_1', \dots, s_{i-1}', s_{i+1}', \dots, s_n')$. Then we have the following definitions:

1. Strategy s_i'' *strictly dominates* s_i' if and only if

$$V_i(s_i'', s_{-i}) > V_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

In other words, s_i'' yields a strictly higher payoff than s_i' , regardless of the strategies used by the other $n - 1$ players.

2. s_i'' is the *dominant strategy* if and only if

$$V_i(s_i'', s_{-i}) > V_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}, \text{ for all } s_i \neq s_i''.$$

That is, s_i'' strictly dominates every other strategy for player i .

3. Strategy s_i'' *weakly dominates* s_i' if and only if

$$\begin{aligned} V_i(s_i'', s_{-i}) &\geq V_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}, \text{ and} \\ V_i(s_i'', s_{-i}) &> V_i(s_i', s_{-i}) \text{ for some } s_{-i} \in S_{-i}. \end{aligned}$$

That is, s_i'' yields at least as high a payoff as s_i' for all strategies of the other players and yields a strictly higher payoff for some strategies of the other players.

3.7 Appendix: Rationalizability

WE BEGAN THIS CHAPTER with the statement that a rational player would not use a strictly dominated strategy. Since rationality means choosing what is best, given expectations of other players' strategies, and since the existence of a strictly dominated strategy implies the existence of another strategy that gives a strictly higher payoff regardless of what the other players' strategies are, it logically follows that a rational player will not use a strictly dominated strategy. If, in addition, all players believe that all players are rational, then each player believes that

no other player will use any strictly dominated strategies. This logic led us to eliminate strictly dominated strategies from the game that is derived by deleting strictly dominated strategies from the original game. Continuing in this manner, we derived the strategies that survive the IDSDS.

As just described, the IDSDS eliminates what players would *not do* if rationality were common knowledge. But what is it that they *would do*? If, after using the IDSDS, only one strategy remains for each player, then this procedure gives us a clear and definitive description as to how players will behave. For once all that is "impossible" has been eliminated, then "whatever remains, however improbable, must be the truth." But suppose multiple strategies survive the IDSDS? Although we eliminated what is inconsistent with rationality being common knowledge, is all that remains *consistent* with rationality being common knowledge?

Remember that rationality means acting optimally, given one's beliefs about what other players will do. Thus, a strategy is consistent with rationality only if at least some beliefs about the other players' strategies make that strategy the best one. Let's try working directly with that definition and see what happens.

Consider a two-player game. If player 1 is rational, then she chooses a strategy that maximizes her payoff, given her beliefs as to the strategy of player 2. But what are reasonable beliefs for player 1 to hold about player 2's strategy? If player 1 believes that player 2 is rational, then she will expect player 2 to use a strategy that maximizes *his* payoff, given *his* beliefs about her strategy. Should we allow player 1 to expect that player 2 would hold just any old beliefs as to what player 1 will do? Not if rationality is common knowledge. If player 1 believes that player 2 believes that player 1 is rational, then player 1 believes that player 2's beliefs about player 1's strategy ought to be consistent with player 2's believing that player 1 is rational, which means that player 2 believes that player 1 plays a strategy that is optimal for some beliefs about 2's strategy. Of course, it doesn't end there, so let's jump to a more general statement.

✦ **DEFINITION 3.4** A strategy is **rationalizable** if it is consistent with rationality being common knowledge, which means that the strategy is optimal for a player, given beliefs that are themselves consistent with rationality being common knowledge.

That's not a user-friendly definition, so my plan of explaining rationalizability with more generality may have backfired. So let's move in the other direction and work with a particular example. Consider the game shown in FIGURE A3.1.

FIGURE A3.1 Solving a Game for the Rationalizable Strategies

		Player 2		
		<i>x</i>	<i>y</i>	<i>z</i>
Player 1	<i>a</i>	3,1	1,2	1,3
	<i>b</i>	1,2	0,1	2,0
	<i>c</i>	2,0	3,1	5,0
	<i>d</i>	1,1	4,2	3,3

Think about determining whether strategy *a* is rationalizable for player 1. Are there beliefs about what player 2 will do that would make *a* optimal for player 1? Yes, since *a* is best if and only if player 1 believes that player 2 will use *x*. But does player 2 have beliefs about what player 1 will do that makes it optimal for player 2 to use *x*? If not, then it doesn't make much sense for player 1 to believe that player 2 will use *x* (since player 1 believes that player 2 is rational), and without such a belief, there's not much of an argument for player 1 to use *a*. In fact, *x* is optimal for player 2 if and only if player 2 believes that player 1 will use *b*.

Let's summarize thus far: It makes sense for player 1 to use a if she believes that player 2 will use x . It makes sense for player 1 to believe that player 2 will use x if player 2 believes that player 1 will use b . But then, this just begs another question: Is it reasonable for player 2 to believe that player 1 will use b ? If b is a poor strategy for player 1, then the belief supporting player 2's using x is undermined and, with it, the argument for player 1 to use a . In fact, b is strictly dominated by c for player 1. If player 1 believes that player 2 believes that player 1 is rational, then player 1 should not believe that player 2 believes that player 1 will use b , and thus we cannot rationalize player 2's using x and thus cannot rationalize player 1's using a . Strategy a is not rationalizable, because it is not optimal for player 1 on the basis of beliefs that are themselves consistent with rationality being common knowledge.

Now consider strategy c , and let us argue that it is rationalizable. Strategy c is optimal for player 1 if she believes that player 2 will use z . But is z optimal for player 2, given some beliefs about what player 1 will do? We need that to be the case in order for player 1 to believe that player 2 will use z , because, recall that player 1 believes that player 2 is rational and rational players only use a strategy that is optimal, given their beliefs. If player 2 believes that player 1 will use d , then playing z is indeed best for player 2. But is it reasonable for player 2 to believe that player 1 will use d ? Yes it is, because d is optimal if player 1 believes that player 2 will play y . Hence, it is reasonable for player 1 to believe that player 2 believes that player 1 will play d when player 1 believes that player 2 believes that player 1 believes that player 2 will play y . But is *that* belief consistent with rationality being common knowledge? (You might think that this could never end, but just hang in there for one more round of mental gymnastics.) If player 2 believes that player 1 will play c , then playing y is optimal for player 2. Hence, it is reasonable for player 1 to believe that player 2 believes that player 1 believes that player 2 will play y when player 1 believes that player 2 believes that player 1 believes that player 2 believes that player 1 will play c . Now, what about *that* belief? Well take note that we're back to where we started from, with player 1 playing c . We can then repeat the argument *ad infinitum*:

1. Player 1's playing c is optimal when player 1 believes that player 2 will play z .
2. Player 2's playing z is optimal when player 2 believes that player 1 will play d .
3. Player 1's playing d is optimal when player 1 believes that player 2 will play y .
4. Player 2's playing y is optimal when player 2 believes that player 1 will play c .
5. Player 1's playing c is optimal when player 1 believes that player 2 will play z .
6. Repeat steps 2-5.

After intense use of our "little gray cells" (as the detective Hercule Poirot would say), we conclude that strategy c is rationalizable for player 1 because it is optimal for player 1 given beliefs as to what 2 will do and those beliefs are consistent with rationality being common knowledge. Furthermore, all strategies in that cycle are rationalizable using those beliefs. For example, z is optimal for 2 if 2 believes 1 will use d , and 1 using d is optimal if 1 believes 2 will use y , and y is optimal for 2 if 2 believes 1 will use c , and 1 using c is optimal if 1 believes 2 will use z , at which point we're back where we started from. Hence, strategies c and d are rationalizable for player 1 and y and z for player 2. In fact, one can show that these are the only rationalizable strategies.

If you were to apply the IDSDS to this game, you'd find that those strategies which survive the IDSDS are exactly the same as the rationalizable strategies just derived. Interesting? Coincidence? Not quite. First note that a rationalizable strategy also survives the IDSDS because being rational implies not using a strictly dominated strategy. But can a strategy survive the IDSDS and *not* be rationalizable? Yes, it is possible, although the technical nature of that difference is not one that will concern us in this book. Furthermore, in a wide class of circumstances, the two concepts deliver the same answer. As you can imagine, the IDSDS is vastly easier to understand and use, which are good enough reasons for me to make it the focus of our attention. Nevertheless, it is important to keep in mind that it is the concept of rationalizability that directly encompasses what it means for a strategy to be consistent with rationality being common knowledge.

3.8 Appendix: Strict Dominance with Randomization*

For the game in Figure A3.2, which strategies are strictly dominated for player 1? It would appear none of them. Strategy *a* is the optimal strategy for player 1 when player 2 plays *x* or *y*; hence, *a* is not strictly dominated. Strategy *b* is player 1's optimal strategy when 2 plays *z* and thus is not strictly dominated. While strategy *c* is never the best strategy for player 1 to use, that does not imply it is strictly dominated, and, in fact, neither strategy *a* nor *b* always yield a strictly higher payoff than *c*. Strategy *c* does better than *a* when player 2 chooses *z*, and *c* does better than *b* when 2 chooses *x* or *y*.

Suppose we now expand the set of feasible strategies so that players can randomize in their selection. Such a strategy is referred to as a mixed strategy (which is extensively explored in Chapter 7). For player 1, a mixed strategy is any randomization over strategies *a*, *b*, and *c*, which are now referred to as pure strategies. That is, a player decides on some random device that attaches a probability to each of her pure strategies and then lets the random device determine which pure strategy is played. For example, one mixed strategy is to choose strategy *a* with probability .4, strategy *b* with probability .35, and strategy *c* with probability .25. Another mixed strategy is to choose strategy *a* with probability .8 and strategy *c* with probability .2, so strategy *b* has no chance of being chosen. There are as many mixed strategies as there are ways in which to allocate probabilities over *a*, *b*, and *c*.

By allowing for mixed strategies, more strategies could now strictly dominate a pure strategy. For the game in Figure A3.2, we can show that if player 1 has mixed strategies in her arsenal, strategy *c* is, in fact, strictly dominated. Consider the mixed strategy of choosing *a* with probability 1/2 and *b* with probability 1/2. For each of the possible strategies of player 2, the expected payoff to player 1 from this mixed strategy is provided in Table A3.1. For example, if player 2 chooses *x*, then player 1 will receive a payoff of 5 with probability 1/2 (because she'll choose strategy *a* with probability 1/2) and a payoff of 1 with probability

*This section uses probabilities and expectations. A reader who is unfamiliar with these concepts should first read Section 7.2.

FIGURE A3.2

		Player 1		
		c	b	a
Player 2	z	5,1	4,2	0,1
	y	1,2	0,4	6,3
	x	2,3	1,2	2,1

REFERRE

1. This ap Molotc 44 (196
2. This gr Verlag,
3. B. A. B: users o
4. See, for of Oper
5. Terri Kn Evidenc

TABLE

Player 2:	
x	
y	
z	

TABLE A3.1

Player 2's Strategy	Player 1's Expected Payoff from Choosing <i>a</i> with Probability 1/2 and <i>b</i> with Probability 1/2		Player 1's Payoff from Strategy <i>c</i>
x	$(.5 \times 5) + (.5 \times 1) = 3$	>	2
y	$(.5 \times 4) + (.5 \times 0) = 2$	>	1
z	$(.5 \times 0) + (.5 \times 6) = 3$	>	2

1/2 (because she'll choose strategy *b* with probability 1/2), which results in an expected payoff equal to 3. In comparing these expected payoffs with those from choosing strategy *c* for sure, note that the mixed strategy yields a strictly higher payoff regardless of what player 2 does. Thus, a mixed strategy that chooses *a* half of the time and *b* half of the time strictly dominates strategy *c*. By expanding the set of possible strategies to include randomized ones, strategy *c* can be eliminated on the grounds of strict dominance. The lesson here is that if a pure strategy is not optimal for any strategies of the other players but is not strictly dominated, then it is a candidate for elimination through randomization.

REFERENCES

1. This application is based on Eugene Smolensky, Selwyn Becker, and Harvey Molotch, "The Prisoner's Dilemma and Ghetto Expansion," *Land Economics*, 44 (1968), 419-30.
2. This game is from Steven Brams, *Superior Beings* (New York: Springer-Verlag, 1983). Brams has established himself as one of the most creative users of game theory.
3. B. A. Baldwin and G. B. Meese, "Social Behaviour in Pigs Studied by Means of Operant Conditioning," *Animal Behaviour*, 27 (1979), 947-57.
4. See, for example, Edward J. Bird and Gert G. Wagner, "Sport as a Common Property Resource: A Solution to the Dilemmas of Doping," *Journal of Conflict Resolution*, 41 (1997), 749-66.
5. Terri Kneeland, "Rationality and Consistent Beliefs: Theory and Experimental Evidence," Department of Economics, University of British Columbia, February 7, 2013.

those strategies
le strategies just
a rationalizable
ies not using a
DS and *not* be
f that difference
ide class of cir-
an imagine, the
ugh reasons for
tant to keep in
passes what it
on knowledge.

ization*

d for player 1?
y for player 1
strictly domi-
ategy when 2
While strategy
use, that does
fact, neither
er payoff than
r 2 chooses z,
or y.
ble strategies
ction. Such a
hich is exten-
mixed strat-
referred to as
at attaches a
device deter-
is to choose
ategy *c* with
probability .8
eing chosen.
locate prob-

ly dominate
player 1 has
d. Consider
bability 1/2.
to player 1
if player 2
/2 (because
probability