

Games Solvable by Backward Reasoning

IT'S YOUR MOVE, CHARLIE BROWN

In a recurring theme in the comic strip *Peanuts*, Lucy holds a football on the ground and invites Charlie Brown to run up and kick it. At the last moment, Lucy pulls the ball away. Charlie Brown, kicking only air, lands on his back, and this gives Lucy great perverse pleasure.

Anyone could have told Charlie that he should refuse to play Lucy's game. Even if Lucy had not played this particular trick on him last year (and the year before and the year before that), he knows her character from other contexts and should be able to predict her action.

At the time when Charlie is deciding whether or not to accept Lucy's invitation, her action lies in the future. However, just because it lies in the future does not mean Charlie should regard it as uncertain. He should know that of the two possible outcomes—letting him kick and seeing him fall—Lucy's preference is for the latter. Therefore he should forecast that when the time comes, she is going to pull the ball away. The logical possibility that Lucy will let him kick the ball is realistically irrelevant. Reliance on it would be, to borrow Dr. Johnson's characterization of remarriage, a triumph of hope over experience. Charlie should disregard it, and forecast that acceptance will inevitably land him on his back. He should decline Lucy's invitation.

TWO KINDS OF STRATEGIC INTERACTIONS

The essence of a game of strategy is the interdependence of the players' decisions. These interactions arise in two ways. The first is *sequential*, as in the Charlie Brown story. The players make alternating moves. Charlie, when it is his turn, must look ahead to how his current actions will affect the future actions of Lucy, and his own future actions in turn.

The second kind of interaction is *simultaneous*, as in the prisoners' dilemma tale of chapter 1. The players act at the same time, in ignorance of the others' current actions. However, each must be aware that there are other active players, who in turn are similarly aware, and so on. Therefore each must figuratively put himself in the shoes of all and try to calculate the outcome. His own best action is an integral part of this overall calculation.

When you find yourself playing a strategic game, you must determine whether the interaction is simultaneous or sequential. Some games, such as football, have elements of both, in which case you must fit your strategy to the context. In this chapter, we develop, in a preliminary way, the ideas and rules that will help you play sequential games; simultaneous-move games are the subject of chapter 3. We begin with really simple, sometimes contrived, examples, such as the Charlie Brown story. This is deliberate; the stories are not of great importance in themselves, and the right strategies are usually easy to see by simple intuition, allowing the underlying ideas to stand out much more clearly. The examples get increasingly realistic and more complex in the case studies and in the later chapters.



The First Rule of Strategy

The general principle for sequential-move games is that each player should figure out the other players' future responses and use them in calculating his own best current move. This idea is so important that it is worth codifying into a basic rule of strategic behavior:

RULE 1: Look forward and reason backward.

Anticipate where your initial decisions will ultimately lead and use this information to calculate your best choice.


In the Charlie Brown story, this was easy to do for anyone (except Charlie Brown). He had just two alternatives, and one of them led to Lucy's decision between two possible actions. Most strategic situations involve a longer sequence of decisions with several alternatives at each. A tree diagram of the choices in the game sometimes serves as a visual aid for correct reasoning in such games. Let us show you how to use these trees.

DECISION TREES AND GAME TREES

A sequence of decisions, with the need to look forward and reason backward, can arise even for a solitary decision maker not involved in a game of strategy with others. For Robert Frost in the yellow wood:

Two roads diverged in a wood, and I—
I took the road less traveled by,
And that has made all the difference.¹

We can show this schematically.




This need not be the end of the choice. Each road might in turn have further branches. The road map becomes correspondingly complex. Here is an example from our own experience.

Travelers from Princeton to New York have several choices. The first decision point involves selecting the mode of travel: bus, train, or car. Those who drive then have to choose among the Verrazano-Narrows Bridge, the Holland Tunnel, the Lincoln Tunnel, and the George Washington Bridge. Rail commuters must decide whether to switch to the PATH train at Newark or continue to Penn Station. Once in New York, rail and bus commuters must choose among going by foot, subway (local or express), bus, or taxi to get to their final destination. The best choices depend on many factors, including price, speed, expected congestion, the final destination in New York, and one's

aversion to breathing the air on the New Jersey Turnpike.

This road map, which describes one's options at each junction, looks like a tree with its successively emerging branches—hence the term. The right way to use such a map or tree is not to take the route whose first branch looks best—for example, because you would prefer driving to taking the train when all other things are equal—and then “cross the Verrazano Bridge when you get to it.” Instead, you anticipate the future decisions and use them to make your earlier choices. For example, if you want to go downtown, the PATH train would be superior to driving because it offers a direct connection from Newark.




We can use just such a tree to depict the choices in a game of strategy, but one new element enters the picture. A game has two or more players. At various branching points along the tree, it may be the turn of different players to make the decision. A person making a choice at an earlier point must look ahead, not just to his own future choices but to those of others. He must forecast what the others will do, by putting himself figuratively in their shoes, and thinking as they would think. To remind you of the difference, we will call a tree showing the decision sequence in a game of strategy a *game tree*, reserving *decision tree* for situations in which just one person is involved.

Charlie Brown in Football and in Business


The story of Charlie Brown that opened this chapter is absurdly simple, but

you can become familiar with game trees by casting that story in such a picture. Start the game when Lucy has issued her invitation, and Charlie faces the decision of whether to accept. If Charlie refuses, that is the end of the game. If he accepts, Lucy has the choice between letting Charlie kick and pulling the ball away. We can show this by adding another fork along this road.



As we said earlier, Charlie should forecast that Lucy will choose the upper branch. Therefore he should figuratively prune the lower branch of her choice from the tree. Now if he chooses his own upper branch, it leads straight to a nasty fall. Therefore his better choice is to follow his own lower branch. We show these selections by making the branches thicker and marking them with arrowheads.

Are you thinking that this game is too frivolous? Here is a business version of it. Imagine the following scenario. Charlie, now an adult, is vacationing in the newly reformed formerly Marxist country of Freedonia. He gets into a conversation with a local businessman named Fredo, who talks about the wonderful profitable opportunities that he could develop given enough capital, and then makes a pitch: “Invest \$100,000 with me, and in a year I will turn it into \$500,000, which I will share equally with you. So you will more than double your money in a year.” The opportunity Fredo describes is indeed attractive, and he is willing to write up a proper contract under Freedonian law. But how secure is that law? If at the end of the year Fredo absconds with all the money, can Charlie, back in the United States, enforce the contract in Freedonian courts? They may be biased in favor of their national, or too slow, or bribed by Fredo. So Charlie is playing a game with Fredo, and the tree is as shown here. (Note that if Fredo honors the contract, he pays Charlie \$250,000; therefore Charlie’s profit is that minus the initial investment of \$100,000—that is, \$150,000.)



What do you think Fredo is going to do? In the absence of a clear and strong reason to believe his promise, Charlie should predict that Fredo will abscond, just as young Charlie should have been sure that Lucy would pull the ball away. In fact the trees of the two games are identical in all essential respects. But how many Charlies have failed to do the correct reasoning in such games?

What reasons can there be for believing Fredo's promise? Perhaps he is engaged in many other enterprises that require financing from the United States or export goods to the United States. Then Charlie may be able to retaliate by ruining his reputation in the United States or seizing his goods. So this game may be part of a larger game, perhaps an ongoing interaction, that ensures Fredo's honesty. But in the one-time version we showed above, the logic of backward reasoning is clear.

We would like to use this game to make three remarks. First, different games may have identical or very similar mathematical forms (trees, or the tables used for depictions in later chapters). Thinking about them using such formalisms highlights the parallels and makes it easy to transfer your knowledge about a game in one situation to that in another. This is an important function of the "theory" of any subject: it distills the essential similarities in apparently dissimilar contexts and enables one to think about them in a unified and therefore simplified manner. Many people have an instinctive aversion to theory of any kind. But we think this is a mistaken reaction. Of course, theories have their limitations. Specific contexts and experiences can often add to or modify the prescriptions of theory in substantial ways. But to abandon theory altogether would be to abandon a valuable starting point for thought, which may be a beachhead for conquering the problem. You should make game theory your friend, and not a bugbear, in your strategic thinking.

The second remark is that Fredo should recognize that a strategic Charlie would be suspicious of his pitch and not invest at all, depriving Fredo of the opportunity to make \$250,000. Therefore Fredo has a strong incentive to make his promise credible. As an individual businessman, he has little influence over

Freedonia's weak legal system and cannot allay the investor's suspicion that way. What other methods may be at his disposal? We will examine the general issue of credibility, and devices for achieving it, in chapters 6 and 7.

The third, and perhaps most important, remark concerns comparisons of the different outcomes that could result based on the different choices the players could make. It is not always the case that more for one player means less for the other. The situation where Charlie invests and Fredo honors the contract is better for both than the one where Charlie does not invest at all. Unlike sports or contests, games don't have to have winners and losers; in the jargon of game theory, they don't have to be zero-sum. Games can have win-win or lose-lose outcomes. In fact, some combination of commonality of interest (as when Charlie and Fredo can both gain if there is a way for Fredo to commit credibly to honoring the contract) and some conflict (as when Fredo can gain at Charlie's expense by absconding after Charlie has invested) coexist in most games in business, politics, and social interactions. And that is precisely what makes the analysis of these games so interesting and challenging.

More Complex Trees

We turn to politics for an example of a slightly more complex game tree. A caricature of American politics says that Congress likes pork-barrel expenditures and presidents try to cut down the bloated budgets that Congress passes. Of course presidents have their own likes and dislikes among such expenditures and would like to cut only the ones they dislike. To do so, they would like to have the power to cut out specific items from the budget, or a line-item veto. Ronald Reagan in his State of the Union address in January 1987 said this eloquently: "Give us the same tool that 43 governors have—a line-item veto, so we can carve out the boon-doggles and pork, those items that would never survive on their own."

At first sight, it would seem that having the freedom to veto parts of a bill can only increase the president's power and never yield him any worse outcomes. Yet it is possible that the president may be better off without this tool. The point is that the existence of a line-item veto will influence the Congress's strategies in passing bills. A simple game shows how.


For this purpose, the essence of the situation in 1987 was as follows. Suppose there were two items of expenditure under consideration: urban renewal (U) and an antiballistic missile system (M). Congress liked the former and the president liked the latter. But both preferred a package of the two to the status

quo. The following table shows the ratings of the possible scenarios by the two players, in each case 4 being best and 1, worst.


Outcomes	Congress	President
Both U and M	3	3
U only	4	1
M only	1	4
Neither	2	2

The tree for the game when the president does not have a line-item veto is shown on the following page. The president will sign a bill containing the package of U and M, or one with M alone, but will veto one with U alone. Knowing this, the Congress chooses the package. Once again we show the selections at each point by thickening the chosen branches and giving them arrowheads. Note that we have to do this for all the points where the president might conceivably be called upon to choose, even though some of these are rendered moot by Congress's previous choice. The reason is that Congress's actual choice is crucially affected by its calculation of what the president would have done if Congress had counterfactually made a different choice; to show this logic we must show the president's actions in all logically conceivable situations.

Our analysis of the game yields an outcome in which both sides get their second best preference (rating 3).



Next, suppose the president has a line-item veto. The game changes to the following:



Now Congress foresees that if it passes the package, the president will selectively veto U, leaving only M. Therefore Congress's best action is now either to pass U only to see it vetoed, or pass nothing. Perhaps it may have a preference for the former, if it can score political points from a presidential veto, but perhaps the president may equally score political points by this show of budgetary discipline. Let us suppose the two offset each other, and Congress is indifferent with respect to the two choices. But either gives each party only their third-best outcome (rating 2). Even the president is left worse-off by his extra freedom of choice.²

This game illustrates an important general conceptual point. In single-person decisions, greater freedom of action can never hurt. But in games, it can hurt because its existence can influence other players' actions. Conversely, tying your own hands can help. We will explore this "advantage of commitment" in chapters 6 and 7.

We have applied the method of backward reasoning in a game tree to a very trivial game (Charlie Brown), and extended it to a slightly more complicated game (the line-item veto). The general principle remains applicable, no matter how complicated the game may be. But trees for games where each player has several choices available at any point, and where each player gets several turns to move, can quickly get too complicated to draw or use. In chess, for example, 20 branches emerge from the root—the player with the white pieces can move any of his/her eight pawns forward one square or two, or move one of his two

knights in one of two ways. For each of these, the player with the black pieces has 20 moves, so we are up to 400 distinct paths already. The number of branches emerging from later nodes in chess can be even larger. Solving chess fully using the tree method is beyond the ability of the most powerful computer that exists or might be developed during the next several decades, and other methods of partial analysis must be sought. We will discuss later in the chapter how chess experts have tackled this problem.

Between the two extremes lie many moderately complex games that are played in business, politics, and everyday life. Two approaches can be used for these. Computer programs are available to construct trees and compute solutions.³ Alternatively, many games of moderate complexity can be solved by the logic of tree analysis, without drawing the tree explicitly. We illustrate this using a game that was played in a TV show that is all about games, where each player tries to “outplay, outwit, and outlast” the others.

STRATEGIES FOR “SURVIVORS”

CBS’s *Survivor* features many interesting games of strategy. In the sixth episode of *Survivor: Thailand*, the two teams or tribes played a game that provides an excellent example of thinking forward and reasoning backward in theory and in practice.⁴ Twenty-one flags were planted in the field of play between the tribes, who took turns removing the flags. Each tribe at its turn could choose to remove 1 or 2 or 3 flags. (Thus zero—passing up one’s turn—was not permitted; nor was it within the rules to remove four or more at one turn.) The team to take the last flag, whether standing alone or as a part of a group of 2 or 3 flags, won the game.⁵ The losing tribe had to vote out one of its own members, thus weakening it in future contests. In fact the loss proved crucial in this instance, and a member of the other tribe went on to win the ultimate prize of a million dollars. Thus the ability to figure out the correct strategy for this game would prove to be of great value.

The two tribes were named Sook Jai and Chuay Gahn, and Sook Jai had the first move. They started by taking 2 flags and leaving 19. Before reading on, pause a minute and think. If you were in their place, how many would you have chosen?

Write down your choice somewhere, and read on. To understand how the game should be played, and compare the correct strategy with how the two tribes actually played, it helps to focus on two very revealing incidents. First, each tribe had a few minutes to discuss the game among its own members before the

play started. During this discussion within Chuay Gahn, one of the members, Ted Rogers, an African American software developer, pointed out, “At the end, we must leave them with four flags.” This is correct: if Sook Jai faces 4 flags, it must take 1 or 2 or 3, leaving Chuay Gahn to take the remaining 3 or 2 or 1, respectively, at its next turn and win the game. Chuay Gahn did in fact get and exploit this opportunity correctly; facing 6 flags, they took 2.

But here is the other revealing incident. At the previous turn, just as Sook Jai returned from having taken 3 flags out of the 9 facing them, the realization hit one of their members, Shii Ann, a feisty and articulate competitor who took considerable pride in her analytical skills: “If Chuay Gahn now takes two, we are sunk.” So Sook Jai’s just-completed move was wrong. What should they have done?

Shii Ann or one of her Sook Jai colleagues should have reasoned as Ted Rogers did but carried the logic of leaving the other tribe with 4 flags to its next step. How do you ensure leaving the other tribe with 4 flags at its next turn? By leaving it with 8 flags at its previous turn. When it takes 1 or 2 or 3 out of eight, you take 3 or 2 or 1 at your next turn, leaving them with 4 as planned. Therefore Sook Jai should have turned the tables on Chuay Gahn and taken just 1 flag out of the 9. Shii Ann’s analytical skill kicked into high gear one move too late! Ted Rogers perhaps had the better analytical insights. But did he?

How did Sook Jai come to face 9 flags at its previous move? Because Chuay Gahn had taken 2 from 11 at *its* previous turn. Ted Rogers should have carried his own reasoning one step further. Chuay Gahn should have taken 3, leaving Sook Jai with 8, which would be a losing position.

The same reasoning can be carried even farther back. To leave the other tribe with 8 flags, you must leave them with 12 at their previous turn; for that you must leave them with 16 at the turn before that and 20 at the turn before that. So Sook Jai should have started the game by taking just 1 flag, not 2 as it actually did. Then it could have had a sure win by leaving Chuay Gahn with 20, 16,...4 at their successive turns.*

Now think of Chuay Gahn’s very first turn. It faced 19 flags. If it had carried its own logic back far enough, it would have taken 3, leaving Sook Jai with 16 and already on the way to certain defeat. Starting from any point in the middle of the game where the opponent has played incorrectly, the team with the turn to move can seize the initiative and win. But Chuay Gahn did not play the game perfectly either.*

The table below shows the comparison between the actual and the correct moves at each point in the game. (The entry “No move” means that all moves are losing moves if the opponent plays correctly.) You can see that almost all the

choices were wrong, except Chuay Gahn's move when facing 13 flags, and that must have been accidental, because at their next turn they faced 11 and took 2 when they should have taken 3.

Tribe	No. of flags before move	No. of flags taken	Move to put team on path to sure victory
Sook Jai	21	2	1
Chuay Gahn	19	2	3
Sook Jai	17	2	1
Chuay Gahn	15	1	3
Sook Jai	14	1	2
Chuay Gahn	13	1	1
Sook Jai	12	1	No move
Chuay Gahn	11	2	3
Sook Jai	9	3	1
Chuay Gahn	6	2	2
Sook Jai	4	3	No move
Chuay Gahn	1	1	1

Before you judge the tribes harshly, you should recognize that it takes time and some experience to learn how to play even very simple games. We have played this game between pairs or teams of students in our classes and found that it takes Ivy League freshmen three or even four plays before they figure out the complete reasoning and play correctly all the way through from the first move. (By the way, what number did you choose when we asked you to initially, and what was your reasoning?) Incidentally, people seem to learn faster by watching others play than by playing themselves; perhaps the perspective of an observer is more conducive to seeing the game as a whole and reasoning about it coolly than that of a participant.

TRIP TO THE GYM NO. 1

Let us turn the flag game into hot potato: now you win by forcing the other team to take the last flag. It's your move and there are 21 flags. How many do you take?

To fix your understanding of the logic of the reasoning, we offer you the first of our Trips to the Gym—questions on which you can exercise and hone your

developing skills in strategic thinking. The answers are in the Workouts section in the end of the book.

Now that you are invigorated by this exercise, let us proceed to think about some general issues of strategy in this whole class of games.

WHAT MAKES A GAME FULLY SOLVABLE BY BACKWARD REASONING?

The 21-flags game had a special property that made it fully solvable, namely the absence of uncertainty of any kind: whether about some natural chance elements, the other players' motives and capabilities, or their actual actions. This seems a simple point to make, but it needs some elaboration and clarification.

First, at any point in the game when one tribe had the move, it knew exactly what the situation was, namely how many flags remained. In many games there are elements of pure chance, thrown up by nature or by the gods of probability. For example, in many card games, when a player makes a choice, he/she does not know for sure what cards the other players hold, although their previous actions may give some basis for drawing some inferences about that. In many subsequent chapters, our examples and analysis will involve games that have this natural element of chance.

Secondly, the tribe making its choice also knew the other tribe's objective, namely to win. And Charlie Brown should have known that Lucy enjoyed seeing him fall flat on his back. Players have such perfect knowledge of the other player's or players' objectives in many simple games and sports, but that is not necessarily the case in games people play in business, politics, and social interactions. Motives in such games are complex combinations of selfishness and altruism, concern for justice or fairness, short-run and long-run considerations, and so on. To figure out what the other players will choose at future points in the game, you need to know what their objectives are and, in the case of multiple objectives, how they will trade one off against the other. You can almost never know this for sure and must make educated guesses. You must not assume that other people will have the same preferences as you do, or as a hypothetical "rational person" does, but must genuinely think about their situation. Putting yourself in the other person's shoes is a difficult task, often made more complicated by your emotional involvement in your own aims and pursuits. We will have more to say about this kind of uncertainty later in this chapter and at various points throughout the book. Here we merely point out that the uncertainty about other players' motives is an issue for which it may be

useful to seek advice from an objective third party—a strategic consultant.

Finally, players in many games must face uncertainty about other players' choices; this is sometimes called strategic uncertainty to distinguish it from the natural aspects of chance, such as a distribution of cards or the bounce of a ball from an uneven surface. In 21-flags there was no strategic uncertainty, because each tribe saw and knew exactly what the other had done previously. But in many games, players take their actions simultaneously or in such rapid sequence that one cannot see what the other has done and react to it. A soccer goalie facing a penalty kick must decide whether to move to his/her own right or left without knowing which direction the shooter will aim for; a good shooter will conceal his/her own intentions up to the last microsecond, by which time it is too late for the goalie to react. The same is true for serves and passing shots in tennis and many other sports. Each participant in a sealed-bid auction must make his/her own choice without knowing what the other bidders are choosing. In other words, in many games the players make their moves simultaneously, and not in a preassigned sequence. The kind of thinking that is needed for choosing one's action in such games is different from, and in some respects harder than, the pure backward reasoning of sequential-move games like 21-flags; each player must be aware of the fact that others are making conscious choices and are in turn thinking about what he himself is thinking, and so on. The games we consider in the next several chapters will elucidate the reasoning and solution tools for simultaneous-move games. In this tools chapter, however, we focus solely on sequential-move games, as exemplified by 21-flags and, at a much higher level of complexity, chess.

Do People Actually Solve Games by Backward Reasoning?

Backward reasoning along a tree is the correct way to analyze and solve games where the players move sequentially. Those who fail to do so either explicitly or intuitively are harming their own objectives; they should read our book or hire a strategic consultant. But that is an advisory or normative use of the theory of backward reasoning. Does the theory have the usual explanatory or positive value that most scientific theories do? In other words, do we observe the correct outcomes from the play of actual games? Researchers in the new and exciting fields of behavioral economics and behavioral game theory have conducted experiments that yield mixed evidence.

What seems to be the most damaging criticism comes from the ultimatum game. This is the simplest possible negotiation game: there is just one take-it-or-

leave-it offer. The ultimatum game has two players, a “proposer,” say A, a “responder,” say B, and a sum of money, say 100 dollars. Player A begins the game by proposing a division of the 100 dollars between the two. Then B decides whether to agree to A’s proposal. If B agrees, the proposal is implemented; each player gets what A proposed and the game ends. If B refuses, then neither player gets anything, and the game ends.

A QUICK TRIP TO THE GYM: REVERSE ULTIMATUM GAME

In this variant of the ultimatum game, A makes an offer to B about how to divide up the 100 dollars. If B says yes, the money is divided up and the game is over. But if B says no, then A must decide whether to make another offer or not. Each subsequent offer from A must be more generous to B. The game ends when either B says yes or A stops making offers. How do you predict this game will end up?

In this case, we can suppose that A will keep on making offers until he has proposed 99 to B and 1 for himself. Thus, according to tree logic, B should get almost all of the pie. If you were B, would you hold out for 99:1? We’d advise against it.

Pause a minute and think. If you were playing this game in the A role, what division would you propose?

Now think how this game would be played by two people who are “rational” from the point of view of conventional economic theory—that is, each is concerned only with his or her self-interest and can calculate perfectly the optimal strategies to pursue that interest. The proposer (A) would think as follows. “No matter what split I propose, B is left with the choice between that and nothing. (The game is played only once, so B has no reason to develop a reputation for toughness, or to engage in any tit-for-tat response to A’s actions.) So B will accept whatever I offer. I can do best for myself by offering B as little as possible—for example, just one cent, if that is the minimum permissible under the rules of the game.” Therefore A would offer this minimum and B would accept.*

Pause and think again. If you were playing this game in the B role, would

you accept one cent?

Numerous experiments have been conducted on this game.⁶ Typically, two dozen or so subjects are brought together and are matched randomly in pairs. In each pair, the roles of proposer and responder are assigned, and the game is played once. New pairs are formed at random, and the game played again. Usually the players do not know with whom they are matched in any one play of the game. Thus the experimenter gets several observations from the same pool in the same session, but there is no possibility of forming ongoing relationships that can affect behavior. Within this general framework, many variations of conditions are attempted, to study their effects on the outcomes.

Your own introspection of how you would act as proposer and as responder has probably led you to believe that the results of actual play of this game should differ from the theoretical prediction above. And indeed they differ, often dramatically so. The amounts offered to the responder differ across proposers, but one cent or one dollar, or in fact anything below 10 percent of the total sum at stake, is very rare. The median offer (half of the proposers offer less than that and half offer more) is in the 40–50 percent range; in many experiments a 50:50 split is the single most frequent proposal. Proposals that would give the responder less than 20 percent are rejected about half the time.

IRRATIONALITY VERSUS OTHER-REGARDING RATIONALITY

Why do proposers offer substantial shares to the responders? Three reasons suggest themselves. First, the proposers may be unable to do the correct backward reasoning. Second, the proposers may have motives other than the pure selfish desire to get as much as they can; they act altruistically or care about fairness. Third, they may fear that responders would reject low offers.

The first is unlikely because the logic of backward reasoning is so simple in this game. In more complex situations, players may fail to do the necessary calculations fully or correctly, especially if they are novices to the game being played, as we saw in 21-flags. But the ultimatum game is surely simple enough, even for novices. The explanation must be the second, the third, or a combination thereof.

Early results from ultimatum experiments favored the third. In fact, Harvard's Al Roth and his coauthors found that, given the pattern of rejection thresholds that prevailed in their subject pool, the proposers were choosing their offers to achieve an optimal balance between the prospect of obtaining a greater share for themselves against the risk of rejection. This suggests a remarkable

conventional rationality on part of the proposers.

However, later work to distinguish the second and the third possibilities led to a different idea. To distinguish between altruism and strategy, experiments were done using a variant called the dictator game. Here the proposer dictates how the available total is to be split; the other player has no say in the matter at all. Proposers in the dictator game give away significantly smaller sums on average than they offer in the ultimatum game, but they give away substantially more than zero. Thus there is something to both of those explanations; proposers' behavior in the ultimatum game has both generous and strategic aspects.

Is the generosity driven by altruism or by a concern for fairness? Both explanations are different aspects of what might be called a regard for others in people's preferences. Another variation of the experiment helps tell these two possibilities apart. In the usual setup, after the pairs are formed, the roles of proposer and responder are assigned by a random mechanism like a coin toss. This may build in a notion of equality or fairness in the players' minds. To remove this, a variant assigns the roles by holding a preliminary contest, such as a test of general knowledge, and making its winner the proposer. This creates some sense of entitlement to the proposer, and indeed leads to offers that are on average about 10 percent smaller. However, the offers remain substantially above zero, indicating that proposers have an element of altruism in their thinking. Remember that they do not know the identity of the responders, so this must be a generalized sense of altruism, not concern for the well-being of a particular person.

A third variation of individual preferences is also possible: contributions may be driven by a sense of shame. Jason Dana of the University of Illinois, Daylian Cain of Yale School of Management, and Robyn Dawes of Carnegie-Mellon University performed an experiment with the following variation of the dictator game.⁷ The dictator is asked to allocate \$10. After the allocation is made, but before it is delivered to the other party, the dictator is given the following offer: You can have \$9, the other party will get nothing, and they will never know that they were part of this experiment. Most dictators accept this offer. Thus they would rather give up a dollar to ensure that the other person never knows how greedy they were. (An altruistic person would prefer keeping \$9 and giving away \$1 to keeping \$9 while the other person gets nothing.) Even when a dictator had offered \$3, he would rather take that away to keep the other person in the dark. This is much like incurring a large cost to cross the street to avoid making a small donation to a beggar.

Observe two things about these experiments. First, they follow the standard

methodology of science: hypotheses are tested by designing appropriate variations of controls in the experiment. We mention a few prominent variations of this kind here. (Many more are discussed in Colin Camerer's book cited in chapter 2, note 6.) Second, in the social sciences, multiple causes often coexist, each contributing part of the explanation for the same phenomenon. Hypotheses don't have to be either fully correct or totally wrong; accepting one need not mean rejecting all others.

Now consider the behavior of the responders. Why do they reject an offer when they know that the alternative is to get even less? The reason cannot be to establish a reputation for being a tough negotiator that may bear fruit in future plays of this game or other games of division. The same pair does not play repeatedly, and no track record of one player's past behavior is made available to future partners. Even if a reputational motive is implicitly present, it must take a deeper form: a general rule for action that the responder follows without doing any explicit thinking or calculation in each instance. It must be an instinctive action or an emotion-driven response. And that is indeed the case. In a new emerging line of experimental research called neuroeconomics, the subjects' brain activity is scanned using functional magnetic resonance imaging (fMRI) or positron emission tomography (PET) while they make various economic decisions. When ultimatum games are played under such conditions, it is found that the responders' anterior insula shows more activity as the proposers' offers become more unequal. Since the anterior insula is active for emotions, such as anger and disgust, this result helps explain why second movers reject unequal offers. Conversely, the left-side pre-frontal cortex is more active when an unequal offer is accepted, indicating that conscious control is being exercised to balance between acting on one's disgust and getting more money.⁸

Many people (especially economists) argue that while responders may reject small shares of the small sums that are typically on offer in laboratory experiments, in the real world, where stakes are often much larger, rejection must be very unlikely. To test this, ultimatum game experiments have been conducted in poorer countries where the amounts were worth several months' income for the participants. Rejection does become somewhat less likely, but offers do not become significantly less generous. The consequences of rejection become more serious for the proposers just as they do for the responders, so proposers fearing rejection are likely to behave more cautiously.

Although behavior can be explained in part by instincts, hormones, or emotions hardwired into the brain, part of it varies from one culture to another. In experiments conducted across many countries, it was found that the perception of what constitutes a reasonable offer varied by up to 10 percent

across cultures, but properties like aggressiveness or toughness varied less. Only one group differed substantially from the rest: among the Machiguenga of the Peruvian Amazon, the offers were much smaller (average 26 percent) and only one offer was rejected. Anthropologists explain that the Machiguenga live in small family units, are socially disconnected, and have no norms of sharing. Conversely, in two cultures the offers exceeded 50 percent; these have the custom of lavish giving when one has a stroke of good luck, which places an obligation on the recipients to return the favor even more generously in the future. This norm or habit seems to carry over to the experiment even though the players do not know whom they are giving to or receiving from.⁹

Evolution of Altruism and Fairness

What should we learn from the findings of these experiments on the ultimatum game, and others like them? Many of the outcomes do differ significantly from what we would expect based on the theory of backward reasoning with the assumption that each player cares only about his or her own reward. Which of the two—correct backward calculation or selfishness—is the wrong assumption, or is it a combination? And what are the implications?

Consider backward reasoning first. We saw the players on *Survivor* fail to do this correctly or fully in 21-flags. But they were playing the game for the first time, and even then, their discussion revealed glimpses of the correct reasoning. Our classroom experience shows that students learn the full strategy after playing the game, or watching it played, just three or four times. Many experiments inevitably and almost deliberately work with novice subjects, whose actions in the game are often steps in the process of learning the game. In the real world of business, politics, and professional sports, where people are experienced at playing the games they are involved in, we should expect that the players have accumulated much more learning and that they play generally good strategies either by calculation or by trained instinct. For somewhat more complex games, strategically aware players can use computers or consultants to do the calculations; this practice is still somewhat rare but is sure to spread. Therefore, we believe that backward reasoning should remain our starting point for analysis of such games and for predicting their outcomes. This first pass at the analysis can then be modified as necessary in a particular context, to recognize that beginners may make mistakes and that some games may become too complex to be solved unaided.

We believe that the more important lesson from the experimental research is

that people bring many considerations and preferences into their choices besides their own rewards. This takes us beyond the scope of conventional economic theory. Game theorists should include in their analysis of games the players' concerns for fairness or altruism. "Behavioral game theory *extends* rationality rather than abandoning it."¹⁰

This is all to the good; a better understanding of people's motives enriches our understanding of economic decision making and strategic interactions alike. And that is already happening; frontier research in game theory increasingly includes in the players' objectives their concerns for equity, altruism, and similar concerns (and even a "second-round" concern to reward or punish others whose behavior reflects or violates these precepts).¹¹

But we should not stop there; we should go one step further and think about why concerns for altruism and fairness, and anger or disgust when someone else violates these precepts, have such a strong hold on people. This takes us into the realm of speculation, but one plausible explanation can be found in evolutionary psychology. Groups that instill norms of fairness and altruism into their members will have less internal conflict than groups consisting of purely selfish individuals. Therefore they will be more successful in taking collective action, such as provision of goods that benefit the whole group and conservation of common resources, and they will spend less effort and resources in internal conflict. As a result, they will do better, both in absolute terms and in competition with groups that do not have similar norms. In other words, some measure of fairness and altruism may have evolutionary survival value.

Some biological evidence for rejecting unfair offers comes from an experiment run by Terry Burnham.¹² In his version of the ultimatum game, the amount at stake was \$40 and the subjects were male Harvard graduate students. The divider was given only two choices: offer \$25 and keep \$15 or offer \$5 and keep \$35. Among those offered only \$5, twenty students accepted and six rejected, giving themselves and the divider both zero. Now for the punch line. It turns out that the six who rejected the offer had testosterone levels 50 percent higher than those who accepted the offer. To the extent that testosterone is connected with status and aggression, this could provide a genetic link that might explain an evolutionary advantage of what evolutionary biologist Robert Trivers has called "moralistic aggression."

In addition to a potential genetic link, societies have nongenetic ways of passing on norms, namely the processes of education and socialization of infants and children in families and schools. We see parents and teachers telling impressionable children the importance of caring for others, sharing, and being

nice; some of this undoubtedly remains imprinted in their minds and influences their behavior throughout their lives.

Finally, we should point out that fairness and altruism have their limit. Long-run progress and success of a society need innovation and change. These in turn require individualism and a willingness to defy social norms and conventional wisdom; selfishness often accompanies these characteristics. We need the right balance between self-regarding and other-regarding behaviors.

VERY COMPLEX TREES

When you have acquired a little experience with backward reasoning, you will find that many strategic situations in everyday life or work lend themselves to “tree logic” without the need to draw and analyze trees explicitly. Many other games at an intermediate level of complexity can be solved using computer software packages that are increasingly available for this purpose. But for complex games such as chess, a complete solution by backward reasoning is simply not feasible.

In principle, chess is the ideal game of sequential moves amenable to solution by backward reasoning.¹³ The players alternate moves; all previous moves are observable and irrevocable; there is no uncertainty about the position or the players’ motives. The rule that the game is a draw if the same position is repeated ensures that the game ends within a finite total number of moves. We can start with the terminal nodes (or endpoints) and work backward. However, practice and principle are two different things. It has been estimated that the total number of nodes in chess is about 10^{120} , that is, 1 with 120 zeroes after it. A supercomputer 1,000 times as fast as the typical PC would need 10^{103} years to examine them all. Waiting for that is futile; foreseeable progress in computers is not likely to improve matters significantly. In the meantime, what have chess players and programmers of chess-playing computers done?

Chess experts have been successful at characterizing optimal strategies near the end of the game. Once the chessboard has only a small number of pieces on it, experts are able to look ahead to the end of the game and determine by backward reasoning whether one side has a guaranteed win or whether the other side can obtain a draw. But the middle of the game, when several pieces remain on the board, is far harder. Looking ahead five pairs of moves, which is about as much as can be done by experts in a reasonable amount of time, is not going to simplify the situation to a point where the endgame can be solved completely from there on.

The pragmatic solution is a combination of forward-looking analysis and value judgment. The former is the science of game theory—looking ahead and reasoning backward. The latter is the art of the practitioner—being able to judge the value of a position from the number and interconnections of the pieces without finding an explicit solution of the game from that point onward. Chess players often speak of this as “knowledge,” but you can call it experience or instinct or art. The best chess players are usually distinguished by the depth and subtlety of their knowledge.

Knowledge can be distilled from the observation of many games and many players and then codified into rules. This has been done most extensively with regard to openings, that is, the first ten or even fifteen moves of a game. There are hundreds and hundreds of books that analyze different openings and discuss their relative merits and drawbacks.

How do computers fit into this picture? At one time, the project of programming computers to play chess was seen as an integral part of the emerging science of artificial intelligence; the aim was to design computers that would think as humans do. This did not succeed for many years. Then the attention shifted to using computers to do what they do best, crunch numbers. Computers can look ahead to more moves and do this more quickly than humans can.* Using pure number crunching, by the late 1990s dedicated chess computers like Fritz and Deep Blue could compete with the top human players. More recently, computers have been programmed with some knowledge of midgame positions, imparted by some of the best human players.

Human players have ratings determined by their performances; the best-ranked computers are already achieving ratings comparable to the 2800 enjoyed by the world’s strongest human player, Garry Kasparov. In November 2003, Kasparov played a four-game match against the latest version of the Fritz computer, X3D. The result was one victory each and two draws. In July 2005, the Hydra chess computer demolished Michael Adams, ranked number 13 in the world, winning five games and drawing one in a six-game match. It may not be long before the rival computers rank at the top and play each other for world championships.

What should you take away from this account of chess? It shows the method for thinking about any highly complex games you may face. You should combine the rule of look ahead and reason back with your experience, which guides you in evaluating the intermediate positions reached at the end of your span of forward calculation. Success will come from such synthesis of the science of game theory and the art of playing a specific game, not from either alone.

BEING OF TWO MINDS

Chess strategy illustrates another important practical feature of looking forward and reasoning backward: you have to play the game from the perspective of both players. While it is hard to calculate your best move in a complicated tree, it is even harder to predict what the other side will do.

If you really could analyze all possible moves and countermoves, and the other player could as well, then the two of you would agree up front as to how the entire game would play out. But once the analysis is limited to looking down only some branches of the tree, the other player may see something you didn't or miss something you've seen. Either way, the other side may then make a move you didn't anticipate.

To really look forward and reason backward, you have to predict what the other players will actually do, not what you would have done in their shoes. The problem is that when you try to put yourself in the other players' shoes, it is hard if not impossible to leave your own shoes behind. You know too much about what you are planning to do in your next move and it is hard to erase that knowledge when you are looking at the game from the other player's perspective. Indeed, that explains why people don't play chess (or poker) against themselves. You certainly can't bluff against yourself or make a surprise attack.

There is no perfect solution to this problem. When you try to put yourself in the other players' shoes, you have to know what they know and not know what they don't know. Your objectives have to be their objectives, not what you wish they had as an objective. In practice, firms trying to simulate the moves and countermoves of a potential business scenario will hire outsiders to play the role of the other players. That way, they can ensure that their game partners don't know too much. Often the biggest learning comes from seeing the moves that were not anticipated and then understanding what led to that outcome, so that it can be either avoided or promoted.

To end this chapter, we return to Charlie Brown's problem of whether or not to kick the football. This question became a real issue for football coach Tom Osborne in the final minutes of his championship game. We think he too got it wrong. Backward reasoning will reveal the mistake.

CASE STUDY: THE TALE OF TOM OSBORNE AND THE '84 ORANGE BOWL

In the 1984 Orange Bowl the undefeated Nebraska Cornhuskers and the

once-beaten Miami Hurricanes faced off. Because Nebraska came into the Bowl with the better record, it needed only a tie in order to finish the season with the number-one ranking.

Coming into the fourth quarter, Nebraska was behind 31–17. Then the Cornhuskers began a comeback. They scored a touchdown to make the score 31–23. Nebraska coach Tom Osborne had an important strategic decision to make.

In college football, a team that scores a touchdown then runs one play from a hash mark 2 1/2 yards from the goal line. The team has a choice between trying to run (or pass) the ball into the end zone, which scores two additional points, or trying the less risky strategy of kicking the ball through the goalposts, which scores one extra point.

Coach Osborne chose to play it safe, and Nebraska successfully kicked for the one extra point. Now the score was 31–24. The Cornhuskers continued their comeback. In the waning minutes of the game they scored a final touchdown, bringing the score to 31–30. A one-point conversion would have tied the game and landed them the title. But that would have been an unsatisfying victory. To win the championship with style, Osborne recognized that he had to go for the win.

The Cornhuskers went for the win with a two-point conversion attempt. Irving Fryar got the ball but failed to score. Miami and Nebraska ended the year with equal records. Since Miami beat Nebraska, it was Miami that was awarded the top place in the standings.

Put yourself in the cleats of Coach Osborne. Could you have done better?

Case Discussion

Many Monday morning quarterbacks fault Osborne for going for the win rather than the tie. But that is not our bone of contention. Given that Osborne was willing to take the additional risk for the win, he did it the wrong way. He would have done better to first try the two-point conversion. If it succeeded, then go for the one-point; if it failed, attempt a second two-pointer.

Let us look at this more carefully. When down by 14 points, he knew that he needed two touchdowns plus three extra points. He chose to go for the one-point and then the two. If both attempts succeeded, the order in which they were made becomes irrelevant. If the one-point conversion was missed but the two-point was successful, here too the order is irrelevant and the game ends up tied with Nebraska getting the championship. The only difference occurs if Nebraska

misses the two-point conversion. Under Osborne's plan, that results in the loss of the game and the championship. If, instead, they had tried the two-point conversion first, then if it failed they would not necessarily have lost the game. They would have been behind 31–23. When they scored their next touchdown this would have brought them to 31–29. A successful two-point attempt would tie the game and win the number-one ranking!*

We have heard the counterargument that if Osborne went for the two-pointer first and missed, his team would have been playing for the tie. This would have provided less inspiration and perhaps they might not have scored the second touchdown. Moreover, by waiting until the end and going for the desperation win-lose two-pointer his team would rise to the occasion knowing everything was on the line. This argument is wrong for several reasons. Remember that if Nebraska waits until the second touchdown and then misses the two-point attempt, they lose. If they miss the two-point attempt on their first try, there is still a chance for a tie. Even though the chance may be diminished, something is better than nothing. The momentum argument is also flawed. While Nebraska's offense may rise to the occasion in a single play for the championship, we expect the Hurricanes' defense to rise as well. The play is important for both sides. To the extent that there is a momentum effect, if Osborne makes the two-point conversion on the first touchdown, this should increase the chance of scoring another touchdown. It also allows him to tie the game with two field goals.

One of the general morals of this story is that if you have to take some risks, it is often better to do so as quickly as possible. This is obvious to those who play tennis: everyone knows to take more risk on the first serve and hit the second serve more cautiously. That way, if you fail on your first attempt, the game won't be over. You may still have time to take some other options that can bring you back to or even ahead of where you were. The wisdom of taking risks early applies to most aspects of life, whether it be career choices, investments, or dating.

For more practice using the principle of look forward, reason backward, have a look at the following case studies in chapter 14: "Here's Mud in Your Eye"; "Red I Win, Black You Lose"; "The Shark Repellent That Backfired"; "Tough Guy, Tender Offer"; "The Three-Way Duel"; and "Winning without Knowing How."