

Case Studies

THE OTHER PERSON'S ENVELOPE IS ALWAYS GREENER

The inevitable truth about gambling is that one person's gain must be another person's loss. Thus it is especially important to evaluate a gamble from the other side's perspective before accepting. If they are willing to gamble, they expect to win, which means they expect you to lose. Someone must be wrong, but who? This case study looks at a bet that seems to profit both sides. That can't be right, but where's the flaw?

There are two envelopes, each containing an amount of money; the amount of money is either \$5, \$10, \$20, \$40, \$80, or \$160, and everybody knows this. Furthermore, we are told that one envelope contains exactly twice as much money as the other. The two envelopes are shuffled, and we give one envelope to Ali and one to Baba. After both the envelopes are opened (but the amounts inside are kept private), Ali and Baba are given the opportunity to switch. If both parties want to switch, we let them.

Suppose Baba opens his envelope and sees \$20. He reasons as follows: Ali is equally likely to have \$10 or \$40. Thus my expected reward if I switch envelopes is $\$(10 + 40)/2 = \$25 > \$20$. For gambles this small, the risk is unimportant, so it is in my interest to switch. By a similar argument, Ali will want to switch whether she sees \$10 (since she figures that he will get either \$5 or \$20, which has an average of \$12.50) or \$40 (since she figures to get either \$20 or \$80, which has an average of \$50).

Something is wrong here. Both parties can't be better off by switching envelopes, since the amount of money to go around is not getting any bigger by

switching. What is the mistaken reasoning? Should Ali and/or Baba offer to switch?

Case Discussion

A switch should never occur if Ali and Baba are both rational and assume that the other is too. The flaw in the reasoning is the assumption that the other side's willingness to switch envelopes does not reveal any information. We solve the problem by looking deeper into what each side thinks about the other's thought process. First we take Ali's perspective about what Baba thinks. Then we use this from Baba's perspective to imagine what Ali might be thinking about him. Finally, we go back to Ali and consider what she should think about how Baba thinks Ali thinks about Baba. Actually, this all sounds much more complicated than it is. Using the example, the steps are easier to follow.

Suppose that Ali opens her envelope and sees \$160. In that case, she knows that she has the greater amount and hence is unwilling to participate in a trade. Since Ali won't trade when she has \$160, Baba should refuse to switch envelopes when he has \$80, for the only time Ali might trade with him occurs when Ali has \$40, in which case Baba prefers to keep his original \$80. But if Baba won't switch when he has \$80, then Ali shouldn't want to trade envelopes when she has \$40, since a trade will result only when Baba has \$20. Now we have arrived at the case in hand. If Ali doesn't want to switch envelopes when she has \$40, then there is no gain from trade when Baba finds \$20 in his envelope; he doesn't want to trade his \$20 for \$10. The only person who is willing to trade is someone who finds \$5 in the envelope, but of course the other side doesn't want to trade with him.

HERE'S MUD IN YOUR EYE

One of our colleagues decided to go to a Jackson Browne concert at Saratoga Springs. He was one of the first to arrive and scouted the area for the best place to sit. It had rained recently and the area in front of the stage was all muddy. Our colleague settled on the front row closest to the stage yet still behind the muddied area. Where did he go wrong?

Case Discussion

No, the mistake wasn't in picking Jackson Browne. His 1972 hit song "Doctor My Eyes" is still a classic. The mistake was in not looking ahead. As the crowd arrived, the lawn filled up until there was nowhere behind him left to sit. At that point, latecomers ventured into the muddled region. Of course nobody wanted to sit down there. So they stood. Our colleague's view was completely blocked and his blanket equally darkened by the masses of muddled feet.

Here's a case where look forward and reason backward would have made all the difference. The trick is to not choose the best place to sit independently of what others are doing. You have to anticipate where the late arrivals are going to go, and based on this prediction, choose what you anticipate will be the best seat. As the Great Gretzky said in another context, you have to skate to where the puck will be, not where it is.

RED I WIN, BLACK YOU LOSE

While we might never get the chance to skipper in an America's Cup race, one of us found himself with a very similar problem. At the end of his academic studies, Barry celebrated at one of Cambridge University's May Balls (the English equivalent of a college prom). Part of the festivities included a casino. Everyone was given £20 worth of chips, and the person who had amassed the greatest fortune by evening's end would win a free ticket to next year's ball. When it came time for the last spin of the roulette wheel, by a happy coincidence, Barry led with £700 worth of chips, and the next closest was a young Englishwoman with £300. The rest of the group had been effectively cleaned out. Just before the last bets were to be placed, the woman offered to split next year's ball ticket, but Barry refused. With his substantial lead, there was little reason to settle for half.

To better understand the next strategic move, we take a brief detour to the rules of roulette. The betting in roulette is based on where a ball will land when the spinning wheel stops. There are typically numbers 0 through 36 on the wheel. When the ball lands on 0, the house wins. The safest bet in roulette is to bet on even or odd (denoted by black or red). These bets pay even money—a one-dollar bet returns two dollars—while the chance of winning is only 18/37. Even betting her entire stake would not lead to victory at these odds; therefore, the woman was forced to take one of the more risky gambles. She bet her entire stake on the chance that the ball would land on a multiple of three. This bet pays two to one (so her £300 bet would return £900 if she won) but has only a 12/37

chance of winning. She placed her bet on the table.

At that point it could not be withdrawn. What should Barry have done?

Case Discussion

Barry should have copied the woman's bet and placed £300 on the chance that the ball would land on a multiple of three. This would have guaranteed that he stayed ahead of her by £400 and won the ticket: either they both would lose the bet and Barry would win £400 to £0, or they both would win the bet and Barry would end up ahead £1,300 to £900. The woman had no other choice. If she did not bet, she would have lost anyway; whatever she bet on, Barry could have followed her and stayed ahead.*

Her only hope was that Barry would bet first. If Barry had been first to place £200 on black, what should she have done? She should have bet her £300 on red. Betting her stake on black would do her no good, since she would win only when Barry won (and she would place second with £600, compared with Barry's £900). Winning when Barry lost would be her only chance to take the lead, and that dictated a bet on red. The strategic moral is the opposite to that of our tales of Martin Luther and Charles de Gaulle. In this tale of roulette, the person who moved first was at a disadvantage. The woman, by betting first, allowed Barry to choose a strategy that would guarantee victory. If Barry had bet first, the woman could have chosen a response that offered an even chance of winning. The general point is that in games it is not always an advantage to seize the initiative and move first. This reveals your hand, and the other players can use this to their advantage and your cost. Second movers may be in the stronger strategic position.

THE SHARK REPELLENT THAT BACKFIRED

Corporations have adopted many new and innovative ways, often called shark repellent, to prevent outside investors from taking over their company. Without commenting on the efficiency or even morality of these ploys, we present a new and as yet untested variety of shark repellent and ask you to consider how to overcome it.

The target company is Piper's Pickled Peppers. Although now publicly held, the old family ties remain, as the five-member board of directors is completely controlled by five of the founder's grandchildren. The founder recognized the

possibility of conflict between his grandchildren as well as the threat of outsiders. To guard against both family squabbles and outsider attacks, he first required that the board of director elections be staggered. This means that even someone who owns 100 percent of the shares cannot replace the entire board—rather, only the members whose terms are expiring. Each of the five members had a staggered five-year term. An outsider could hope to get at most one seat a year. Taken at face value, it appeared that it would take someone three years to get a majority and control of the company.

The founder was worried that his idea of staggered terms would be subject to change if a hostile party wrested control of the shares. A second provision was therefore added. The procedure for board election could be changed *only* by the board itself. Any board member could make a proposal without the need for a seconder. But there was a major catch. The proposer would be required to vote for his own proposal. The voting would then proceed in clockwise order around the boardroom table. To pass, a proposal needed at least 50 percent of the total board (absences were counted as votes against). Given that there were only five members, that meant at least 3 out of 5. Here's the rub. Any person who made a proposal to change either the membership of the board or the rules governing how membership was determined would be deprived of his position on the board and his stock holdings *if his proposal failed*. The holdings would be distributed evenly among the remaining members of the board. In addition, any board member who voted for a proposal that failed would also lose his seat on the board and his holdings.

For a while this provision proved successful in fending off hostile bidders. But then Sea Shells by the Sea Shore Ltd. bought 51 percent of the shares in a hostile takeover attempt. Sea Shells voted itself one seat on the board at the annual election. But it did not appear that loss of control was imminent, as Sea Shells was one lone voice against four.

At their first board meeting, Sea Shells proposed a radical restructuring of the board membership. This was the first such proposal that the board had ever voted on. Not only did the Sea Shells proposal pass; amazingly, it passed unanimously! As a result, Sea Shells got to replace the entire board immediately. The old directors were given a lead parachute (which is still better than nothing) and then were shown the door.

How did Sea Shells do it? Hint: It was pretty devious. Backward reasoning is the key. First work on a scheme to get the resolution to pass, and then you can worry about unanimity. To ensure that the Sea Shells proposal passes, start at the end and make sure that the final two voters have an incentive to vote for the proposal. This will be enough to pass the resolution, since Sea Shells starts the

process with a first yes vote.

Case Discussion

Many proposals do the trick. Here's one of them. Sea Shells's restructuring proposal has the following three cases:

1. If the proposal passes unanimously, then Sea Shells chooses an entirely new board. Each board member replaced is given a small compensation.
2. If the proposal passes 4 to 1, then the person voting against is removed from the board, and no compensation is made.
3. If the proposal passes with a vote of 3 to 2, then Sea Shells transfers the entirety of its 51 percent share of Piper's Pickled Peppers to the other two yes voters in equal proportion. The two no voters are removed from the board with no compensation.

At this point, backward reasoning finishes the story. Imagine that the vote comes down to the wire: the last voter is faced with a 2–2 count. If he votes yes, it passes and he gets 25.5 percent of the company's stock. If it fails, Sea Shells's assets (and the other yesvoter's shares) are distributed evenly among the three remaining members, so he gets $(51 + 12.25)/3 = 21.1$ percent of the company's stock. He'll say yes.

Everyone can thereby use backward reasoning to predict that if it comes down to a 2–2 tiebreaking vote, Sea Shells will win when the final vote is cast. Now look at the fourth voter's dilemma. When it is his turn to vote, the other votes are:

- i. 1 yes (Sea Shells)
- ii. 2 yes

or

- iii. 3 yes.

If there are three yes votes, the proposal has already passed. The fourth voter would prefer to get something over nothing and therefore votes yes. If there are two yes votes, he can predict that the final voter will vote yes even if he votes no. The fourth voter cannot stop the proposal from passing. Hence, again it is better to be on the winning side, so he will vote yes. Finally, if he sees only one yes vote, then he would be willing to bring the vote to a 2–2 tie. He can safely predict that the final voter will vote yes, and the two of them will make out very nicely indeed.

The first two Piper's board members are now in a true pickle. They can predict that even if they both vote no, the last two will go against them and the proposal will pass. Given that they can't stop it from passing, it is better to go along and get something.

This case demonstrates the power of backward reasoning. Of course it helps to be devious too.

TOUGH GUY, TENDER OFFER

When Robert Campeau made his first bid for Federated Stores (and its crown jewel, Bloomingdales), he used the strategy of a *two-tiered* tender offer. A two-tiered bid typically offers a high price for the first shares tendered and a lower price to the shares tendered later. To keep numbers simple, we look at a case in which the pre-takeover price is \$100 per share. The first tier of the bid offers a higher price, \$105 per share to the first shareholders until half of the total shares are tendered. The next 50 percent of the shares tendered fall into the second tier; the price paid for these shares is only \$90 per share. For fairness, shares are not placed in the different tiers based on the order in which they are tendered. Rather, everyone gets a blended price: all the shares tendered are placed on a prorated basis into the two tiers. Those who don't tender find all of their shares end up in the second tier if the bid succeeds.¹

We can express the average payment for shares by a simple algebraic

expression: if fewer than 50 percent tender, everyone gets \$105 per share; if an amount $X\% \geq 50\%$ of the company's total stock gets tendered, then the average price paid per share is

$$\$150\left(\frac{50}{X}\right) + \$90\left(\frac{X+50}{X}\right) = \$90 + \$15\left(\frac{50}{X}\right).$$

One thing to notice about the way the two-tiered offer is made is that it is unconditional; even if the raider does not get control, the tendered shares are still purchased at the first-tier price. The second feature to note about the way this two-tiered offer works is that if *everyone* tenders, then the average price per share is only \$97.50. This is less than the price before the offer. It's also worse than what they expect should the takeover fail; if the raider is defeated, shareholders expect the price to return to the \$100 level. Hence they hope that the offer is defeated or that another raider comes along.

In fact, another raider did come along, namely Macy's. Imagine that Macy's makes a conditional tender offer: it offers \$102 per share *provided* it gets a majority of the shares. To whom do you tender, and which (if either) offer do you expect to succeed?

Case Discussion

Tendering to the two-tiered offer is a dominant strategy. To verify this, we consider all the possible cases. There are three possibilities to check.

The two-tiered offer attracts less than 50 percent of the total shares and fails.

The two-tiered offer attracts some amount above 50 percent and succeeds.

The two-tiered offer attracts exactly 50 percent. If you tender, the offer will succeed, and without you it fails.

In the first case, the two-tiered offer fails, so that the post-tender price is either \$100 if both offers fail or \$102 if the competing offer succeeds. But if you tender you get \$105 per share, which is bigger than either alternative. In the second case, if you don't tender you get only \$90 per share. Tendering gives you at

worst \$97.50. So again it is better to tender. In the third case, while other people are worse off if the offer succeeds, you are privately better off. The reason is that since there are exactly 50 percent tendered, you will be getting \$105 per share. This is worthwhile. Thus you are willing to push the offer over.

Because tendering is a dominant strategy, we expect everyone to tender. When everyone tenders, the average blended price per share may be below the pre-bid price and even below the expected future price should the offer fail. Hence the two-tiered bid enables a raider to pay less than the company is worth. The fact that shareholders have a dominant strategy does not mean that they end up ahead. The raider uses the low price of the second tier to gain an unfair advantage. Usually the manipulative nature of the second tier is less stark than in our example because the coercion is partially hidden by the takeover premium. If the company is really worth \$110 after the takeover, then the raider can still gain an unfair advantage by using a second tier below \$110 but above \$100. Lawyers view the two-tiered bid as coercive and have successfully used this as an argument to fight the raider in court. In the battle for Bloomingdales, Robert Campeau eventually won, but with a modified offer that did not include any tiered structure.

We also see that a conditional bid is not an effective counter-strategy against an unconditional two-tiered bid. In our example, the bid by Macy's would be much more effective if its offer of \$102 per share were made unconditionally. An unconditional bid by Macy's destroys the equilibrium in which the two-tiered bid succeeds. The reason is that if people thought that the two-tiered bid were certain to succeed, they would expect a blended price of \$97.50, which is less than they would receive by tendering to Macy's. Hence it cannot be that shareholders expect the two-tiered bid to succeed and still tender to it.*

In late 1989, Campeau's operations unraveled because of excessive debt. Federated Stores filed for reorganization under Chapter 11 of the bankruptcy law. When we say Campeau's strategy was successful, we merely mean that it achieved the aim of winning the takeover battle. Success in running the company was a different game.

THE SAFER DUEL

As pistols become more accurate, does that change the deadliness of a duel?

Case Discussion

At first glance, the answer would seem to be obvious: yes. But recall that the players will adapt their strategies to the new situation. Indeed, the answer is easier to see if we flip the question: suppose we try to make dueling safer by reducing the accuracy of the pistols. The new outcome is that the adversaries will come closer to one another before firing.

Recall our discussion of the duel on chapter 10. Each player waits to shoot until the point where his probability of hitting the other side is just equal to the other side's chance of missing. Note that the accuracy of the pistols doesn't enter into the equation. All that matters is the ultimate chance of success.

To illustrate this point with some numbers, suppose that the adversaries are equally good shots. Then the optimal strategy is for the two to keep on approaching each other until the moment that the probability of hitting reaches $1/2$. At that point one duelist takes a shot. (It doesn't matter which person shoots, as the chance of success is a half for the shooter and a half for the person who is being shot at.) The probability each player will survive is the same ($1/2$) irrespective of the accuracy of the pistols. A change in the rules need not affect the outcome; all the players will adjust their strategies to offset it.

THE THREE-WAY DUEL

Three antagonists, Larry, Moe, and Curly, are engaged in a three-way duel. There are two rounds. In the first round, each player is given one shot: first Larry, then Moe, and then Curly. After the first round, any survivors are given a second shot, again beginning with Larry, then Moe, and then Curly. For each duelist, the best outcome is to be the sole survivor. Next best is to be one of two survivors. In third place is the outcome in which no one gets killed. Dead last is that you get killed.

Larry is a poor shot, with only a 30 percent chance of hitting a person at whom he aims. Moe is a much better shot, achieving 80 percent accuracy. Curly is a perfect shot—he never misses. What is Larry's optimal strategy in the first round? Who has the greatest chance of survival in this problem?

Case Discussion

Although backward reasoning is the safe way to solve this problem, we can jump ahead a little by using some forward-looking arguments. We start by

examining each of Larry's options in turn. What happens if Larry shoots at Moe? What happens if Larry shoots at Curly?

If Larry shoots at Moe and hits, then he signs his own death warrant. It becomes Curly's turn to shoot, and he never misses. Curly will not pass at the chance to shoot Larry, as this leads to his best outcome. Larry shooting at Moe does not seem to be a very attractive option.

If Larry shoots at Curly and hits, then it is Moe's turn. Moe will shoot at Larry. (Think about how we know this to be true.) Hence, if Larry hits Curly, his chance of survival is less than 20 percent, the chance that Moe misses.

So far, neither of these options looks to be very attractive. In fact, Larry's best strategy is to fire up in the air! In this case, Moe will shoot at Curly, and if he misses, Curly will shoot and kill Moe. Then it becomes the second round and it is Larry's turn to shoot again. Since only one other person remains, he has at least a 30 percent chance of survival, since that is the probability that he kills his one remaining opponent.

The moral here is that small fish may do better by passing on their first chance to become stars. We see this every four years in presidential campaigns. When there is a large number of contenders, the leader of the pack often gets derailed by the cumulative attacks of all the medium-sized fish. It can be advantageous to wait, and step into the limelight only after the others have knocked each other and themselves out of the running.

Your chances of survival depend on not only your own ability but also whom you threaten. A weak player who threatens no one may end up surviving if the stronger players kill each other off. Curly, although he is the most accurate, has the lowest chance of survival—only 14 percent. So much for survival of the fittest! Moe has a 56 percent chance of winning. Larry's best strategy turns his 30 percent accuracy into a 41.2 percent chance of winning.²

THE RISK OF WINNING

One of the more unusual features of a Vickrey sealed-bid auction is that the winning bidder does not know how much she will have to pay until the auction is over and she has won. Remember, in a Vickrey auction the winning bidder pays only the second highest bid. In contrast, there is no uncertainty in the more standard sealed-bid auction, in which the winner pays her bid. Since everyone knows her own bid, no one has any doubts as to how much she will have to pay if she wins.

The presence of uncertainty suggests that we might want to consider the

effect of risk on the participants' bidding strategies. The typical response to uncertainty is negative: the bidders are worse off in a Vickrey auction because they do not know how much they will have to pay if they have submitted the winning bid. Is it reasonable that a bidder will respond to this uncertainty or risk by lowering her bid below the true valuation?

Case Discussion

It is true that the bidders dislike the uncertainty associated with how much they might have to pay if they win. Each is in fact worse off. Yet, in spite of the risk, participants should still bid their true valuations. The reason is that a truthful bid is a dominant strategy. As long as the selling price is below the valuation, the bidder wants to buy the good. The only way to ensure that you win whenever the price is below your value is to bid the true value.

In a Vickrey auction, bidding the true valuation doesn't make you pay more—except when someone else would have outbid you, in which case you would have wanted to raise your bid until the selling price exceeded your valuation. The risk associated with a Vickrey auction is limited; the winner is never forced to pay an amount greater than her bid. While there is uncertainty about what the winner will pay, this uncertainty is only over the degree of good news. Even though the good news might be variable, the best strategy is to win the auction whenever it's profitable. That means bidding your true value. You never miss a profitable opportunity, and whenever you win you pay less than your true value.

BUT ONE LIFE TO LAY DOWN FOR YOUR COUNTRY

How can the commanders of an army motivate its soldiers to risk their lives for their country? Most armies would be finished if each soldier on the battlefield started to make a rational calculation of the costs and the benefits of risking his own life. What are the various devices that can motivate and incentivize soldiers to risk their lives?

Case Discussion

First look at some devices that transform the soldiers' self-regarding rationality. The process begins in boot camp. Basic training in the armed forces

everywhere is a traumatic experience. The new recruit is maltreated, humiliated, and put under such immense physical and mental strain that the few weeks quite alter his personality. An important habit acquired in this process is an automatic, unquestioning obedience. There is no reason why socks should be folded, or beds made, in a particular way, except that the officer has so ordered. The idea is that the same obedience will occur when the order is of greater importance. Trained not to question orders, the soldier becomes a fighting machine; commitment is automatic.

Many armies got their soldiers drunk before battle. This may have reduced their fighting efficiency, but it also reduced their capacity for rational calculation of self-preservation.

The seeming irrationality of each soldier turns into strategic rationality. Shakespeare knew this perfectly well; in *Henry V*, the night before the battle of Agincourt (fought on St. Crispin's day, October 25, 1415), King Henry prays (emphasis added):

O God of battles! steel my soldiers' hearts;
Possess them not with fear; *take from them now*
The sense of reckoning, if th'opposed numbers
Pluck their hearts from them

Just before the battle, Henry does something that may at first seem to defeat his purpose. Instead of enforcing any compulsion to fight, he declares:

...he which hath no stomach to this fight,
Let him depart; his passport shall be made,
And crowns for convoy put into his purse:
We would not die in that man's company That fears his fellowship to die
with us.

The catch is that anyone who wants to take up this offer has to do so in full view of all of his companions. Of course everyone is too ashamed to do so. And the action (actually, inaction) of publicly declining the offer changes soldiers' preferences, even personalities, irrevocably. By their act of rejecting the offer, the soldiers have psychologically burned their ships home. They have established an implicit contract with each other not to flinch from death if the

time comes.*

Next consider incentives to act. These can be material: in the old days, victorious soldiers had the opportunity to loot from the property and even the bodies of the enemy. Generous death benefits can be promised for next-of-kin if the worst happens. But the incentives to fight and risk lives are mostly nonmaterial: medals, honor, and glory come to the brave whether they live or die in battle; the lucky survivors can boast of their exploits for years to come. Here is Shakespeare's King Henry V again:

He that shall live this day, and see old age,
Will yearly on the vigil feast his neighbours,...he'll remember with
advantages What feats he did that day...
And Crispin Crispian shall ne'er go by,
From this day to the ending of the world,
But we in it shall be remember'd;
We few, we happy few, we band of brothers;
For he to-day that sheds his blood with me Shall be my brother;...
And gentlemen in England now a-bed Shall think themselves accursed they
were not here,
And hold their manhoods cheap whiles any speaks That fought with us
upon Saint Crispin's day.

Being the king's brother; others holding their manhoods cheap when you speak: what powerful incentives! But think a moment. What does it really mean to be the king's brother? Suppose you live and return to England with the victorious army. Is the king going to say: "Ah, my brother! Come and live with me at the palace." No. You will return to the same old life of poverty that you had before. In concrete terms, the incentive is empty. It is like the "cheap talk" we mentioned in connection with credibility. But it works. The science of game theory cannot fully explain why. Henry's speech is the art of strategy at its best.

There is a related subtext. The night before the battle, Henry goes wandering in disguise among his troops to find out what they are really thinking and feeling. He discovers one disconcerting fact: they are afraid of being killed or captured, and they believe that he does not face the same risk. Even if the enemy gets to him, they will not kill him. It will be more profitable to hold him for ransom and this will then be paid. Henry must dispel this fear if he is to command the soldiers' loyalty and solidarity. It would not do in his speech the


following morning to say: “Hey, guys; I hear some of you think that I am not risking my life with you. Let me assure you most earnestly that I am.” That would be worse than useless; it would have the effect of reinforcing the soldier’s worst suspicions, rather like Richard Nixon’s declaration “I am not a crook” during the Watergate crisis. No; in his speeches Henry simply takes it for granted that he is risking his life and turns the question around: “Are you risking your life *with me*?” That is how we should interpret the phrases “we would not die in that man’s company” and “he that sheds his blood with me.” Once again, it is a beautiful example of the art of strategy.

Of course this is not actual history but Shakespeare’s fictionalization of it. However, we think that artists often have more perceptive insights about human emotions, reasoning, and motivation than do psychologists, let alone economists. Therefore we should be willing to learn lessons on the art of strategy from them.

WINNING WITHOUT KNOWING HOW


Chapter 2 introduced games in which players move in sequence and which always end after a finite number of moves. In theory, we could examine every possible sequence of moves and thereby discover the best strategy. This is relatively easy for tic-tac-toe and impossible (at present) for chess. In the game below, the best strategy is unknown. Yet, even without knowing what it is, the very fact that it exists is enough to show that it must lead to a win for the first player.

ZECK is a dot game for two players. The object is to force your opponent to take the last dot. The game starts with dots arranged in any rectangular shape, for example 7×4 :



Each turn, a player removes a dot and with it *all* remaining dots to the northeast. If the first player chooses the fourth dot in the second row, this leaves his

opponent with




Each period, at least one dot must be removed. The person who is forced to take the last dot loses.

For any shaped rectangle with more than one dot, the first player must have a winning strategy. Yet this strategy is not currently known. Of course we can look at all the possibilities and then figure it out for any particular game, such as the 7×4 above—but we don't know the best strategy for all possible configurations of dots. How can we show who has the winning strategy without knowing what it is?

Case Discussion

If the second player has a winning strategy, that means that for *any* opening move of the first player, the second has a response that puts him in a winning position. In particular, this means that the second player must have a winning response even if the first player just takes the upper-right-hand dot.



But no matter how the second player responds, the board will be left in a configuration that the first player could have created in his first move. If this is truly a winning position, the first player should have and could have opened the

game this way. There is nothing the second player can do to the first that the first player can't do unto him beforehand.

A BURQA FOR PRICES

Hertz and Avis advertise that you can rent a car for \$19.95/day. But that car rental price typically leaves out the inflated cost of filling up the tank at the return, often twice the price at the pump. Ads for hotel room rates don't mention the \$2/minute charge for long-distance calls. When choosing between HP and Lexmark printers, who has the cheaper cost per page? It is hard to tell when the toner cartridges don't let you know how many pages you'll get. Cell phone companies offer plans with a fixed number of minutes per month. Minutes you don't use are lost, and if you go over, there is a steep charge.* The ad promising 800 minutes for \$40/month will almost always cost more than 5¢/minute. As a result, it becomes difficult, if not impossible, to understand or compare the real cost. Why does this practice persist?

Case Discussion

Consider what would happen if one car rental company decided to advertise its all-in price. This maverick would have to set a higher daily rental price in order to make up for the lost revenue from overcharging for gas. (That would still be a good idea: wouldn't you rather pay an extra \$2/day and then not have to worry about finding a place to fill up as you dash back to the airport? This might save you from missing the flight or even save your marriage.) The problem is that the company who plays it straight puts itself at a disadvantage compared to its rivals. The one honest firm would seem to be charging the *highest* price when customers do a comparison on Expedia. There isn't an asterisk that says, "We don't rip you off on gas like everyone else does."

The problem is that we are stuck in a bad equilibrium, much like the one involving the QWERTY keyboard. Customers assume that the prices will include lots of hidden extras. Unless a firm can cut through the clutter and convince customers that they aren't playing the same game, the honest firm will just seem to be too expensive. Worse still, since customers don't know the true cost at the rival firms, they don't know how much they should pay. Imagine that a cell phone company offered a single flat price per minute. Does 8¢/minute beat \$40 for 800 minutes (with a 35¢ per minute surcharge for going over)? Who

knows?

The bottom line is companies go on advertising just one component of the total price. The parts they don't mention are then priced at exorbitant levels. But that doesn't mean that firms end up making more money. Because each company can anticipate making high profits on the back end, they are willing to go to extraordinary lengths to attract or steal customers. Thus laser printers are practically given away, as are most cell phones. The firms compete away all of their future profits in the battle to attract customers. The end result is too much switching and the loss of customer loyalty.

If society wants to improve matters for consumers, one way would be to legislate a change in the convention: require that hotels, car rental companies, and cell phone providers advertise the all-in price paid by the average customer. Comparison shopping sites now do this for books sold online, where the all-in price comparison includes the cost of shipping and handling.³

KING SOLOMON'S DILEMMA REDUX

King Solomon wanted to find a way to obtain some information: who was the real mother? The two women who possessed the information had conflicting incentives about revealing it. Mere words would not suffice; strategic players would willingly manipulate answers in their own interests. What is needed is some way to make the players put their money, or, more generally, something they value, where their mouths are. How could a game theory king have persuaded the two women to tell the truth?

Case Discussion

Of several devices that work even when both women play strategically, here is the simplest.⁴ Call the two women Anna and Bess. Solomon sets up the following game:


Move 1: Solomon decides on a fine or punishment.

Move 2: Anna is asked to either give up her claim, in which case Bess gets the child and the game ends, or to assert her claim, in which case we go on to...

Move 3: Bess can either accept Anna's claim, in which case Anna gets the child and the game ends, or challenge Anna's claim. In the latter case, Bess must put in a bid B of her own choosing for the child, and Anna must pay the fine F to Solomon. We go on to...

Move 4: Anna can either match Bess's bid, in which case Anna gets the child and pays B to Solomon, while Bess pays the fine F to Solomon; or Anna does not match, in which case Bess gets the child and pays her bid B to Solomon.

Here is the game in tree form:



As long as the true mother values the child more than the false claimant, in the subgame perfect equilibrium the true mother gets the child. Solomon does not have to know these values. No fines or bids are actually paid; their sole purpose is to avoid any false claims by either woman.

The reasoning is simple. First suppose Anna is the true mother. Bess knows in move 3 that, unless she bids more than the child is worth to her, Anna will match her bid in move 4, and she (Bess) will end up paying the fine and not getting the child. So Bess will not bid. Knowing this, Anna in move 2 will claim the child and get it. Next suppose Bess is the true mother. Then Anna knows in move 2 that Bess in move 3 will choose a bid that is not worth Anna's while to match in move 4, so she (Anna) is simply going to end up paying the fine F and not getting the child. So in move 2 Anna does best for herself by renouncing her claim.

At this point you are no doubt criticizing us for reducing everything to the sordid world of money. We respond by pointing out that in the actual play that results in the equilibrium of this game, the bids are not actually paid, and neither

is the fine. Their only purpose is as a threat; they make it costly for either woman to lie. In this respect, they are similar to the threat of cutting the child in two and, we would argue, a lot less gruesome.

One potential difficulty remains. For the device to work, it must be the case that the true mother is able to bid at least as much as the false claimant. Presumably she loves and values the child at least as much in a subjective sense, but what if she does not have as much money to back up her value? In the original story, the two women came from the same household (actually the book says that they were both prostitutes), so Solomon could reasonably regard their abilities to pay as approximately equal. Even otherwise, the difficulty can be resolved. The bids and fines need not be monetary sums at all. Solomon can specify them in some other “currency” that the two women should be expected to possess in nearly equal amounts, for example having to perform a certain number of days of community service.

BAY BRIDGE

The morning traffic from Oakland to San Francisco across the Bay Bridge gets backed up from 7:30 A.M. to 11:00 A.M. Until the jam clears at 11:00, each additional car that enters the traffic makes all those who come later wait just a little longer. The right way to measure this cost is to sum up the additional waiting times across everyone who is delayed. What is the total waiting-time cost imposed by one additional car that crosses the bridge at 9:00 A.M.?

You may be thinking you don’t know enough information. A remarkable feature of this problem is that the externality can be calculated based on the little you’ve been told. You don’t need to know how long it takes the cars to cross the toll plaza, nor the distribution of cars that arrive after 9:00. The answer is the same whether the length of the traffic jam stays constant or varies widely until it clears.

Case Discussion

The trick is to see that all that matters is the sum of the waiting time. We are not concerned with who waits. (In other circumstances, we might want to weigh the waiting times by the monetary value of time for those caught in the jam.) The simplest way to figure out the total extra waiting time is to shuffle around who waits, putting all the burden on one person. Imagine that the extra driver, instead

of crossing the bridge at 9:00 A.M., pulls his car over to the side and lets all the other drivers pass. If he passes up his turn in this way, the other drivers are no longer delayed by the extra car. Of course, he has to wait two hours before the traffic clears. But these two hours exactly equal the total waiting time imposed on all the other drivers if he were to cross the bridge rather than wait on the sidelines. The reason is straightforward. The total waiting time is the time it takes for everyone to cross the bridge. Any solution that involves everyone crossing the bridge gives the same total waiting time, just distributed differently. Looking at the solution in which the extra car does all the extra waiting is the easiest way to add up the new total waiting time.

WHAT PRICE A DOLLAR?

Professor Martin Shubik of Yale University designed the following game of entrapment. An auctioneer invites bids for a dollar. Bidding proceeds in steps of five cents. The highest bidder gets the dollar, but *both* the highest *and* the second highest bidders pay their bids to the auctioneer.⁵

Professors have made tidy profits—enough for a lunch or two at the faculty club—from unsuspecting undergraduates playing this game in classroom experiments. Suppose the current highest bid is 60 cents and you are second with 55. The leader stands to make 40 cents, but you stand to lose your 55. By raising to 65, you can put the boot on the other foot. The logic is no different when the leading bid is \$3.60 and yours is \$3.55. If you do not raise the bidding still further, the “winner” loses \$3.60, but you lose \$3.55.

How would you play this game?

Case Discussion

This is an example of the slippery slope. Once you start sliding, it is hard to recover. It is better not to take the first step unless you know where you are going.

The game has one equilibrium, in which the first bid is a dollar and there are no further bids. But what happens if the bidding starts at less than a dollar? The escalation has no natural limit other than the amount of money in your wallet: the bidding must stop when you run out of money. That is all we need to apply Rule 1: Look forward and reason backward.

Imagine that Eli and John are the two students in Shubik’s auction of a

dollar. Each has \$2.50 in his wallet, and each knows the other's cash supply.⁶ To keep things simple, bidding takes place in dime units.

To start at the end, if Eli ever bids \$2.50, he'll win the dollar (and be down \$1.50). If he bids \$2.40, then John must bid \$2.50 in order to win. Since it is not worth spending a dollar to win a dollar, an Eli bid of \$2.40 will win if John's current bid is at \$1.50 or less.

The same argument works if Eli bids \$2.30. John can't bid \$2.40 and expect to win, because Eli would counter with \$2.50. To beat \$2.30, John needs to go all the way up to \$2.50. Hence a \$2.30 bid beats \$1.50 and below. So does a \$2.20 bid, a \$2.10 bid, all the way down to a \$1.60 bid. If Eli bids \$1.60, John should predict that Eli won't give up until the bidding reaches \$2.50. Eli's \$1.60 is already lost, but it is worth his while to spend another 90 cents to capture the dollar.

The first person to bid \$1.60 wins, because that establishes a credible commitment to go up to \$2.50. In our mind, we should think of \$1.60 as the same sort of winning bid as \$2.50. In order to beat \$1.50, it suffices to bid \$1.60, and nothing less will do. That means \$1.50 will beat all bids at 60 cents and below. Even a bid of 70 cents will beat all bids at 60 cents and below. Why? Once someone bids 70 cents, it is worthwhile for them to go up to \$1.60 and be guaranteed victory. With this commitment, no one with a bid of 60 cents or less finds it worthwhile to challenge.

We expect that either John or Eli will bid 70 cents and the bidding will end. Although the numbers will change, the conclusion does not depend on there being just two bidders. Given that budgets differ, backward reasoning can still find the answer. But it is critical that everyone know everyone else's budget. When budgets are unknown, as one would expect, an equilibrium will exist only in mixed strategies.

Of course there is a much simpler and more profitable solution for the students: collusion. If the bidders agree among themselves, a designated person will bid a dime, no one else will bid at all, and the class will share the profit of 90 cents.

You may take this story as proof of the folly of Yale undergraduates. But was the escalation of the superpowers' nuclear arms arsenals all that different? Both incurred costs in the trillions of dollars in quest of the "dollar" of victory. Collusion, which in this case means peaceful coexistence, is a much more profitable solution.

THE KING LEAR PROBLEM

Tell me, my daughters
Since now we will divest us both of rule,
Interest of territory, cares of state,
Which of you shall we say doth love us most?
That we our largest bounty may extend
Where nature doth with merit challenge.
—Shakespeare, *King Lear*

King Lear was worried about how his children would treat him in his old age. Much to his regret, he discovered that children do not always deliver what they promise. In addition to love and respect, children are also motivated by the possibility of an inheritance. Here we look at how a strategic use of inheritance can manipulate children to visit their parents.

Imagine that parents want each of their children to visit once and phone twice a week. To give their children the right incentives, they threaten to disinherit any child who fails to meet this quota. The estate will be evenly divided among all the children who meet this quota. (In addition to motivating visits, this scheme has the advantage of avoiding the incentive for children to suffocate their parents with attention.)

The children recognize that their parents are unwilling to disinherit all of them. As a result, they get together and agree to cut back the number of visits, potentially down to zero.

The parents call you in and ask for some help in revising their will. Where there is a will, there is a way to make it work. But how? You are not allowed to disinherit all of the children.

Case Discussion

As before, any child who fails to meet the quota is disinherited. The problem is what to do if all of them are below the quota. In that case, give *all* of the estate to the child who visits the most. This will make the children's reduced visiting cartel impossible to maintain. We have put the children into a multiperson dilemma. The smallest amount of cheating brings a massive reward. A child who makes just one more phone call increases his or her inheritance from an equal

share to 100 percent. The only escape is to go along with the parents' wishes. (Obviously, this strategy fails with only children. There is no good solution for couples with an only child. Sorry.)

UNITED STATES V. ALCOA

An established firm in an industry stands to gain by keeping out new competition. Then it can raise prices to monopoly levels. Since monopoly is socially harmful, the antitrust authorities try to detect and prosecute firms that employ strategies to deter rivals from entering the business.

In 1945, the Aluminum Corporation of America (Alcoa) was convicted of such a practice. An appellate panel of circuit court judges found that Alcoa had consistently installed more refining capacity than was justified by demand. In his opinion, Judge Learned Hand said:

It was not inevitable that it [Alcoa] should always anticipate increases in the demand for ingot and be prepared to supply them. Nothing compelled it to keep doubling and redoubling its capacity before others entered the field. It insists that it never excluded competitors; but we can think of no more effective exclusion than progressively to embrace each new opportunity as it opened and to face every newcomer with new capacity already geared into a great organization.

This case has been debated at length by scholars of antitrust law and economics.² Here we ask you to consider the conceptual basis of the case. How could the construction of excess capacity deter new competitors?

Case Discussion

An established firm wants to convince potential new competitors that the business would not be profitable for them. This basically means that if they entered, the price would be too low to cover their costs. Of course the established firm could simply put out the word that it would fight an unrelenting price war against any newcomers. But why would the newcomers believe such a verbal threat? After all, a price war is costly to the established firm too.

Installing capacity in excess of the needs of current production gives

credibility to the established firm's threat. When such capacity is in place, output can be expanded more quickly and at less extra cost. It remains only to staff the equipment and get the materials; the capital costs have already been incurred and are bygones. A price war can be fought more easily, more cheaply, and therefore more credibly.

ARMS ACROSS THE OCEAN

In the United States many homeowners own guns for self-defense. In Britain almost no one owns a gun. Cultural differences provide one explanation. The possibility of strategic moves provides another.

In both countries, a majority of homeowners prefer to live in an unarmed society. But they are willing to buy a gun if they have reason to fear that criminals will be armed.* Many criminals prefer to carry a gun as one of the tools of their trade.

The table below suggests a possible ranking of outcomes. Rather than assign specific monetary payoffs to each possibility, the outcomes are ranked 1, 2, 3, and 4 from best to worst for each side.

		Criminals	
		No guns	Guns
Homeowners	No guns	1 2	4 1
	Guns	2 4	3 3

If there were no strategic moves, we would analyze this as a game with simultaneous moves and use the techniques from chapter 3. We first look for dominant strategies. Since the criminals' grade in column 2 is always higher than that in a corresponding row in column 1, criminals have a dominant strategy: they prefer to carry guns whether or not homeowners are armed. Homeowners do not have a dominant strategy; they prefer to respond in kind. If criminals are unarmed, a gun is not needed for self-defense.

What is the predicted outcome when we model the game in this way? Following Rule 2, we predict that the side with a dominant strategy uses it; the other side chooses its best response to the dominant strategy of its opponent. Since Guns is the dominant strategy for criminals, this is their predicted course

of action. Homeowners choose their best response to Guns; they too will own a gun. The resulting equilibrium is ranked (3, 3), the third-best outcome for both parties.

In spite of their conflicting interests, the two sides can agree on one thing. They both prefer the outcome in which neither side carries guns (1, 2) to the case in which both sides are armed (3, 3). What strategic move makes this possible, and how could it be credible?

Case Discussion

Imagine for a moment that criminals are able to preempt the simultaneity and make a strategic move. They would commit not to carry guns. In this now sequential game, homeowners do not have to predict what criminals will do. They would see that the criminals' move has been made, and they are not carrying guns. Homeowners then choose their best response to the criminals' commitment; they too go unarmed. This outcome is ranked (1, 2), an improvement for *both* sides.

It is not surprising that criminals do better by making a commitment.* But homeowners are better off, too. The reason for the mutual gain is that both sides place a greater weight on the others' move than their own. Homeowners can reverse the criminals' move by allowing them to make an unconditional move.†

In reality, homeowners do not constitute one united player, and neither do criminals. Even though criminals as a class may gain by taking the initiative and giving up guns, any one member of the group can get an additional advantage by cheating. This prisoners' dilemma would destroy the credibility of the criminals' initiative. They need some way to bond themselves together in a joint commitment.

If the country has a history of strict gun control laws, guns will be unavailable. Homeowners can be confident that criminals will be unarmed. Britain's strict control of guns allows criminals to commit to work unarmed. This commitment is credible, as they have no alternative. In the United States, the greater prevalence of guns denies criminals an ability to commit to work unarmed. As a result, many homeowners are armed for self-defense. Both sides are worse off.

Clearly this argument oversimplifies reality; one of its implications is that criminals should support gun control legislation. Even in Britain, this commitment is difficult to maintain. The political strife over Northern Ireland had the indirect effect of increasing the availability of guns to the criminal

population. As a consequence, any commitment from criminals not to carry guns has begun to break down.

In looking back, note that something unusual happened in the transition from a simultaneous-move to a sequential-move game. Criminals chose to forego what was their dominant strategy. In the simultaneous-move game it was dominant for them to carry guns. In the sequential-move game, they chose not to. The reason is that in a sequential-move game, their course of action affects the homeowners' choice. Because of this interaction, they can no longer take the homeowners' response as beyond their influence. They move first, so their action affects the homeowners' choice. Carrying a gun is no longer a dominant strategy in the sequential representation of the game.

FOOLING ALL THE PEOPLE SOME OF THE TIME: THE LAS VEGAS SLOTS

Any gambling guide should tell you that slot machines are your worst bet. The odds are way against you. To counter this perception and encourage slot machine play, some Las Vegas casinos have begun to advertise the payback ratio for their machines—the fraction of each dollar bet returned in prize money. Going one step further, some casinos guarantee that they have machines that are set to a payback ratio greater than 1! These machines actually put the odds in your favor. If you could only find those machines and play them, you would expect to make money. The trick, of course, is that they don't tell you which machines are which. When they advertise that the average payback is 90 percent and that some machines are set at 120 percent, that also means that other machines must be set somewhere below 90 percent. To make it harder for you, there is no guarantee that machines are set the same way each day—today's favorable machines could be tomorrow's losers. How might you go about guessing which machines are which?

Case Discussion

Since this is our final case, we can admit that we do not have the answer—and if we did, we probably wouldn't share it. Nonetheless, strategic thinking can help you make a more educated guess. The trick is to put yourself into the casino owners' shoes. They make money only when people play the disadvantageous machines at least as much as the favorable or loose machines as they are known.

Is it really possible that the casinos could “hide” the machines that are offering the favorable odds? If people play the machines that pay out the most, won’t they find the best ones? Not necessarily, and especially not necessarily in time! The payoff of the machine is in large part determined by the chance of a jackpot prize. Look at a slot machine that takes a quarter a pull. A jackpot prize of \$10,000 with a 1 in 40,000 chance would give a payoff ratio of 1. If the casino raised the chance to 1 in 30,000, then the payoff ratio would be very favorable at 1.33. But people watching others play the machine would almost always see a person dropping quarter after quarter with no success. A natural conclusion would be that this is one of the least favorable machines. Eventually, when the machine pays its jackpot prize, it could be retooled and set at a lower rate.

In contrast, the least favorable machines could be set to pay back a small prize with a high frequency, and basically eliminate the hope of the big jackpot. Look at a machine set with a payback of 80 percent. If it provided a \$1 prize on roughly every fifth draw, then this machine would make a lot of noise, attracting attention and possibly more gamblers’ money. Are these the machines they put at the end of the aisles or near the buffet?

Perhaps the experienced slot players have figured all this out. But if so, you can bet that the casinos are just doing the reverse. Whatever happens, the casinos can find out at the end of the day which machines were played the most. They can make sure that the payoff patterns that attract the most play are actually the ones with the lower payoff ratio. For while the difference between a payoff ratio of 1.20 and 0.80 may seem large—and determines the difference between making money and losing money—it can be extremely hard to distinguish based on the number of pulls any one slot player can afford to make. The casinos can design the payoffs to make these inferences harder and even go the wrong way most of the time.

The strategic insight is to recognize that unlike the United Way, Las Vegas casinos are not in the business to give out money. In their search for the favorable machines, the majority of the players can’t be right. For if the majority of the people were able to figure it out, the casino would discontinue their offer rather than lose money. So, don’t wait in line. You can bet that the most heavily played machines are not the ones with the highest payback.

FURTHER READING

PIONEERING BOOKS are often enjoyable to read. In this spirit, we recommend John von Neumann and Oscar Morgenstern's *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1947), even though the mathematics may be hard to follow in places. Thomas Schelling's *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960) is more than just a pioneering book; it continues to provide instruction and insight.

For an entertaining exposition of zero-sum games, J. D. Williams's *The Compleat Strategyst*, rev. ed. (New York: McGraw-Hill, 1966) still cannot be beat. The most thorough and highly mathematical treatment of pre-Schelling game theory is in Duncan Luce and Howard Raiffa, *Games and Decisions* (New York: Wiley, 1957). Among general expositions of game theory, Morton Davis, *Game Theory: A Nontechnical Introduction*, 2nd ed. (New York: Basic Books, 1983), is probably the easiest to read.

In terms of biographies, surely the most famous book on game theory is Sylvia Nasar, *A Beautiful Mind: The Life of Mathematical Genius and Nobel Laureate John Nash* (New York: Touchstone, 2001). The book is even better than the movie. William Poundstone's *Prisoner's Dilemma* (New York: Anchor, 1993) goes beyond a description of the eponymous game to offer a first-rate biography of John von Neumann, the polymath who invented the modern computer along with game theory.

In terms of textbooks, we are naturally partial to two of our own. Avinash Dixit and Susan Skeath, *Games of Strategy*, 2nd ed. (New York: W. W. Norton & Company, 2004), is designed for undergraduates. Barry Nalebuff and Adam Brandenburger's *Co-opetition* (New York: Doubleday, 1996) offers an application of game theory for MBAs and managers more broadly.

Other excellent textbooks include Robert Gibbons, *Game Theory for Applied Economists* (Princeton, NJ: Princeton University Press, 1992); John McMillan, *Games, Strategies, and Managers: How Managers Can Use Game Theory to Make Better Business Decisions* (New York: Oxford University Press, 1996); Eric Rasmusen, *Games and Information* (London: Basil Blackwell, 1989); Roger B. Myerson, *Game Theory: Analysis of Conflict* (Cambridge, MA: Harvard University Press, 1997); Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory* (Cambridge, MA: MIT Press, 1994); and Martin J. Osborne, *An Introduction to Game Theory* (New York: Oxford University Press, 2003). We always look forward to Ken Binmore's books. *Playing for Real: A Text on Game Theory* (New York: Oxford University Press, 2007) is the much-anticipated revision of his *Fun and Games* (Lexington, MA: D. C. Heath, 1992). (Warning: The title is a bit misleading. The book is actually quite challenging, both conceptually and mathematically. But it is very rewarding for the well prepared.) Binmore's latest offering is *Game Theory: A Very Short Introduction* (New York: Oxford University Press, 2008).

The following books are much more advanced and are used largely in graduate courses. They are strictly for the very ambitious: David Kreps, *A Course in Microeconomic Theory* (Princeton, NJ: Princeton University Press, 1990) and Drew Fudenberg and Jean Tirole, *Game Theory* (Cambridge, MA: MIT Press, 1991).

One of our sins of omission is a discussion of "cooperative games." Here players choose and implement their actions jointly and produce equilibria like the Core or the Shapley Value. This was done because we think cooperation should emerge as the equilibrium outcome of a noncooperative game in which actions are chosen separately. That is, individuals' incentive to cheat on any agreement should be recognized and made a part of their strategy choice. Interested readers can find treatments of cooperative games in the books by Davis and by Luce and Raiffa mentioned above and more extensively in Martin Shubik's *Game Theory in the Social Sciences* (Cambridge, MA: MIT Press, 1982).

There are several terrific books applying game theory to specific contexts. One of the most powerful applications is to auction design. Here there is no better source than Paul Klemperer's *Auctions: Theory and Practice*, The

Toulouse Lectures in Economics (Princeton, NJ: Princeton University Press, 2004). Professor Klemperer was behind the design of many of the spectrum auctions, including the UK auction, which helped bring in some £34 billion and nearly bankrupted the telecom industry in the process. For game theory applied to law, see Douglas Baird, Robert Gertner, and Randal Picker, *Game Theory and the Law* (Cambridge, MA: Harvard University Press, 1998). One of their many contributions is the idea of an information escrow, which turns out to be a particularly useful tool in negotiation.* In the field of politics, noteworthy books include Steven Brams, *Game Theory and Politics* (New York: Free Press, 1979), and his more recent *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures* (Princeton, NJ: Princeton University Press, 2007); William Riker, *The Art of Political Manipulation* (New Haven, CT: Yale University Press, 1986); and the more technical approach of Peter Ordeshook's *Game Theory and Political Theory* (New York: Cambridge University Press, 1986). For applications to business, Michael Porter's *Competitive Strategy* (New York: Free Press, 1982); R. Preston McAfee's *Competitive Solutions: The Strategist's Toolkit* (Princeton, NJ: Princeton University Press, 2005); and Howard Raiffa's *The Art and Science of Negotiation* (Cambridge, MA: Harvard University Press, 1982) are excellent resources.

On the web, www.gametheory.net has the best collection of links to books, movies, and reading lists on game theory and its application.