

players sometimes don't believe big bets, and don't fold as a result. If their hand is better, the bluff comes straight out of the liar's pocket. So the central feature of a game like five-card draw is not figuring out the probability of drawing an inside straight or three of a kind, although that's certainly useful too. It's about convincing others that your hand is stronger than it really is. Part of the key to accumulating bargaining chips, whether in poker or diplomacy, is engineering the future by exploiting leverage that really does not exist. Along with taking prudent risks, creating leverage is one of the most important features in changing outcomes. Of course, that is just a polite way of saying that it's good to know when and how to lie.

Betting, whether with chips, stockholders' money, perjury charges, or soldiers, can lead others to wrong inferences that benefit the bettor; but gambling always suffers from two limitations. First, it can be expensive to bet more than a hand is worth. Second, everyone has an interest in trying to figure out who is bluffing and who is being honest. Raising the stakes helps flush out the bluffers. The bigger the cumulative bet, the costlier it is to pretend to have the resolve to see a dispute through when the cards really are lousy. How much pain anyone is willing to risk on a bluff, and how similar their wagering is when they are bluffing and when they are really holding good cards, is crucial to the prospects of winning or of being unmasked. That, of course, is why diplomats, lawyers, and poker players need a good poker face, and it is why, for example, you take your broker's advice more seriously if she invests a lot of her own money in a stock she's recommending.

Getting the best results comes down to matching actions to beliefs. Gradually, under the right circumstances, exploiting information leads to consistency between what people see, what they think, and what they do, just as it does in Mastermind. Convergence in thinking facilitates deals, bargains, and the resolution of disputes.

With that, we've just completed the introductory course in game theory. Nicely done! Now we're ready to go on to the more advanced course. In the next chapter we look in more depth at how the very fact of our being strategic changes everything going on around us. That will set the stage for working out how we can use strategy to change things to be better for ourselves and those we care about and, if we are altruistic enough, maybe even for almost everyone.

GAME THEORY 102

GAME THEORY 101 started us off thinking about how different people are from particles. In short, we are strategists. We calculate before we interact. And with 101 under our belts, we know enough to delve more closely into the subtleties of strategizing.

Of the many lessons game theory teaches us, one of particular import is that the future—or at least its anticipation—can cause the past, perhaps even more often than the past causes the future. Sound crazy? Ask yourself, do Christmas tree sales cause Christmas? This sort of reverse causality is fundamental to how game theorists work through problems to anticipate outcomes. It is very different from conventional linear thinking. Let me offer an example where failing to recognize how the future shapes the past can lead to really bad consequences.

Many believe that arms races cause war.¹ With that conviction in mind, policy makers vigorously pursue arms control agreements to improve the prospects of peace. To be sure, controlling arms means that if there is war, fewer people are killed and less property is destroyed. That is certainly a good thing, but that is not why people advocate arms control. They want to make war less likely. But reducing the amount or lethality of arms just does not do that.

The standard account of how arms races cause war involves what game theorists call a hand wave—that is, at some point the analyst waves his hands in the air instead of providing the logical connection from argument to conclusions. The arms-race hand wave goes like this:

When a country builds up its arms it makes its adversaries fear that their security is at risk. In response, they build up their own arms to defend themselves. The other side looks at that buildup—seeing their own as purely defensive—and tries to protect itself by developing still more and better weapons. Eventually the arms race produces a massive overcapacity to kill and destroy. Remember how many times over the U.S. and Soviet nuclear arsenals could destroy the world! So, as the level of arms—ten thousand nuclear-tipped missiles, for instance—grows out of proportion to the threat, things spiral out of control (that's the hand wave—*why do things spiral out of control?*), and war starts.

Wait a moment, let's slow down and think about that. The argument boils down to claiming that when the costs of war get to be really big—arms are out of proportion to the threat—war becomes more likely. That's really odd. Common sense and basic economics teach us that when the cost of anything goes up, we generally buy less, not more. Why should that be any less true of war?

True, just about every war has been preceded by a buildup in weapons, but that is not the relevant observation. It is akin to looking at a baseball player's positive test for steroids as proof that he cheats. What we want to know is how often the acquisition of lots more weapons leads to war, not how often wars are preceded by the purchase of arms. The answer to the question we care about is, *not very often*.

By looking at wars and then asking whether there had been an arms race, we confuse cause and effect. We ignore all the instances in which arms may successfully deter fighting exactly because the anticipated destruction is so high. Big wars are very rare precisely because when we expect high costs we look for ways to compromise. That, for instance, is why the 1962 Cuban Missile Crisis ended peacefully. That is why every major crisis between the United States and the Soviet Union throughout the cold war ended without the initiation of a hot war. The fear of nuclear annihilation kept it cold. That is why lots of events that could have ignited world wars ended peacefully and are now all but forgotten.

So, in war and especially in peace, reverse causality is at work. When policy makers turn to arms control deals, thinking they are promoting peace, they are taking much bigger risks than they seem to realize. Failing to think about reverse causation leads to poor predictions of what is likely

to happen, and that can lead to dangerous decisions and even to catastrophic war.

We will see many more instances of this kind of reasoning in later chapters. We will examine, for example, why most corporate fraud probably is not sparked by executive greed and why treaties to control greenhouse gas emissions may not be the best way to fight global warming. Each example reinforces the idea that correlation is not causation. They also remind us that the logic of reverse causation—called endogeneity in game theory—means that what we actually “observe”—such as arms races followed by war—are often biased samples.

The fact that decisions can be altered by the expectation of their consequences has lots of implications. In Game Theory 101 we talked about bluffing. Working out when promises or threats should be taken seriously and when they are (in game-theory-speak) “cheap talk” is fundamental to solving complicated situations in business, in politics, and in our daily encounters. Sorting out when promises or threats are sincere and when they are just talk is the problem of determining whether commitments are credible.

LET'S PLAY GAMES

In predicting and engineering the future, part of getting things right is working out what stands in the way of this or that particular outcome. Even after pots of money are won at cards, or hands are shaken and contracts or treaties are signed, we can't be sure of what will actually get implemented. We always have to ask about commitments. Deals and promises, however sincerely made, can unravel for lots of reasons. Economists have come up with a superbly descriptive label for a problem in enforcing contracts. They ask, is the contract “renegotiation-proof”?³ This question is at the heart of litigiousness in the United States.

I once worked on a lawsuit involving two power companies. One produced excess electricity and sold it to a different electric company in another state. As it happened, the price for electricity shot way up after the contract was signed. The contract called for delivery at an agreed-upon lower price. The power seller stopped delivering the promised electricity to

the buyer, demanding more money for it. Naturally, the buyer objected, pointing out that the contract did not provide for changing the price just because market conditions changed. That was a risk that the buyer and seller agreed to take when they signed their contract. Still, the seller refused to deliver electricity. The seller was sued and defended itself vigorously so that legal costs racked up on both sides. All the while that bitter accusations flew back and forth, the seller kept offering to make a new deal with the plaintiff. The deal involved renegotiating their contract to make adjustments for extreme changes in market prices. The plaintiff resisted, always pointing—rightly—to the contract. But the plaintiff also really needed the electricity and couldn't get it anywhere else for a better price than the seller, my client, was willing to take—and my client knew that. Eventually, the cost of not providing the necessary electricity to their own clients became so great that the plaintiff caved in and took the deal they were offered.

Here was nasty, avaricious human nature hard at work in just the way game theorists think about it. Yes, there was a contract, and its terms were clear enough, but the cost of fighting to enforce the contract became too great. However much the plaintiff declared its intent to fight the case in court, the defendant knew it was bluffing. The plaintiff's need for electricity and the cost of battling the case out in court were greater than the cost of accepting a new deal. And so it was clear that the terms of the contract were not renegotiation-proof. The original deal was set aside and a new one was struck. The original deal really was not a firm commitment to sell (or probably, for that matter, to buy) electricity at a specified price over a specified time period when the market price moved markedly from the price stipulated in the agreement. Justice gave way, as it so often does in our judicial system, to the relative ability of plaintiffs and defendants to endure pain.

Commitment problems come in other varieties. The classic game theory illustration of a commitment problem is seen in the game called the prisoner's dilemma, which is played out on almost every cop show on TV every night of the week. The story is that two criminals (I'll call them Chris and Pat) are arrested. Each is held in a separate cell, with no communication between them. The police and the DA do not have enough evidence to convict them of the serious crime they allegedly committed. But they do have enough evidence to convict them of a lesser offense. If

Chris and Pat cooperate with each other by remaining silent, they'll be charged and convicted of the lesser crime. If they both confess, they'll each receive a stiff sentence. However, if one confesses and the other does not, then the one who confesses—ratting out the other—will get off with time served, and the other will be put away for life without a chance for parole.

It is possible, maybe even likely, that Chris and Pat, our two crooks, made a deal beforehand, promising to remain silent if they are caught. The problem is that their promise to each other is not credible because it's always in their interest—if the game is not going to be repeated an indefinite number of times—to renege, talking a blue streak to make a deal with the prosecutor. Here's how it works:

THE PRISONER'S DILEMMA

Pat's Choices→	Don't confess (stay faithful to Chris)	Confess (rat out Chris)
Chris's Choices ↓		
Don't confess (stay faithful to Pat)	Chris and Pat get 5 years	Chris gets life; Pat gets time served
Confess (rat out Pat)	Chris gets time served; Pat gets life	Chris and Pat get 15 years

After Chris and Pat are arrested, neither knows whether the other will confess or really will stay silent as promised. What Chris knows is that if Pat is true to his word and doesn't talk, Chris can get off with time served by betraying Pat. If instead Chris stays faithful to her promise and keeps silent too, she can expect to get five years. Remember, game theory reasoning takes a dim view of human nature. Each of the crooks looks out for numero uno. Chris cares about Chris; Pat looks out only for Pat. So if Pat is a good, loyal buddy—that is, a sucker—Chris can take advantage of the chance she's been given to enter a plea. Chris would walk and Pat would go to prison for life.

Of course, Pat works out this logic too, so maybe instead of staying silent, Pat decides to talk. Even then, Chris is better off confessing than

she would be by keeping her mouth shut. If Pat confesses and Chris stays silent, Pat gets off easy—that's neither here nor there as far as Chris is concerned—and Chris goes away for a long time, which is everything to her. If Chris talks too, her sentence is lighter than if she stayed silent while Pat confessed. Sure, Chris (and Pat) gets fifteen years, but Chris is young, and fifteen years, with a chance for parole, certainly beats life in prison with no chance for parole. In fact, whatever Chris thinks Pat will do, Chris's best bet is to confess.

This produces the dilemma. If both crooks kept quiet they would each get a fairly light sentence and be better off than if both confessed (five years each versus fifteen). The problem is that neither one benefits from taking a chance, knowing that it's always in the other guy's interest to talk. As a consequence, Chris's and Pat's promises to each other notwithstanding, they can't really commit to remaining silent when the police interrogate them separately.

IT'S ALL ABOUT THE DOG THAT DIDN'T BARK

The prisoner's dilemma illustrates an application of John Nash's greatest contribution to game theory. He developed a way to solve games. All subsequent, widely used solutions to games are offshoots of what he did. Nash defined a game's equilibrium as the planned choice of actions—the strategy—of each player, requiring that the plan of action is designed so that no player has any incentive to take an action not included in the strategy. For instance, people won't cooperate or coordinate with each other unless it is in their individual interest. No one in the game-theory world willingly takes a personal hit just to help someone else out. That means we all need to think about what others would do if we changed our plan of action. We need to sort out the "what ifs" that confront us.

Historians spend most of their time thinking about what happened in the world. They want to explain events by looking at the chain of things that they can observe in the historical record. Game theorists think about what *did not* happen and see the anticipated consequences of what didn't happen as an important part of the cause of what did happen. The central characteristic of any game's solution is that each and every player expects to be worse off by choosing differently from the way they did. They've

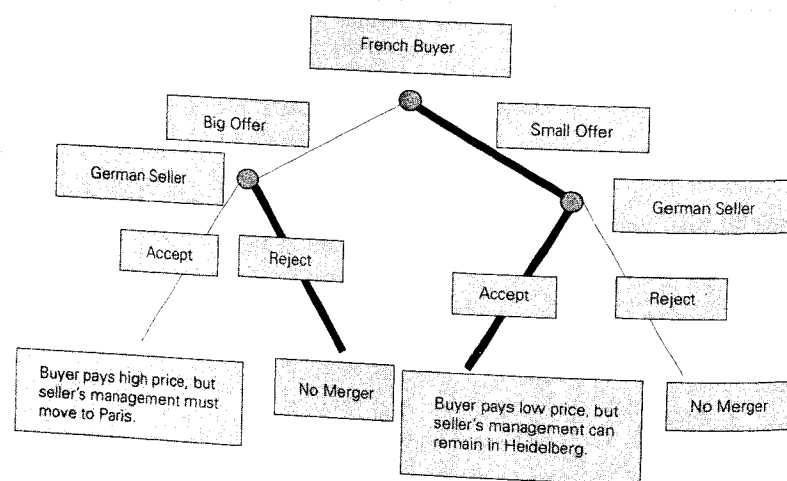
pondered the counterfactual—what would my world look like if I did this or I did that?—and did whatever they believed would lead to the best result for them personally.

Remember the very beginning of this book, when we pondered why Leopold was such a good king in Belgium and such a monster in the Congo? This is part of the answer. The real Leopold would have loved to do whatever he wanted in Belgium, but he couldn't. It was not in his interest to act like an absolute monarch when he wasn't one. Doing some counterfactual reasoning, he surely could see that if he tried to act like an absolute ruler in Belgium, the people probably would put someone else on the throne or get rid of the monarchy altogether, and that would be worse for him than being a constitutional monarch. Seeing that prospect, he did good works at home, kept his job, and freed himself to pursue his deepest interests elsewhere. Not facing such limitations in the Congo, there he did whatever he wanted.

This counterfactual thinking becomes especially clear if we look at a problem or game as a sequence of moves. In the prisoner's dilemma table I showed what happens when the two players choose without knowing what the other will do. Another way to see how games are played is to draw a tree that shows the order in which players make their moves. Who gets to move first matters a lot in many situations, but it does not matter in the prisoner's dilemma because each player's best choice of action is the same—confess—whatever the other crook does. Let's have a look at a prospective corporate acquisition I worked on (with the details masked to maintain confidentiality). In this game, anticipating what the other player will do is crucial to getting a good outcome.

The buyer, a Paris-based bank, wanted to acquire a German bank. The buyer was prepared to pay a big premium for the German firm but was insistent on moving all of the German executives to the corporate headquarters in Paris. As we analyzed the prospect of the acquisition, it became apparent that the price paid was not the decisive element for the Heidelberg-based bank. Sure, everyone wanted the best price they could get, but the Germans loved living in Heidelberg and were not willing to move to Paris just for money. Paris was not for them. Had the French bankers pushed ahead with the offer they had in mind, the deal would have been rejected, as can be seen in the game tree below. But because their attention was drawn to the importance the Germans attached to

FIG. 3.1. Pay Less to Buy a Bank



where they lived, the offer was changed from big money to a more modest amount—fine enough for the French—but with assurances that the German executives could remain in Heidelberg for at least five years, which wasn't ideal for the French, but necessary for their ends to be realized.

The very thick, dark lines in the figure show what the plans of action were for the French buyer and the German seller. There is a plan of action for every contingency in this game. One aspect of the plan of action on the part of the executives in Heidelberg was to say *nein* to a big-money offer that required them to move to Paris. This never happened, exactly because the French bankers asked the right "what if" question. They asked, What happens if we make a big offer that is tied to a move to Paris, and what happens if we make a more modest money offer that allows the German bank's management to stay in Heidelberg? Big money in Paris, as we see with the thick, dark lines, gets *nein* and less money in Heidelberg encourages the seller to say *ja wohl*. Rather than not make the deal at all, the French chose the second-best outcome from their point of view. They made the deal that allowed the German management to stay put for five years. The French wisely put themselves in their German counterparts' shoes and acted accordingly.

By thinking about the strategic interplay between themselves and the

German executives, the French figured out how to make a deal they wanted. They concentrated on the all-important question, "What will the Germans do if we insist they move to Paris?" No one actually moved to Paris. Historians don't usually ask questions about things that did not happen, so they would probably overlook the consequences of an offer that insisted the German management relocate to France. They might even wonder why the Germans sold so cheaply. In the end, the Germans stayed in Heidelberg.

Why should we care about their moving to Paris when in fact they didn't? The reason they stayed in Heidelberg while agreeing to the merger is precisely because of what would have happened had the French insisted on moving them to France: no deal would have been struck, and so there would have been no acquisition for anyone to study.

The two games I have illustrated in the preceding pages are very simple. They involve only two players, and each game has only one possible rational pair of strategies leading to an equilibrium result. Even a simple two-player game, however, can involve more than one set of sensible plans of action that lead to different possible ends of the game. We'll solve an example of such a game in the last chapter. Of course, with more players and more choices of actions, many complicated games involve the possibility of many different strategies and many different outcomes. Part of my task as a consultant is to work out how to get players to select strategies that are more beneficial for my client than some other way of playing the game. That's where trying to shape information, beliefs, and even the game itself become crucial, and in this next section I'd like to show you just what I mean.

WANT TO BE A CEO?

As we all know, great jobs are getting harder to come by, and reaching the top is as competitive as ever. Merit may be necessary, but, as many of us can attest, it's unlikely to be sufficient. There are, after all, many more well-qualified people than there are high-level jobs to fill.

That being said, even if you've managed to mask or overcome your personal limitations and have been blessed with great timing and good luck such that you now find yourself in the rarefied air of the boardroom,

there's something worth knowing that might have escaped you, something that might still prevent you from grabbing that cherished top spot: the selection process.

That's right, understanding and shaping the process by which a CEO or other leaders are chosen can tip the competition in your favor. It's funny that few of us pay much attention, in a strategic sense, to something as prosaic as how votes are counted, whether in the boardroom or national elections. And yet the method used to translate what people want into what they get can turn a losing candidacy into a winning one.³

When I talk about shaping outcomes based on voting, I don't mean anything like miscounting or cheating. I don't mean relying on hanging chads or anything like that. I'm just thinking about the many regular, commonly used ways of arriving at a choice based on what voters or shareholders or board members want.

Few board members or shareholders pause to think about how the votes are going to be counted when they select a new CEO. Hardly anyone asks whether it really matters if we require a candidate to get a majority or a plurality; if we count just votes for people's first choice or we allow them to express their first and second (or even more) preferences; if in decisions with many candidates we vote on all of them at once or we pair them up in head-to-head contests. And yet you can bet your bottom dollar that these decisions really can change the results.

Just think back to the hotly contested 2008 Democratic Party primaries. The Democrats allocated delegates from each state roughly proportionally to the candidates based on their share of the popular vote in each primary. Barack Obama won a majority of delegates that way, and was ultimately elected president. If the Democrats had used the Republicans' winner-take-all rule in each primary, Hillary Clinton would have won enough delegates to be the nominee, and she too probably would have gone on to beat John McCain. That's a pretty big consequence of a seemingly inconsequential rule.

There is, of course, no right way to count votes. Every method has advantages and disadvantages. So we might as well use voting rules when we can to help the candidates we favor. Generally we don't have the opportunity to change how votes are counted in government elections, but we sure do when it comes to corporate decisions.

In fact, I've twice used the range of boardroom voting procedures to

help shape corporate choices of CEOs. Once the effort was entirely successful, and the second time, well, the candidate that my partner and I helped rose from obscurity to be treated as a very serious candidate. He ultimately lost, but he did so much better than anyone expected that he was quickly hired away from his company to become the CEO of a different firm—and he was a great success there.

How was the CEO selection process modeled? Let's take a look at my first experience on this front (it has the nice feature that even the person who was chosen as CEO didn't know—and probably still does not know—how he won). Here's what happened:

The retiring CEO of the company in question—certainly it must be obvious that there is a need for anonymity here—didn't have strong feelings about who he wanted to replace him. He did, however, have very strong feelings about who he did *not* want to replace him, and it just so happens that the person he didn't want was the leading candidate for succession. The retiring CEO truly despised this person, who had been his nemesis for many years, and he hired me in secret to help engineer the CEO selection. The modeling job: figure out how to beat the detested front-runner.

As with any analysis, the first step was to figure out just what were the real issues that had to be resolved. In this case, the big questions were simple enough. They required working out who the prospective candidates were and how they stacked up against each other. Let's call the candidates Larry, Moe, Curly, Mutt, and Jeff—with Mutt being the guy to beat.

The problem was best analyzed with a bunch of beauty contests. Each beauty contest asked how the stakeholders with a say in selecting the CEO felt about Larry vs. Moe, Larry vs. Curly, Larry vs. Mutt, Larry vs. Jeff, Moe vs. Curly, and so forth.

Once the issues—the beauty contests—are specified, we need to know each stakeholder's position, or which candidate he or she favors in each head-to-head contest, and by how much. (I will go into further detail in subsequent chapters as to the particular methodology behind these assessments.)

Thankfully, in this particular case, we have a good source of information: the outgoing CEO. He knew who the players were and he knew how they felt about each candidate. And you can be assured he didn't get to be CEO without knowing which of his colleagues had real clout and who would just go along with the wishes of others.

The existing procedure for succession in the CEO's company did not involve head-to-head contests, ranking of candidates, runoffs, or a host of other common voting rules. Instead, they normally voted for all the candidates simultaneously, just as is done in American presidential elections. Whoever got the most votes would be the winner. Now that would have been very bad news for my client, the retiring CEO. It was clear that such a procedure, a fine, upstanding, legitimate procedure, would result in the election of Mutt, the detested candidate. What to do?

The first thing was to sort out who was likely to win each of the beauty contests. The stakeholders consisted of the members of the company's CEO selection committee. Let's say the committee was made up of fifteen individuals, with one vote each. The outgoing CEO's information on the comparisons of candidates in pairs allowed me to tease out the strict order in which different committee members preferred the candidates. It was apparent that the fifteen committee members divided equally into five voting blocs based on their preferences, with three members in each group. Here are the five different preference orderings held by the members of the selection committee, with candidates listed from most preferred to least preferred:

1. Mutt, Jeff, Larry, Curly, Moe
2. Mutt, Moe, Curly, Larry, Jeff
3. Moe, Mutt, Curly, Larry, Jeff
4. Jeff, Moe, Curly, Larry, Mutt
5. Larry, Jeff, Curly, Moe, Mutt

In a contest in which everyone got to vote just one time, such as is used in the United States to pick the president, the detested Mutt would get 6 votes (two blocs held him in first place and each had three members), Moe 3, Jeff 3, Larry 3, and poor Curly 0. Mutt wins. That is exactly what had to be prevented.

However, under another voting system, if committee members got to cast 4 votes for their first choice candidate, 3 for second, 2 for third, 1 for fourth, and no points for their last choice (a method known as a Borda count), then Mutt and Moe would receive 33 weighted votes each, Jeff would get 30, and Larry and Curly would bring up the rear with 27 each. If they then held a tie-breaking runoff between Mutt and Moe, Moe

would pick up the votes of the third, fourth, and fifth blocs. They each favored him over Mutt. Moe would be the new CEO by this procedure.

So already we can see that there is a rule that could beat Mutt. However, it was clear enough that this procedure would be tough to get through the committee. It was just too complicated to ask members first to rank candidates and then to hold a runoff once they discovered there was a tie. Such a convoluted election process would easily arouse suspicion among committee members. They might have wondered why the retiring CEO was asking them to do something so elaborate when they could just vote straight up for any candidate they wanted.

Even if this complicated procedure could get committee approval, we would not have been home free. The procedure itself might be thwarted if a Mutt supporter caught on. For instance, if even one member of the second bloc worked out the results ahead of time, that member—who really wanted Mutt—could strategically (that is, by lying) decide to rank Jeff second and Moe last. Sure, that would have been a misrepresentation, but then the voter's interest was in the final choice, not any intermediate decision. Acting strategically by inflating Jeff's ranking, the bloc member would have ensured that Moe's total in the rankings would come to 30 instead of 33, and so would Jeff's, greatly complicating the process by creating a tie for second place and increasing the odds that Mutt would win. After all, in a runoff against Jeff (likely since more blocs preferred Jeff to Moe than the other way around), Mutt would be the winner. By acting strategically, then, it was possible for one or more members of the second bloc to ensure the election of their most preferred candidate, the detested Mutt. That was a chance I wasn't willing to take. Instead I decided to get Curly elected.

Poor Curly—he was at a huge disadvantage. No one viewed him as their first-place choice. Nobody even thought of him as second choice. In fact, he was barely on anyone's radar screen. I know that for sure because after working out how to get him elected, I had a conversation with a member of the selection committee. My role in the process was a secret, known only by the then CEO and me. I asked the committee member who he thought would be chosen, and he mentioned Mutt and maybe Moe. I nonchalantly asked about Jeff, Larry . . . and Curly. He took Jeff and Larry seriously, although he didn't think they could win. Then he told me that neither he nor anyone else on the committee understood why

Curly had put himself forward. After all, he said, he just doesn't have a prayer, no one favors him. Sure, they liked him well enough, but they just didn't seem to think of him as CEO material. Curly's relative obscurity was, in this case, his great advantage. It was unlikely that anyone paid enough attention to Curly's candidacy to maneuver strategically to thwart his prospects, since they didn't think he had any.

Okay, so now the fun begins. The outgoing CEO was well liked and highly respected. He had done a good job. The beauty contests revealed enough to show how to get Curly elected (do you see how?), but I needed to analyze one more issue first. The question was whether the retiring CEO had enough clout to persuade the selection committee to follow the winning voting procedure. The analysis of that question showed that indeed he could get the committee to follow the voting rule he suggested, provided it wasn't too complicated. Fortunately, the procedure my analysis suggested was an eminently reasonable rule. It wasn't particularly complicated, and it capitalized on the committee's majority not being keen to elect Mutt in the first place (remember, he had 6 first-place votes; 9 first-place votes were distributed among the others).

Agenda control—determining the order of decision making—can be everything. In this case it was. By setting the right agendas we could create a series of winning coalitions, each made up of different members from the one before, ending with a winning coalition supporting Curly and leaving no other candidate up for consideration.

The committee members understood that the real contest was between Mutt and Moe—or so they thought. To reinforce their view, the outgoing CEO persuaded the committee to use an agenda—a sequence of choices—that was made up of a specific sequence of head-to-head elimination contests. Of course there were too many candidates to ask the committee to compare each candidate to each other candidate, two at a time. That would have meant ten votes. Instead, the retiring CEO persuaded the committee to vote on Mutt versus Moe, with the loser of that contest being eliminated from consideration and the winner then going up against Jeff. Whoever lost that contest would be dropped from consideration, and the winner (who at this point in principle could have been Mutt or Moe or Jeff) would then be voted on against Larry, and the winner of that vote would finally be voted on against Curly. Whoever was left standing after those four votes would be deemed the winner.

This seemed like a good idea to the selection committee. They thought that by leading with strength—Mutt vs. Moe—they would quickly arrive at the one of those two who overall was most desired as CEO. How wrong they were. To be sure, anyone paying close attention to the five voting blocs' preferences could have worked out how the retiring CEO's agenda would work out, but it was unlikely that the committee members knew the full preference ordering of their compatriots. They, after all, were unlikely to conduct the sort of expert interviews called for by the model. Since they weren't asked to announce their candidate rankings, the rule proposed by the retiring CEO did not compel them to reveal to each other their full ranking of candidates. Probably on their own they had not probed one another beyond second-place preferences. That, presumably, was why they paid so little attention to Curly. So here is what happened:

Moe beat Mutt right off the bat—by a vote of 9 to 6 (blocs 1 and 2 voting for Mutt and the rest voting for Moe, as you can see from the rankings listed earlier). Mutt, being the loser, was dropped from consideration under the seemingly reasonable supposition that more people wanted Moe than Mutt (9, as we saw, to 6). Fair enough. Everything after that was gravy, because my client's main concern was to beat Mutt. But then my client also liked the idea of choosing Curly. He thought that would make him look even better in retrospect, and besides, he was fond of Curly and thought being CEO would be a nice way to cap Curly's career.

The selection committee then considered Moe and Jeff in accordance with the agreed-upon agenda. Jeff beat Moe as handily as Moe had beaten Mutt. Bloc 1 wanted Mutt most of all, but now, confronted with a choice between Moe and Jeff, they went for Jeff. He was their second-place choice, while Moe came in last for bloc 1. Blocs 4 and 5 also thought Jeff was a better prospective CEO than Moe. Only blocs 2 and 3 favored Moe over Jeff. That gave Jeff 9 votes to Moe's 6. Mutt having already been eliminated, no one on the committee stepped back to ask what would happen if they took the opportunity to choose between Jeff and Mutt. As you can see from the bloc preferences, there was yet another winning coalition (blocs 1, 2, and 3) with whose support Mutt would have beaten Jeff, but again, Mutt had been taken out of the picture by Moe in accordance with the elimination rules agreed to.

Mutt and Moe, the apparent front-runners, were now out of the race. Moe beat Mutt, and Jeff beat Moe. Jeff, Larry, and Curly were still stand-

ing. Jeff and Larry each had first-place supporters, so they were run against each other next. Blocs 2, 3, and 5 favored Larry over Jeff. Jeff was out, leaving a final choice between Larry and Curly. Of course you can easily see that Curly is going to defeat Larry. Blocs 2, 3, and 4, in the ever-shifting winning coalition of voters, favored Curly over Larry. Curly, being the last man standing, was the new CEO much to (almost) everyone's surprise. Still, they felt the process had been fair and square, and in its own way it was.

No one seemed to notice that the agenda had decided the outcome. It so happens that Larry was the only candidate that Curly could have beaten. If the agenda had been different, Curly would have lost. Just as Curly could not beat anyone other than Larry, so too could Larry not defeat anyone other than Jeff. Moving Larry up in the agenda would have wiped him out and Curly with him. In fact, because preferences went around in circles (or, to put it technically, they were intransitive), an agenda could be put forward to make any of the candidates into the winner fair and square.

When the vote was over, the committee member with whom I had talked earlier invited me to lunch. He had one question for me: Did you have anything to do with picking our CEO? I smiled and changed the subject. He was sure I did, and I knew he knew, but I was sworn to secrecy. The lunch was great.

SORRY, EINSTEIN: GOD DOES ROLL THE DICE

As we see in moving from the prisoner's dilemma to the bank example to the voting strategies, even in relatively simple games involving relatively few players there can be multiple outcomes. This fact adds yet another strategic dimension to decision making, particularly as games, in the real world, are often played over and over between the same players.

Any time a game has more than one possible result, there is a special type of strategy (called a mixed strategy) that can influence what happens. In a mixed strategy, each player chooses actions probabilistically—say, by rolling dice—to influence what other players expect to get out of the game. Einstein's God may not have played dice with the universe, but we mortals definitely roll the dice with each other.

Whenever you watch a football game and complain about a coach's choice of plays, you probably were watching a mixed strategy at work. For instance, when the ball is on the one-yard line, the play that is most likely to get the ball across the goal line is for the fullback to jump over the pile of players in front of him. Yet coaches often have the quarterback pass the ball or hand off to a running back. The reason: if a coach always called for the fullback to go over the top, the defenders would concentrate the defense at that point, and the play would probably fail. By mixing the calls, the offense forces the defense to spread out, thereby improving the odds of success over repetitions of the situation. Interestingly, this sort of mixing of strategies carries important lessons for business, politics, and lots of other parts of life. Rolling the dice is one way to alter how other people perceive a situation.

Using strategies that involve mixing up moves to create a change in expectations is something that comes up all the time. Although applied game theorists often like to ignore these complicated "mixed strategy" approaches to problems, they do so at their own peril. Rolling the dice can really make a difference in how things turn out.

Examples of such gambling are all around us, and some great movies roll the dice very cleverly to create climactic moments. Who can forget in *The Princess Bride* the back-and-forth over which wineglass is poisoned, and the clever resolution (both were poisoned—it pays to build up an immunity to the poison you plan to use on yourself and others). Or how about the fabulous scene from *The Maltese Falcon* in which Sydney Greenstreet's character, Kasper Gutman, desperately wants the jewel-encrusted bird? Only Sam Spade (Humphrey Bogart) knows where it is, and Sam Spade is no fool. Gutman threatens that Joel Cairo (Peter Lorre) will torture Spade to find out where the bird is, but Spade counters, "If you kill me, how are you going to get the bird? If I know you can't afford to kill me till you have it, how are you going to scare me into giving it to you?" Here Spade, like any good game theorist, questions the credibility of Gutman's commitment to make him talk. We know and Sam Spade knows that, without a real commitment to kill him, Gutman can't get him to talk. But Gutman is no fool either. He knows just how to make the dice tumble, creating the prospect that Spade will talk to save himself. After some clever give-and-take, Gutman retorts, "As you know, sir, men are likely to forget in the heat of action where their best interests lie and let their emotions carry them away."

There it is: "Men are likely to forget in the heat of action where their best interests lie and let their emotions carry them away." How beautifully put. He's just explained that Joel Cairo will try to be careful not to kill Spade, but then, Cairo can get emotional, so there is also a chance that if Spade doesn't talk he'll end up dead. In this brief exchange we see three lovely principles of game theory at work: the question of credible commitment; the use of playing probabilistically to alter how others look at the situation; and the pretense of irrationality (the heat of the moment) for strategic advantage. What could be truer to life's fears and calculations? How many of us would dare to stay silent given Sam Spade's gamble: keep the bird and maybe die, or give up the falcon and (maybe) live?

With a bit of luck, it will become apparent that game theory is not limited to parlor tricks, movie scripts, and brainteasers. It is a powerful tool for reshaping the world. In the remaining chapters, we will use these foundations to see just what kind of problems rational choice theory can tackle, and how math, science, and technology now allow us to predict and engineer particular outcomes that we might otherwise assume would only be determined by a random mix of good or bad fortune and a heavy dose of human whim.

BOMBS AWAY

WE NOW UNDERSTAND the basic principles of game theory. That's all well and good, but how do we use these principles to solve the big problems of our time? Well, let's not mess around—let's take a look at North Korean disarmament.

Being born in North Korea pretty much ensures having a miserable life. The average North Korean has to work an entire year to earn what an average American, Irish, or Norwegian worker makes in around four days. Money isn't everything, but it is a pretty good first-cut approximation of quality of life.

Of course, being born in North Korea is not always bad news for everyone. Kim Jong Il, that country's "Dear Leader," seems to do very well indeed. He is estimated to be worth around \$4 billion. That's about one-third of North Korea's annual gross domestic product. (Poor Bill Gates, his wealth is only about 0.4 percent of the U.S. gross domestic product.)

Kim's wealth, womanizing, heavy drinking, and gourmand tastes incline many to think of him as an inconsequential, even frivolous, lightweight dictator. He is frequently described as irrational, erratic, and dangerous. Most assuredly he is the last of those, but irrational and lightweight—I think not. Sure, he came by his position the old-fashioned way—he inherited it (on the death of his father, the "Great Leader," Kim Il Sung). Still, that was a long time ago. No fool stays in power for years on end when there are so many generals, sons, and wives waiting in the wings to launch a coup.