Thinking Strategically

The Competitive Edge in Business, Politics, and Everyday Life

Avinash K. Dixit and Barry J. Nalebuff

W. W. Norton & Company

New York London

Printed in the United States of America.
Typeset by Barry J. Nalebuff using Textures.

Illustration credit: Chapter 2: Cartoon by Charles Schulz.
Reprinted with permission of United Feature Syndicate, Inc.

Reissued in Norton Paperback 1993

Library of Congress Cataloging in Publication Data

Dixit, Avinash K.
Thinking strategically: the competitive edge in business, politics, and everyday life / by Avinash K. Dixit and Barry J. Nalebuff.
p. cm.
Includes Index
1. Game Theory 2. Management 3. Competition
I. Nalebuff, Barry, 1958 – . II. Title.
HD30.28.D59 1991
658.4'012–dc20 90-33760

ISBN-13: 978-0-393-31035-1
ISBN-10: 0-393-31035-3

W. W. Norton & Company, Inc.
500 Fifth Avenue, New York, N.Y. 10110
www.wwnorton.com

W. W. Norton & Company Ltd.
Castle House, 75/76 Wells Street, London W1T 3QT

3 4 5 6 7 8 9 0

Introduction

What Is Strategic Behavior?

How should people behave in society?

Our answer does not deal with ethics or etiquette. Nor do we aim to compete with philosophers, preachers, or even Emily Post. Our theme, although less lofty, affects the lives of all of us just as much as do morality and manners. This book is about strategic behavior. All of us are strategists, whether we like it or not. It is better to be a good strategist than a bad one, and this book aims to help you improve your skills at discovering and using effective strategies.

Work, even social life, is a constant stream of decisions. What career to follow, how to manage a business, whom to marry, how to bring up children, whether to run for president, are just some examples of such fateful choices. The common element in these situations is that you do not act in a vacuum. Instead, you are surrounded by active decision-makers whose choices interact with yours. This interaction has an important effect on your thinking and actions.

To illustrate the point, think of the difference between the decisions of a lumberjack and those of a general. When the lumberjack decides how to chop wood, he does not expect the wood to fight back; his environment is neutral. But when the general tries to cut down the enemy's army, he must anticipate and overcome resistance to his plans. Like the general, you must recognize that your business rivals, prospective spouse, and even your child are intelligent and purposive people. Their aims often conflict with yours, but they include some poten-

tial allies. Your own choice must allow for the conflict, and utilize the cooperation. Such interactive decisions are called strategic, and the plan of action appropriate to them is called a strategy. This book aims to help you think strategically, and then translate these thoughts into action.

The branch of social science that studies strategic decision-making is called *game theory*. The games in this theory range from chess to child-rearing, from tennis to takeovers, and from advertising to arms control. As the Hungarian humorist George Mikes expressed it, "Many continentals think life is a game; the English think cricket is a game." We think both are right.

Playing these games requires many different kinds of skills. Basic skills, such as shooting ability in basketball, knowledge of precedents in law, or a blank face in poker, are one kind; strategic thinking is another. Strategic thinking starts with your basic skills, and considers how best to use them. Knowing the law, you must decide the strategy for defending your client. Knowing *how well* your football team can pass or run, and how well the other team can defend against each choice, your decision as the coach is *whether* to pass or to run. Sometimes, as in the case of superpowers contemplating an adventure that risks nuclear war, strategic thinking also means knowing when not to play.

Our aim is to improve your strategy I.Q. But we have not tried to provide a book of recipes for strategies. We develop the ideas and principles of strategic thinking; to apply them to a specific situation you face and to find the right choice there, you will have to do some more work. This is because the specifics of each situation are likely to differ in some significant aspects, and any general prescriptions for action we might give could be misleading. In each situation, you will have to pull together principles of good strategy we have discussed, and also other principles from other considerations. You must combine them and, where they conflict with each other, evaluate the relative strengths of the different argu-

ments. We do not promise to solve every question you might have. The science of game theory is far from being complete, and in some ways strategic thinking remains an art.

We do provide guidance for translating the ideas into action. Chapter 1 offers several examples showing how strategic issues arise in a variety of decisions. We point out some effective strategies, some less effective ones, and even some downright bad ones. The subsequent chapters proceed to build these examples into a system or a framework of thought. In the later chapters, we take up several broad classes of strategic situations — brinkmanship, voting, incentives, and bargaining — where you can see the principles in action.

The examples range from the familiar, trivial, or amusing — usually drawn from literature, sports, or movies — to the frightening — nuclear confrontation. The former are merely a nice and palatable vehicle for the game-theoretic ideas. As to the latter, at one point many readers would have thought the subject of nuclear war too horrible to permit rational analysis. But as the cold war winds down and the world is generally perceived to be a safer place, we hope that the game-theoretic aspects of the arms race and the Cuban missile crisis can be examined for their strategic logic in some detachment from their emotional content.

The chapters are full of examples, but these serve primarily to develop or illustrate the particular principle being discussed, and many other details of reality that pertain to the example are set aside. At the end of each chapter, we present a "case study," similar to one you might come across in a business-school class. Each case sets out a particular set of circumstances and invites you to apply the principles discussed in that chapter to find the right strategy for that situation. Some cases are open-ended; but that is also a feature of life. At times there is no clearly correct solution, only imperfect ways to cope with the problem. A serious effort to think each case through before reading our discussion is a bet-

ter way to understand the ideas than any amount of reading of the text alone. For more practice, the final chapter is a collection of twenty three more cases, in roughly increasing order of difficulty.

By the end of the book, we hope that you will emerge a more effective manager, negotiator, athlete, politician, or parent. We warn you that some of the strategies that are good for achieving these goals may not earn you the love of your defeated rivals. If you want to be fair, tell them about our book.

1

Ten Tales of Strategy

We begin with ten tales of strategy from different aspects of life and offer preliminary thoughts on how best to play. Many of you will have faced similar problems in everyday life, and will have reached the correct solution after some thought or trial and error. For others, some of the answers may be surprising, but surprise is not the primary purpose of the examples. Our aim is to show that such situations are pervasive, that they amount to a coherent set of questions, and that methodical thinking about them is likely to be fruitful. In later chapters, we develop these systems of thought into prescriptions for effective strategy. Think of these tales as a taste of dessert before the main course. They are designed to whet your appetite, not fill you up.

1. The Hot Hand

Do athletes ever have a "hot hand"? Sometimes it seems that Larry Bird cannot miss a basket, or Wayne Gretzky or Diego Maradona a shot on goal. Sports announcers see these long streaks of consecutive successes and proclaim that the athlete has a "hot hand." Yet according to psychology professors Thomas Gilovich, Robert Vallone, and Amos Tversky, this is a misperception of reality.[1] They point out that if you flip a coin long enough, you will find some very long series of consecutive

heads. The psychologists suspect that sports commentators, short on insightful things to say, are just finding patterns in what amounts to a long series of coin tosses over a long playing season. They propose a more rigorous test. In basketball, they look at all the instances of a player's baskets, and observe the percentage of times that player's next shot is also a basket. A similar calculation is made for the shots immediately following misses. If a basket is more likely to follow a basket than to follow a miss, then there really is something to the theory of the hot hand.

They conducted this test on the Philadelphia 76ers basketball team. The results contradicted the "hot hand" view. When a player made his last shot, he was less likely to make his next; when he missed his previous attempt, he was more likely to make his next. This was true even for Andrew Toney, a player with the reputation for being a streak shooter. Does this mean we should be talking of the "stroboscopic hand," like the strobe light that alternates between on and off?

Game theory suggests a different interpretation. While the statistical evidence denies the presence of streak shooting, it does not refute the possibility that a "hot" player might warm up the game in some other way. The difference between streak shooting and a hot hand arises because of the interaction between the offensive and the defensive strategies. Suppose Andrew Toney does have a truly hot hand. Surely the other side would start to crowd him. This could easily lower his shooting percentage.

That is not all. When the defense focuses on Toney, one of his teammates is left unguarded and is more likely to shoot successfully. In other words, Toney's hot hand leads to an improvement in the 76ers' *team* performance, although there may be a deterioration in Toney's *individual* performance. Thus we might test for hot hands by looking for streaks in team success.

Similar phenomena are observed in many other team sports. A brilliant running-back on a football team improves its pass-

ing game and a great pass-receiver helps the running game, as the opposition is forced to allocate more of its defensive resources to guard the stars. In the 1986 soccer World Cup final, the Argentine star Diego Maradona did not score a goal, but his passes through a ring of West German defenders led to two Argentine goals. The value of a star cannot be assessed by looking only at his scoring performance; his contribution to his teammates' performance is crucial, and assist statistics help measure this contribution. In ice hockey, assists and goals are given equal weight for ranking individual performance.

A player may even assist himself when one hot hand warms up the other. The Boston Celtics star, Larry Bird, prefers shooting with his right hand (though his left hand is still better than most). The defense knows that Bird is right-handed, so they concentrate on defending against right-handed shots. But they do not do so exclusively, since Bird's left-handed shots are too effective to be left unguarded.

What happens when Bird spends his off season working to improve his left-handed shooting? The defense responds by spending more time covering his left-handed shots. The result is that this frees his right hand more often. A better left-handed shot results in a more effective right-handed shot. In this case not only does the left hand know what the right hand is doing, it's helping it out.

Going one step further, in Chapter 7 we show that when the left hand is stronger it may even be used *less* often. Many of you will have experienced this seemingly strange phenomenon when playing tennis. If your backhand is much weaker than your forehand, your opponents will learn to play to your backhand. Eventually, as a result of all this backhand practice, your backhand will improve. As your two strokes become more equal, opponents can no longer exploit your weak backhand. They will play more evenly between forehands and backhands. You get to use your better forehand more often; this could be the real advantage of improving your backhand.

2. To Lead or Not to Lead

After the first four races in the 1983 America's Cup finals, Dennis Conner's *Liberty* led 3–1 in a best-of-seven series. On the morning of the fifth race, "cases of champagne had been delivered to *Liberty*'s dock. And on their spectator yacht, the wives of the crew were wearing red-white-and-blue tops and shorts, in anticipation of having their picture taken after their husbands had prolonged the United States' winning streak to 132 years."[2] It was not to be.

At the start, *Liberty* got off to a 37-second lead when *Australia II* jumped the gun and had to recross the starting line. The Australian skipper, John Bertrand, tried to catch up by sailing way over to the left of the course in the hopes of catching a wind shift. Dennis Conner chose to keep *Liberty* on the right-hand side of the course. Bertrand's gamble paid off. The wind shifted five degrees in *Australia II*'s favor and she won the race by one minute and forty-seven seconds. Conner was criticized for his strategic failure to follow *Australia II*'s path. Two races later, *Australia II* won the series.

Sailboat racing offers the chance to observe an interesting reversal of a "follow the leader" strategy. The leading sailboat usually copies the strategy of the trailing boat. When the follower tacks, so does the leader. The leader imitates the follower even when the follower is clearly pursuing a poor strategy. Why? Because in sailboat racing (unlike ballroom dancing) close doesn't count: only winning matters. If you have the lead, the surest way to stay ahead is to play monkey see, monkey do.*

Stock-market analysts and economic forecasters are not immune to this copycat strategy. The leading forecasters have an incentive to follow the pack and produce predictions similar to

* This strategy no longer applies once there are more than two competitors. Even with three boats, if one boat tacks right and the other tacks left, the leader has to choose which (if either) to follow.

everyone else's. This way people are unlikely to change their perception of these forecasters' abilities. On the other hand, newcomers take the risky strategies: they tend to predict boom or doom. Usually they are wrong and are never heard of again, but now and again they are proven correct and move to the ranks of the famous.

Industrial and technological competitions offer further evidence. In the personal-computer market, IBM is less known for its innovation than for its ability to bring standardized technology to the mass market. More new ideas have come from Apple, Sun, and other start-up companies. Risky innovations are their best and perhaps only chance of gaining market share. This is true not just of high-technology goods. Proctor and Gamble, the IBM of diapers, followed Kimberly Clark's innovation of resealable diaper tape, and recaptured its commanding market position.

There are two ways to move second. You can imitate as soon as the other has revealed his approach (as in sailboat racing) or wait longer until the success or failure of the approach is known (as in computers). The longer wait is more advantageous in business because, unlike sports, the competition is usually not winner-take-all. As a result, market leaders will not follow the upstarts unless they also believe in the merits of their course.

3. Go Directly to Jail

The conductor of an orchestra in the Soviet Union (during the Stalin era) was traveling by train to his next engagement and was looking over the score of the music he was to conduct that night. Two KGB officers saw what he was reading and, thinking that the musical notation was some secret code, arrested him as a spy. He protested that it was only Tchaikovsky's Violin Concerto, but to no avail. On the second day of his im-

prisonment, the interrogator walked in smugly and said, "You had better tell us all. We have caught your friend Tchaikovsky, and he is already talking."

So begins one telling of the prisoners' dilemma, perhaps the best-known strategic game. Let us develop the story to its logical conclusion. Suppose the KGB has actually arrested someone whose only offense is that he is called Tchaikovsky, and are separately subjecting him to the same kind of interrogation. If the two innocents withstand this treatment, each will be sentenced to 3 years' imprisonment.* If the conductor makes a false confession that implicates the unknown "collaborator," while Tchaikovsky holds out, then the conductor will get away with 1 year (and the KGB's gratitude), while Tchaikovsky gets the harsh sentence of 25 years for his recalcitrance. Of course, the tables will be turned if the conductor stands firm while Tchaikovsky gives in and implicates him. If both confess, then both will receive the standard sentence of 10 years.†

Now consider the conductor's thinking. He knows that Tchaikovsky is either confessing or holding out. If Tchaikovsky confesses, the conductor gets 25 years by holding out and 10 years by confessing, so it is better for him to confess. If Tchaikovsky holds out, the conductor gets 3 years if he holds out, and only 1 if he confesses; again it is better for him to confess. Thus confession is clearly the conductor's best action.

In a separate cell in Dzerzhinsky Square, Tchaikovsky is doing a similar mental calculation and reaching the same conclusion. The result, of course, is that both of them confess. Later, when they meet in the Gulag Archipelago, they com-

* There is the story of the newcomer to the Gulag who was asked by the residents, "How long is your sentence?" The answer was "Ten years." "What did you do?" "Nothing." "No, there must be some mistake. The sentence for that is only three years."

† This actually meant 3,653 days: "The three extra days were for leap years." (A. Solzhenitsyn, *One Day in the Life of Ivan Denisovitch*, 1962.)

pare stories and realize that they have been had. If they both had stood firm, they both would have gotten away with much shorter sentences.

If only they had had an opportunity to meet and talk things over before they were interrogated, they could have agreed that neither would give in. But they are quick to realize that in all probability such an agreement would not have done much good. Once they were separated and the interrogations began, each person's private incentive to get a better deal by double-crossing the other would have been quite powerful. Once again they would have met in the Gulag, there perhaps to settle the score of the betrayals (not of the concerto). Can the two achieve enough mutual credibility to reach their jointly preferred solution?

Many people, firms, and even nations have been gored on the horns of the prisoners' dilemma. Look at the life-or-death issue of nuclear arms control. Each superpower liked best the outcome in which the other disarmed, while it kept its own arsenal "just in case." Disarming yourself while the other remains armed was the worst prospect. Therefore no matter what the other side did, each preferred to stay armed. However, they could join in agreeing that the outcome in which both disarm is better than the one in which both are armed. The problem is the interdependence of decisions: the *jointly* preferred outcome arises when each chooses its *individually* worse strategy. Could the jointly preferred outcome be achieved given each side's clear incentive to break the agreement and to arm itself secretly? In this case it needed a fundamental change in Soviet thinking to get the world started on the road to nuclear disarmament.

For one's comfort, safety, or even life itself, one needs to know the ways to get out of the prisoners' dilemma. In Chapter 4 we look at some such avenues, and see when and how well they are likely to work.

The story of the prisoners' dilemma also carries a useful

general point: most economic, political, or social games are different from games such as football or poker. Football and poker are *zero-sum games*: one person's gain is another person's loss. But in the prisoners' dilemma, there are possibilities for mutual advantage as well as conflict of interest; both prisoners prefer the no-confession result to its opposite. Similarly, in employer-union bargaining, there is an opposition of interests in that one side prefers low wages and the other high ones, but there is agreement that a breakdown of negotiations leading to a strike would be more damaging for both sides. In fact such situations are the rule rather than the exception. Any useful analysis of games should be able to handle a mixture of conflict and concurrence of interests. We usually refer to the players in a game as "opponents," but you should remember that on occasion, strategy makes strange bedfellows.

4. HERE I STAND

When the Catholic Church demanded that Martin Luther repudiate his attack on the authority of popes and councils, he refused to recant: "I will not recant anything, for to go against conscience is neither right nor safe." Nor would he compromise: "Here I stand, I cannot do otherwise."[3] Luther's intransigence was based on the divinity of his positions. When defining what was right, there was no room for compromise. His firmness had profound long-term consequences; his attacks led to the Protestant Reformation and substantially altered the medieval Catholic Church.*

Similarly, Charles de Gaulle used the power of intransigence to become a powerful player in the arena of interna-

* Luther's reputation extends beyond the Church and behind the Iron Curtain. The Wartburg, East Germany's domestically produced car, is jokingly referred to as "The Luther": apparently it can be equally immobile.

tional relations. As his biographer Don Cook expressed it, "[De Gaulle] could create power for himself with nothing but his own rectitude, intelligence, personality and sense of destiny."[4] But above all, his was "the power of intransigence." During the Second World War, as the self-proclaimed leader in exile of a defeated and occupied nation, he held his own in negotiations with Roosevelt and Churchill. In the 1960s, his presidential "Non!" swung several decisions France's way in the European Economic Community.

In what way did his intransigence give him power in bargaining? When de Gaulle took a truly irrevocable position, the other parties in the negotiation were left with just two options — to take it or to leave it. For example, he single-handedly kept England out of the European Economic Community, once in 1963 and again in 1968; the other countries were forced either to accept de Gaulle's veto or to break up the EEC. De Gaulle judged his position carefully to ensure that it would be accepted. But that often left the larger (and unfair) division of the spoils to France. De Gaulle's intransigence denied the other party an opportunity to come back with a counteroffer that was acceptable.

In practice, this is easier said than done, for two kinds of reasons. The first kind stems from the fact that bargaining usually involves considerations beside the pie on today's table. The perception that you have been excessively greedy may make others less willing to negotiate with you in the future. Or, next time they may be more firm bargainers as they try to recapture some of their perceived losses. On a personal level, an unfair win may spoil business relations, or even personal relations. Indeed, biographer David Schoenbrun faulted de Gaulle's chauvinism: "In human relations, those who do not love are rarely loved: those who will not be friends end up by having none. De Gaulle's rejection of friendship thus hurt France."[5] A compromise in the short term may prove a better strategy over the long haul.

The second kind of problem lies in achieving the necessary degree of intransigence. Luther and de Gaulle achieved this through their personalities. But this entails a cost. An inflexible personality is not something you can just turn on and off. Although being inflexible can sometimes wear down an opponent and force him to make concessions, it can equally well allow small losses to grow into major disasters.

Ferdinand de Lesseps was a mildly competent engineer with extraordinary vision and determination. He is famous for building the Suez Canal in what seemed almost impossible conditions. He did not recognize the impossible and thereby accomplished it. Later, he tried using the same technique to build the Panama Canal. It ended in disaster.* Whereas the sands of the Nile yielded to his will, tropical malaria did not. The problem for de Lesseps was that his inflexible personality could not admit defeat even when the battle was lost.

How can one achieve selective inflexibility? Although there is no ideal solution, there are various means by which commitment can be achieved and sustained; this is the topic for Chapter 6.

5. Belling the Cat

In the children's story about belling the cat, the mice decide that life would be much safer if the cat were stuck with a bell around its neck. The problem is, who will risk his life to bell the cat?

This is a problem for both mice and men. How can rel-

* The Suez Canal is a sea-level passage. The digging was relatively easy since the land was already low-lying and desert. Panama involved much higher elevations, lakes along the way, and dense jungle. Lesseps' attempt to dig down to sea level failed. Much later, the U.S. Army Corps of Engineers succeeded using a very different method — a sequence of locks, using the lakes along the way.

atively small armies of occupying powers or tyrants control very large populations for long periods? Why is a planeload of people powerless before a single hijacker with a gun? In both cases, a simultaneous move by the masses stands a very good chance of success. But the communication and coordination required for such action is difficult, and the oppressors, knowing the power of the masses, take special steps to keep it difficult. When the people must act individually and hope that the momentum will build up, the question arises, "Who is going to be the first?" Such a leader will pay a very high cost — possibly his life. His reward may be posthumous glory or gratitude. There are people who are moved by considerations of duty or honor, but most find the costs exceed the benefits.

Khrushchev first denounced Stalin's purges at the Soviet Communist Party's 20th Congress. After his dramatic speech, someone in the audience shouted out, asking what Khrushchev had been doing at the time. Khrushchev responded by asking the questioner to please stand up and identify himself. The audience remained silent. Khrushchev replied: "That is what I did, too."

In a sense, we have seen these examples before. They are just a prisoners' dilemma with more than two people; one might call this the hostages' dilemma. Here we want to use this dilemma to make a different point — namely, the frequent superiority of punishment over reward. The dictator might keep the populace peaceful by providing it material and even spiritual comforts, but this can be a very costly proposition. Oppression and terror relying on the Hostages' Dilemma can be a much cheaper alternative.

There are many examples of this principle. In a large taxi fleet, cars are often assigned to drivers by a dispatcher. The fleet has some good cars and some clunkers. The dispatcher can use his assignment power to extract a small bribe from each of the drivers. Any driver who refuses to pay is sure to get a clunker, while those who cooperate are given the luck

of the draw from the remainder.* The dispatcher gets rich, and the drivers as a group end up with the same set of cabs that they would have if no one used bribery. If the drivers acted in collusion, they probably could stop this practice. The problem lies in getting the movement organized. The point is not so much that the dispatcher can reward those who bribe him, but that he can punish severely those who don't.

A similar story can be told about evicting tenants from rent-controlled apartments. If someone buys such a building in New York, he has the right to evict one tenant so as to be able to live in his own building. But this translates into a power to clear the whole. A new landlord can try the following argument with the tenant in Apartment 1A: "I have the right to live in my building. Therefore, I plan to evict you and move into your apartment. However, if you cooperate and leave voluntarily, then I will reward you with $5,000." This is a token amount in relation to the value of the rent-controlled apartment (although it still buys a few subway tokens in New York). Faced with the choice of eviction with $5,000 or eviction without $5,000, the tenant takes the money and runs. The landlord then offers the same deal to the tenant in 1B, and so on.

The United Auto Workers have a similar advantage when they negotiate with the auto manufacturers sequentially. A strike against Ford alone puts it at particular disadvantage when General Motors and Chrysler continue to operate; therefore Ford is more likely to settle quickly on terms favorable to the Union. Such a strike is also less costly to the Union as only one third of their members are out. After winning against Ford, the Union takes on GM and then Chrysler, using each previous success as precedent and fuel for their fire. In contrast, Japanese union incentives work the other way,

* Even if everyone pays, some drivers will end up with a clunker. But if the clunkers are randomly assigned, no driver faces a great chance of the bad draw. In contrast, the first driver who refuses to pay can expect to drive the clunker quite regularly.

too high, and insisted that workers in each of these industries bear the risks of foreign trade just as they would have to bear any other economic risk. Decisions made case by case can lead to undesirable results overall. In fact, a sequence of majority votes can lead to an outcome that everyone regards as worse than the status quo.

The income tax reform of 1985–86 almost collapsed because the Senate initially took a case-by-case approach. In the first round of the Finance Committee's markup sessions, the amended Treasury proposal became so weighted down with special interest provisions that it sank to a merciful death. The senators realized that they were "powerless" to prevent any one organized lobby from getting special treatment. Yet the combination of these lobbyists could destroy the bill, and this would be worse than producing no legislation at all. So Senator Packwood, the committee chairman, made his own lobby: he persuaded a majority of the committee members to vote against *any* amendment to the tax bill, even those amendments that especially favored their own constituents. The reform was enacted. But special provisions are already staging a comeback, one or two at a time.

Along similar lines, the line-item veto would allow the president to veto legislation selectively. If a bill authorized money for school lunches and a new space shuttle, the president would have the option of neither, either, or both, instead of the current neither or both. Although a first reaction is that this allows the president greater control over legislation, the opposite might end up happening as Congress would be more selective about which bills it passes.* While the line-item veto is generally thought to be unconstitutional, this question may have to be resolved by the Supreme Court.

* Professor Douglas Holtz-Eakin of Columbia University has looked at the effects of line-item veto power at the state level. His results fail to detect any differences in spending when a line-item veto is available. This is discussed in greater detail in case study #10, following the chapter on voting.

These problems arise because myopic decision-makers fail to look ahead and see the whole picture. In the case of tax reform, the Senate recovered its vision just in time; the issue of protectionism still suffers. Chapter 2 develops a system for better long-range strategic vision.

7. Look before You Leap

It is all too common for people to get themselves into situations that are difficult to get out of. Once you have a job in a particular city, it is expensive to resettle. Once you buy a computer and learn its operating system, it becomes costly to learn another one and rewrite all your programs. Travelers who join the frequent-flyer program of one airline thereby raise their cost of using another. And, of course, marriage is expensive to escape.

The problem is that once you make such a commitment, your bargaining position is weakened. Companies may take advantage of their workers' anticipated moving costs and give them fewer or smaller salary raises. Computer companies can charge higher prices for new, compatible peripheral equipment knowing that their customers cannot easily switch to a new, incompatible technology. Airlines, having established a large base of frequent flyers, will be less inclined to engage in fare wars. A couple's agreement that they will split the housework 50:50 may become subject to renegotiation once a child is born.

Strategists who foresee such consequences will use their bargaining power while it exists, namely, before they get into the commitment. Typically, this will take the form of a payment up front. Competition among the would-be exploiters can lead to the same result. Companies will have to offer more attractive initial salaries, computer manufacturers will have to charge sufficiently low prices for their central processing units (CPUs), and airline frequent-flyer programs will have to offer

larger signing-on mileage bonuses. As for married couples, exploitation may be a game that two can play.

The same foresight is what prevents many curious but rational people from trying addictive drugs such as heroin. A Tom Lehrer song describes the drug dealer's ploy:

> "He gives the kids free samples
> Because he knows full well
> That today's young innocent faces
> Will be tomorrow's clientele."

Smart kids know it too, and turn down the free samples.

8. Mix Your Plays

Let us return for a moment to the world of sports. In football, before each snap of the ball the offense chooses between passing and running while the defense organizes itself to counter one of these plays. In tennis, the server might go to the forehand or the backhand of the receiver, while the receiver, in turn, can try to return crosscourt or down the line. In these examples, each side has an idea of its own strong points and of its opponent's weaknesses. It will have a preference for the choice that exploits these weaknesses, *but not exclusively*. It is well understood, by players and sports fans alike, that one should mix one's plays, randomly throwing in the unexpected move. The point is that if you do the same thing all the time, the opposition will be able to counter you more effectively by concentrating its resources on the best response to your one strategy.

Mixing your plays does not mean rotating your strategies in a predictable manner. Your opponent can observe and exploit any systematic pattern almost as easily as he can the unchanging repetition of a single strategy. It is *unpredictability* that is important when mixing.

3

Seeing through Your Rival's Strategy

Every week, *Time* and *Newsweek* compete to have the more eye-catching cover story. A dramatic or interesting cover will attract the attention of potential buyers at newsstands. Thus every week the editors of *Time* meet behind closed doors to select their cover story. They do so with the knowledge that the editors of *Newsweek* are meeting elsewhere, also behind closed doors, to select their cover. The editors of *Newsweek* in turn know that the editors of *Time* are making a similar decision, those of *Time* know that those of *Newsweek* know, and so on.

The two newsmagazines are engaged in a strategic game, but this game is quite different in nature from those we have already examined. The games in Chapter 2 had a sequence of alternating moves. Charlie Brown decided whether or not to kick knowing that Lucy's decision whether or not to pull the ball away lay in the future. In chess, White's moves alternated with Black's. By contrast, the actions of *Time* and *Newsweek* are simultaneous. Each must act in ignorance of the other's choice. By the time each discovers what the other has done, it is too late to change anything. Of course the loser for one week might try to respond the next week, but in this fast-moving world a whole new set of stories and a whole new game will probably have emerged by then.

The nature of the strategic thinking and action needed for the two types of games differs markedly. For the sequential-move games discussed in Chapter 2, each player had to look ahead and anticipate his rival's future responses in order to reason back and decide his own current action. There was a linear chain of reasoning: "If I do this, the other player will do that — in which case, I will respond thus," and so on.

For the simultaneous-move games we consider in this chapter, neither player has the benefit of observing the other's completed move before making his own. Here, the interaction reasoning works not by seeing the other's strategy but by seeing *through* it. For this, it is not enough simply to put yourself in your opponent's shoes. What would you find if you did? You'd only discover that your opponent is doing the same thing, that is, thinking what it must be like to be wearing your shoes. Each person has to place himself simultaneously in *both* his own and the other guy's shoes and then figure out the best moves for both sides. Instead of a linear chain of reasoning, there is a circle that goes "If I think that he thinks that I think ..." The trick is to square this circle.

Not surprisingly, Sherlock Holmes and his arch-rival Professor Moriarty, the Napoleon of crime, were masters of this type of reasoning. As Holmes told Watson in *The Final Problem*:

> "All that I have to say has already crossed your mind," said he.
>
> "Then possibly my answer has crossed yours," I replied.

Like Dr. Watson, you may be wondering how Holmes does it. After hearing our explanation, we hope you will agree that it is rather elementary.

How do you see through all the interlocking but invisible strategies? First, you must not regard the unknown actions

chance of rain, and he might use this information to decide whether or not to take an umbrella to work. The probability that *Newsweek* is using a particular cover theme is quite a different matter.

The difference is that the editor of *Time* has a very pertinent piece of information about *Newsweek*: unlike nature, the other magazine's editors are strategic game-players just as *Time*'s own editors are.* Even though one editor cannot actually observe the other magazine's decision, he can think the matter through from its perspective, and try to figure out what it *must* be doing.

In Chapter 2, we could offer a single, unifying principle to devise the best strategies for games with sequential moves. This was our Rule 1: look ahead and reason back. It won't be so simple in this chapter. But the thinking about thinking required for simultaneous moves can be summarized in three simple rules for action. These rules in turn rest on two simple ideas — dominant strategies and equilibrium. As in Chapter 2, we develop such ideas and rules through simple examples.

1. Dominant Strategies

In baseball, when there are two outs and the count stands at three balls and two strikes, any forced base runners should run on the pitch. This can be seen by thinking through all possible cases. In most cases it does not matter what the runners do. If the pitch is not touched by the batter, either the pitch is the fourth ball and the runners advance, or it is the third strike

* Some people believe that nature, too, is a strategic game-player, and a malevolent one that takes pleasure at upsetting our best-laid plans. For example, when you hear that there is a forty percent chance of rain, that means six times out of ten you will remember to take your umbrella and it won't rain, while the other four times out of ten you'll forget your umbrella and there will be a downpour.

and the inning ends. If the pitch is fouled off, the runners simply return to their original bases. If it is foul-tipped and caught, the inning ends. But in one case running has a clear advantage: if the batter hits the pitch into fair territory, the runners have a better chance of advancing or scoring.

We say that running on the pitch is the *dominant strategy* in this situation; it is better in some eventualities, and not worse in any. In general, a player has a dominant strategy when he has one course of action that outperforms all others no matter what the other players do. If a player has such a strategy, his decision becomes very simple; he can choose the dominant strategy without worrying about the rival's moves. Therefore it is the first thing one should seek.

There are interesting examples of dominant strategies everywhere, detectable once you know what to look for. Consider the position of Indiana Jones in the climax of the movie *Indiana Jones and the Last Crusade*. Indiana Jones, his father, and the Nazis have all converged at the site of the Holy Grail. The two Joneses refuse to help the Nazis reach the last step. So the Nazis shoot Indiana's dad. Only the healing power of the Holy Grail can save the senior Dr. Jones from his mortal wound. Suitably motivated, Indiana leads the way to the Holy Grail. But there is one final challenge. He must choose between literally scores of chalices, only one of which is the cup of Christ. While the right cup brings eternal life, the wrong choice is fatal. The Nazi leader impatiently chooses a beautiful golden chalice, drinks the holy water, and dies the sudden death that follows from a wrong choice. Indiana picks a wooden chalice, the cup of a carpenter. Exclaiming "There's only one way to find out" he dips the chalice into the font and drinks what he hopes is the cup of life. Upon discovering that he has chosen wisely, Indiana brings the cup to his father and the water heals the mortal wound.

Although this scene adds excitement, it is somewhat embarrassing (to us) that such a distinguished professor as Dr.

Indiana Jones would overlook his dominant strategy. He should have given the water to his father without testing it first. If Indiana has chosen the right cup, his father is still saved. If Indiana has chosen the wrong cup, then his father dies but Indiana is spared. Testing the cup before giving it to his father doesn't help, since if Indiana has made the wrong choice, there is no second chance — Indiana dies from the water and his father dies from the wound.*

Finding dominant strategies is considerably easier than the search for the Holy Grail. Consider Alfred, Lord Tennyson's familiar line: "'Tis better to have loved and lost than never to have loved at all."[1] In other words, love is a dominant strategy.

2. Over-Cover Warfare

In the competition between *Time* and *Newsweek*, think of a hypothetical week that produces two major news stories: there is an impasse between the House and the Senate on the budget, and a new drug is claimed to be effective against AIDS. The editors' choice of cover story is primarily based on what will attract the most newsstand buyers (subscribers buy the magazine no matter what the cover). Of these newsstand buyers, suppose 30 percent are interested in the budget story and 70 percent in the AIDS story. These people will buy the magazine only if the story that interests them appears on the cover; if both magazines have the same story, the group interested in it splits equally between them.

Now *Time*'s editor can reason as follows. "If *Newsweek* uses the AIDS story, then if I use the budget story I get the whole of the 'budget market' (30 percent of all readers), whereas if I use

* This example also points out one of the weaknesses of game theory. Acts are judged by their consequences alone. No moral value is placed on the act itself. Even though his father is already mortally wounded, Indiana might not want to take responsibility for the act that causes his death.

the AIDS story we share the 'AIDS market' (so I get 35 percent of all readers); so, the AIDS story yields me more sales than the budget story. If *Newsweek* uses the budget story, then I get 15 percent using the budget story, and 70 percent with the AIDS story; once again I do better using the latter. Therefore I have a dominant strategy, namely using the AIDS story. It works better for me than the other strategy regardless of which of the two courses my rival chooses."

We can see the logic of this reasoning much more quickly and clearly from a simple table. We show two columns corresponding to *Newsweek*'s choices, and two rows corresponding to *Time*'s choices. This produces four boxes; each corresponds to one combination of strategies. The entry in each box shows *Time*'s sales, measured in percentage of the total potential readership. The first row shows *Time*'s sales from choosing the AIDS story, as we range over *Newsweek*'s alternative choices. The second row shows *Time*'s sales from choosing the budget story, again as we range over *Newsweek*'s choices. For example, in the bottom left or south-west corner box, *Time* has the budget story and *Newsweek* has the AIDS story, and *Time* gets 30 percent of the market.

***Time*'s Sales**

		Newsweek's Choices: AIDS	*Newsweek*'s Choices: Budget
Time's Choices	AIDS	35	70
	Budget	30	15

The dominant strategy is easy to see. The first row is uni-

formly better than the second row: each entry in the first row is bigger than the entry immediately below it in the second row. This is the criterion for dominance. With the table, you can make a quick visual check of whether or not the criterion is met. You can figuratively lift the first row and lay it over the second, and each number in the second row will be covered by a bigger number in the first. The visual advantage of the table over the verbal reasoning of the previous paragraph grows in more complicated games, in which each side has several strategies.

It so happens that in this game, both players have a dominant strategy. To see this, draw up a table for *Newsweek*'s sales, shown below. The first column of numbers shows *Newsweek*'s sales if it uses the AIDS story, as we range over *Time*'s choices. This column is uniformly better than the second column; once again you can perform the overlaying test in your mind's eye. Therefore the AIDS story is the dominant strategy for *Newsweek*, too.

***Newsweek*'s Sales**

		Newsweek's Choices	
		AIDS	Budget
Time's Choices	AIDS	35	30
	Budget	70	15

Games in which each side has a dominant strategy are the simplest games from the strategic perspective. There is strategic interaction, but with a foregone conclusion. Each player's

choice is his dominant strategy, irrespective of what the other does. That does not make such games uninteresting, either to play or to think about. For example, in the hundred-yard dash, the dominant strategy is to run as fast as you can. But many people enjoy participating in and viewing such races. In Chapter 1's Prisoners' Dilemma, as played in Dzerzhinsky Square, both players have dominant strategies. Yet this compelling force takes them to a mutually disastrous outcome. This raises a very interesting question — how can the players cooperate to get a better result? We will have more confessions to make about this in our next chapter.

Sometimes one player has a dominant strategy but the other does not. We illustrate this with just a slight change in the cover story competition between *Time* and *Newsweek*. Suppose that the readership has a slight bias in favor of *Time*. When the two magazines have the same cover story, 60 percent of the potential buyers who like that story will pick *Time* and 40 percent will pick *Newsweek*. Now the table of *Time*'s sales is as follows:

***Time*'s Sales**

		Newsweek's Choices: AIDS	*Newsweek*'s Choices: Budget
Time's Choices	AIDS	42	70
	Budget	30	18

For *Time*, the AIDS story is still the dominant strategy. But *Newsweek*'s table becomes

***Newsweek*'s Sales**

		Newsweek's Choices: AIDS	*Newsweek*'s Choices: Budget
Time's Choices	AIDS	28	30
	Budget	70	12

If you lift the first column and lay it over the second, 30 gets covered by a smaller number (28), and 12 by a larger (70). Neither strategy dominates the other. In other words, *Newsweek*'s best choice is no longer independent of *Time*'s strategy. If *Time* chooses the AIDS story, *Newsweek* does better by choosing the budget story, and vice versa. For *Newsweek*, the whole of the budget market is now better than the smaller share of the larger AIDS market.

The editors of *Newsweek* do not observe what those of *Time* choose, but they can figure it out. Since *Time* has a dominant strategy, that must be their choice. So *Newsweek*'s editors can confidently assume that those of *Time* have chosen the AIDS story, and pick their own best response, namely the budget story.

Thus games in which only one side has a dominant strategy are also very simple. This side plays its dominant strategy, and the other chooses its best response to that strategy.

Now that we have introduced the idea of a dominant strategy, it is worth emphasizing two points about what a dominant strategy is not. It is easy to get confused about just what it is that a dominant strategy actually dominates.

In 1981, Leonard Silk, writing about the Congressional debate on the Economic Recovery Tax Act, concluded: "Mr. Reagan has sensed that the Republicans have what game theorists call a 'dominant strategy' — one that makes a player better off than his opponent, no matter what strategy his opponent uses."[2] We will look at this game more carefully in Chapter 5, but here we only want to point out that Silk's definition of a dominant strategy is incorrect. The dominance in "dominant strategy" is a dominance of one of your strategies over your other strategies, not of you over your opponent. A dominant strategy is one that makes a player better off than *he* would be *if he used any other strategy*, no matter what strategy his opponent uses. Recall that in the cover picture example, both *Time* and *Newsweek* have a dominant strategy; yet both cannot have higher sales than the other.

A second common misperception is that a dominant strategy requires that the worst possible outcome playing the dominant strategy be better than the best outcome of some other strategy. This happens to be true in the examples above. With the numbers in the original setup, the worst that could happen to *Time* when using the AIDS story was a 35 percent share; the best they could hope for with the budget story was 30 percent. However, this is not a general feature of dominant strategies.

Imagine a price war between *Time* and *Newsweek*. Suppose each issue costs $1 to produce, and there are just two possible pricing choices: $3 (implying a profit margin of $2 per copy) and $2 (implying a profit margin of $1 per copy). Suppose that customers will always buy the lower-priced magazine, and if the prices are equal, they will split equally between the two. The total readership is 5 million if the price is $3, and rises to 8 million if the price is only $2. You can easily calculate *Time*'s profits in the four possible pricing combinations, and produce the following table.

***Time*'s Profits**

Time's Price \ *Newsweek*'s Price	$2	$3
$2	4	8
$3	0	5

Time's dominant strategy is to charge $2 (and so is *Newsweek*'s). The worst that can happen to *Time* from following the dominant strategy is to net $4 million. But the best that can happen from following the other strategy is better, namely $5 million. The point is that the comparison of those two numbers is meaningless. The $5 million arises if both magazines charge $3; then *Time* would do even better ($8 million) by switching to $2.

We can sum up the lessons of these examples into a rule for behavior in games with simultaneous moves:

Rule 2: **If you have a dominant strategy, use it.**

Do not be concerned about your rival's choice. If you do not have a dominant strategy, but your rival does, then anticipate that he will use it, and choose your best response accordingly.

A word of caution. We developed the concept of a dominant strategy for games with simultaneous moves. Care must be taken in using it if moves are sequential. Because the nature of the strategic interaction is different, the idea of a dominant strategy is no longer the same. Suppose we say that you have a dominant strategy if for *each given* choice of the rival, you do better with this strategy than with any other. When moves

are sequential and your rival moves first, you would always choose your dominant strategy. As we just said, it is your best response to *each* of your rival's moves, and therefore to the particular one he has chosen. But *if you move first*, your rival's move is not *given*. He will observe your choice when he makes his, and you have the opportunity to influence his behavior. In some circumstances this may best be done by choosing something other than your dominant strategy. We explain this fully in Chapter 6, when we discuss commitment.

3. Dominated Strategies

Not all games have dominant strategies, even for one player. In fact, dominance is the exception rather than the rule. Although the presence of a dominant strategy leads to very simple rules for action, these rules are inapplicable to many actual games. Other principles must then be brought into action.

Just as a dominant strategy is uniformly better than every other strategy, a dominated strategy is uniformly worse than some other strategy. Just as you choose your dominant strategy if you have one, and can be sure that your rival will choose his if he has one, you should avoid your dominated strategies if you have any, and can be sure that your rival will avoid his, if he has any.

If you have just two alternative strategies, and one of them is dominated, then the other must be dominant. Therefore examples of avoiding dominated strategies that are genuinely different from those of choosing dominant strategies must be based on games in which at least one side has at least three strategies. Let us consider a simple example of this kind.

Think of a play in football in which the offense's sole concern is to gain as many yards as possible, and the defense's sole concern is to hold them to as few yards as possible. For example, with very little time left, the offense may want to

improve its chances of kicking a winning field goal.

Suppose the offense has just two strategies, run and pass, while the defense has three strategies: counter the run, counter the pass, and blitz the quarterback. We can calculate the yards likely to be gained by the offensive team for each of the six strategy combinations. For example, take the case in which the defense blitzes and the offense tries a pass. Suppose there is a 10 percent chance that the quarterback will be sacked for a loss of 10 yards, a 70 percent chance of a quick 10-yard pass, and a 20 percent chance of a longer 20-yard pass. The average works out at

$$0.1 \times (-10) + 0.7 \times 10 + 0.2 \times 20 = -1 + 7 + 4 = 10.$$

The numbers obviously depend on the special skills (or lack thereof) of the two teams; we have chosen particular ones just for illustration.*

We show the outcomes of such calculations for all six possible combinations in the following table.

Offense's Expected Yardage Gain

Offense's Strategies	Defense's Strategies: Counter Run	Counter Pass	Blitz
Run	3	7	15
Pass	9	8	10

The offense tries to achieve the largest possible number in

* In our example the offensive team is strong at passing and weak at running. That is why the pass does better than the run even against a pass defense. The run does better against the blitz only because the defensive backs are not in position.

this table. The defense tries to secure the smallest possible number, so we do not need a separate table from which to determine their actions.*

Neither side has a dominant strategy: there is no row with numbers uniformly higher than those in the other row, and no column with numbers uniformly smaller than those in each of the other columns. But the defense does have a dominated strategy, namely the blitz. The result of a blitz is a yardage loss that is uniformly larger, and thus worse for the defense, than those possible with either of the other strategies. Therefore this defense should not blitz, and the offense can be confident that they will not.

The argument doesn't stop there. The blitz strategy might as well be removed from the defensive coach's playbook, and the game can be treated as if each side had two strategies. In this reduced game, the offense has a dominant strategy, namely pass. Its numbers, 9 and 8, are uniformly higher than those of the run strategy — 3 and 7, respectively. The reason pass was not dominant in the original game was that run had a better yield against the defense's blitz (as the ball-carrier might break into open field with the blitzing defensive safeties out of position), but that has now been removed from consideration. So the offense will choose the pass. The defense in turn should think this through, and choose its best response, namely the pass defense.

The general idea can be summed up into one more rule of behavior for games with simultaneous moves:

> *Rule 3*: **Eliminate any dominated strategies from consideration, and go on doing so successively.**

If, during the process, any dominant strategies emerge in the smaller games, they should be chosen successively. If this pro-

* This is so in all "zero-sum" games, in which one side's gain is exactly the other side's loss.

cedure ends in a unique outcome, you have found the prescriptions of action for the players and the outcome of the game. Even if the procedure does not end in a unique outcome, it will reduce the size and the complexity of the game.

We illustrate the idea of successively eliminating dominated strategies by making up a story of an impending naval engagement in the Persian Gulf.* The grid below shows the positions and the choices of the combatants. An Iraqi ship at the point *I* is about to fire a missile, intending to hit an American ship at *A*. The missile's path is programmed at the launch; it can travel in a straight line, or make sharp right-angled turns every 20 seconds. If the Iraqi missile flew in a straight line from *I* to *A*, American missile defenses could counter such a trajectory very easily. Therefore the Iraqis will try a path with some zigzags. All such paths that can reach *A* from *I* lie along the grid shown. Each length like *IF* equals the distance the missile can travel in 20 seconds.

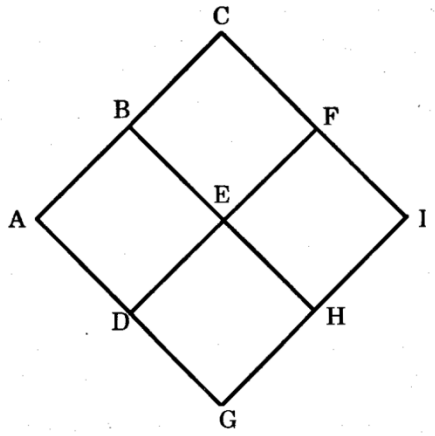

The American ship's radar will detect the launch of the in-

* This story is an updated version of the cat-and-mouse story in J. D. Williams, *The Compleat Strategyst*. Perhaps the cat was Persian.

4

Resolving the Prisoners' Dilemma

Throughout the 1970s, the Organization of Petroleum Exporting Countries (OPEC) colluded to raise the price of crude oil from under $3 per barrel in 1973 to over $30 per barrel in 1980. The world awaited the meeting of each OPEC price-setting meeting with anxiety. By the end of the 1970s, some energy experts were predicting that the price of oil would rise to over $100 per barrel by the end of the century. Then suddenly the cartel seemed to collapse. Prices moved down, briefly touching $10 per barrel in early 1986 before recovering to $18 per barrel in 1987.* As we write this, the Iraqi invasion of Kuwait has shot the price of oil up to $35 per barrel and experts are divided about the future of OPEC.

What governs the success or failure of such cartels? More generally, what governs the balance between cooperation and competition not just in business, but also in politics and social settings? This question can be answered, at least in part, using the prisoners' dilemma that we played out in KGB headquarters in Chapter 1.

The story of OPEC is just such a game. Of course we tell

* Of course it must be remembered that the dollar rose sharply against other currencies from 1981 to 1985, and fell almost equally fast from 1985 to 1987. Therefore neither the drop in oil prices in the first half of the 1980s, nor the recovery since then, were as dramatic in terms of an average of all currencies as they were in dollars alone.

it in a stylized way, highlighting the dilemma and leaving out many historical details. To start with, look at the production decisions of just two members, say Iran and Iraq. For further simplicity, allow each just two production levels, either 2 or 4 million barrels of crude oil a day. Depending on their decisions, the total output on the world market will be 4, 6, or 8 million barrels. Suppose the price will be \$25, \$15, and \$10 per barrel, respectively. Extraction costs are \$2 per barrel in Iran and \$4 per barrel in Iraq. Then we can show the profits (measured in millions of dollars a day) of the two competitors in the usual table. In each box, the top right entry is Iraq's daily profit, the bottom left is Iran's.*

Table of Profits (Iran, Iraq)

Iran's Output	Iraq's Output: 2	Iraq's Output: 4
2	42 / 46	44 / 26
4	22 / 52	24 / 32

Each country has a dominant strategy: produce at the higher of the two available levels. Iran, for example, sees that its profit row corresponding to the production level of 4, namely [\$52 and \$32], is uniformly higher than the one corre-

* This way of representing both players' payoffs in the same matrix is due to Thomas Schelling. With excessive modesty he writes, "If I am ever asked whether I ever made a contribution to game theory, I shall answer yes. Asked what it was, I shall say the invention of staggered payoffs in a matrix. ... I did not suppose that the invention was patentable, so I made it freely available and hardly anybody except my students takes advantage. I offer it to you free of charge."

sponding to the production level of 2, namely [$46 and $26]. When they both choose their dominant strategies, their profits are $32 and $24 million a day, respectively. Nothing to sneeze at, but cooperation would have gotten them more, $46 and $42.

This predicament is called the prisoners' dilemma. Its remarkable feature is that both sides play their dominant strategy, thus maximize their payoff, and yet the outcome is jointly worse than if both followed the strategy of minimizing their payoff. So why don't they follow the minimizing strategy? Look back at the problem for Iran and Iraq. Even if Iran were to follow the minimizing strategy of producing 2 million barrels, Iraq still has an incentive to produce 4 million. Then the outcome would be Iraq's ideal and Iran's worst. If Iran doesn't cooperate and produces 4 million, then Iraq would be foolish to sacrifice its own profits by producing 2 million. The cartel's problem is to find a way to sustain the low-output, high-price strategy that yields the highest joint profit, given the temptation for each to cheat and gain at the expense of the other.

Iran and Iraq's situation is analogous to that of the KGB's prisoners. Each of them found it dominant to confess: if the one held out, the other got a better deal by confessing; if one confessed, the other would be foolish not to. Hence whatever one does, the other wants to confess. But that's true for both. And when both confess, each gets a harsh sentence. Again the selfish pursuit of one's interests leads to an inferior outcome. When neither confesses, the outcome is better for both. The problem is how to attain such cooperation given the competition to obtain an especially good deal for oneself.

The same problem arises when there are several competing firms in the industry. The problem plagues not just businesses, but also students of business. A professor at Texas A&M University had his class of 27 students play a game that trapped them in the prisoners' dilemma.[1] Each student owned a hypothetical firm and had to decide whether to produce 1 and help

keep the price high or produce 2 and gain at the expense of others. Depending on the total number of students producing 1, money would be paid to students according to the following table:

Number of Students Writing "1"	Payoff to Each Student Who Writes "1"	Payoff to Each Student who Writes "2"
0		$0.50
1	$0.04	$0.54
2	$0.08	$0.58
3	$0.12	$0.62
. . .	. . .	. . .
25	$1.00	$1.50
26	$1.04	$1.54
27	$1.08	

This is easier to see and more striking in a chart:

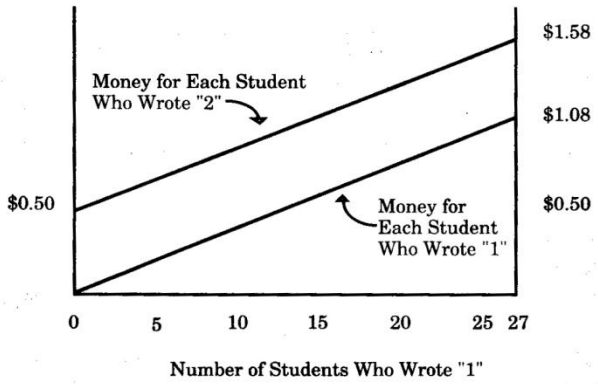

The game is "rigged" so that students who write 2 always get 50 cents more than those who write 1, but the more of

them that write 2, the less their collective gain. Suppose all 27 start planning to write 1, so each would get $1.08. Now one thinks of sneaking a switch to 2. There would be 26 1's, and each would get $1.04 (4 cents less than the original), but the switcher would get $1.54 (46 cents more). The same is true irrespective of the initial number of students thinking of writing 1 versus 2. Writing 2 is a dominant strategy. Each student who switches from writing 1 to writing 2 increases his own payout by 46 cents, but decreases that of each of his 26 colleagues by 4 cents — the group as a whole loses 58 cents. By the time everyone acts selfishly, each maximizing his own payoff, they each get 50 cents. If, instead, they conspired and acted so as to minimize their individual payoff, they would each receive $1.08. How would you play?

In some practice plays of this game, first without classroom discussion and then with some discussion to achieve a "conspiracy," the number of cooperative students writing 1 ranged from 3 to a maximum of 14. In a final binding play, the number was 4. The total payout was $15.82, which is $13.34 less than that from totally successful collusion. "I'll never trust anyone again as long as I live," muttered the conspiracy leader. And how did he vote? "Oh, I voted 2," he replied.

This situation reminds us of Yossarian's position in Joseph Heller's novel *Catch-22*. The Second World War was nearly won, and Yossarian did not want to be among the last to die. His commanding officer asks, "But suppose everyone on our side felt that way?" and Yossarian replies, "Then I'd certainly be a damned fool to feel any other way. Wouldn't I?"

Politicians, too, are prisoners of the same dilemma. In 1984, it was clear to most people that the U.S. federal budget deficit was too large. Expenditure cuts of the required magnitude were politically infeasible, and therefore a significant tax increase was inevitable. But who was going to exercise the political leadership necessary to bring this about? The Democratic presidential candidate, Walter Mondale, tried to

set the stage for such a policy change in his campaign, and was soundly defeated by Ronald Reagan, who promised no tax increase. In 1985, the issue got stalled. No matter how you formed the political divisions — Democrats vs. Republicans, the House of Representatives vs. the Senate, or the Administration vs. the Congress — each side preferred to leave the initiative to the other.

For each, the best outcome was one in which the other proposed the tax increases and expenditure cuts, paying the political price. Conversely, proposing such policies oneself while the other remained passive was the worst outcome. Both sides agreed that the exercise of joint leadership, sharing the credit and the blame, would be better for the country, and even for themselves in the long run, than the combination in which both were passive and the large deficit continued.

We can represent this as a game by drawing up the usual table of strategies and outcomes. The two sides are the Republicans and the Democrats. To show who prefers what, let us rank the outcomes from 1 to 4 from each side's point of view. Low numbers mean better ranking. In each box the lower left number is the Republicans' ranking; the upper right, the Democrats'.

Rankings for Republicans and Democrats

		Democrats: Active	Democrats: Passive
Republicans: Active		2 / 2	1 / 4
Republicans: Passive		4 / 1	3 / 3

You can easily see that for each side, passivity is the dominant strategy. This is just what happened; there was no move

toward tax increase in the 99th Congress. The 99th Congress did pass the Gramm-Rudman-Hollings act, which mandated the deficit-reduction policies to be followed in *future* years. But that was merely a pretense of activity, in fact postponing the hard choices. Its targets are being met more by accounting tricks than by genuine fiscal restraint.

1. How to Achieve Cooperation

Those who find themselves in a prisoners' dilemma will look for ways to escape and achieve the cooperative outcome they jointly prefer. Others may like to see the players remain trapped in the dilemma. For example, buyers benefit from lower prices when the prisoners' dilemma prevents firms in an industry from colluding. In this case society wants to thwart the industry's attempts to resolve the dilemma, and antitrust laws are part of this effort. In either case, whether we seek collusion or its opposite, we must first understand the ways in which the prisoners' dilemma might be averted. Then we can try to facilitate these ways, or to block them, as is appropriate in the case being considered.

The underlying problem is the players' incentive to cheat on any agreement. Therefore the central questions are, How can such cheating be detected? What prospect of punishment will deter it? Let us examine these in turn.

2. Detection of Cheating

A cartel has to find ways to discover if cheating has in fact occurred, and if so, then determine who has cheated. Recognizing that *someone* has cheated is often easy in the examples we have used. In the case of Iran and Iraq's oil production, the price will be $25 only if both countries cooperate and pro-

duce 2 million barrels daily; any price below $25 per barrel is a sure indicator of cheating. In reality, matters are more complicated. The price can be low either because of a fall in demand or because of cheating by a producer. Unless the cartel can sort out these separate influences and determine the truth, it might infer cheating and set in motion its punishment measures when no cheating has in fact occurred, or err the other way around.* This will reduce the accuracy and therefore the efficacy of the measures. A compromise solution is a critical or "trigger" price; if the price falls below this value, the cartel presumes that cheating has occurred and the punishment ensues.

There is yet another complication in reality. Games of this kind often have many dimensions of choice, and the possibility of observing cheating differs among them. For example, firms compete with one another in price, product quality, after-sales service, and many other aspects. The price is relatively easy to observe, although secret discounts or flexibility in pricing trade-ins can cause complications. There are many dimensions of quality that are hard to monitor. Therefore a cartel that tries to enforce collusive high prices finds competition continually breaking out in new dimensions. This happened in the airline industry. During the years of regulation, fares were fixed and entry of new competitors was effectively barred. This was as if the airlines formed a cartel with enforcement provided by the Civil Aeronautics Board. Airlines began to compete, or cheat on the cartel. While they couldn't lower prices, they could provide more valuable services through elaborate meals and beautiful stewardesses. When labor laws forced airlines to hire male stewards and not fire stewardesses over thirty, competition switched to nonstop schedules, seat width, and leg room.

* The statistical literature describes false positives as Type I errors and false negatives as Type II errors. Most common of all is the Type III error: not being able to remember which is which.

Another instance of this process occurred in the area of international trade policy. Tariffs are the most visible tools for restricting trade, and successive rounds of negotiations of the General Agreement on Tariffs and Trade (GATT) achieved large mutual reductions of tariff rates of all industrial countries. But each country still had its domestic political pressures from powerful special interests to restrict imports. Therefore countries gradually switched to other, less visible means, such as voluntary restraint agreements, customs valuation procedures, standards, administrative practices, and complicated quotas.*

The common theme of these examples is that collusion focuses on the more transparent dimensions of choice, and competition shifts to the less observable ones: we call this the Law of Increasing Opaqueness. Though you might not see it clearly, the collusion still hurts you. When quotas on Japanese auto imports went into effect in 1981, not only did the prices of all cars, Japanese and American, go up, but the low-end Japanese models disappeared from the market. Opaque competition was doubly bad: prices were higher, and the balance of the product mix was distorted.

Identifying the cheater may be even more difficult than detecting cheating. With just two players, an honest party knows that the other has cheated. There may still be a problem with getting him to admit his fault. With more than two players, we may know that someone has cheated, but no one (other than the cheater) knows who. In this case, the punishment to deter cheating must be blunt and affect the innocent and the guilty alike.

Finally, cheating may consist of remaining passive and may

* For example, quotas under the multifiber arrangement are levied by extremely complicated categories of garments and countries. This makes it very hard to see the effect of the quota in raising the price of any particular good. Economists have estimated these effects and found price increases as high as 100 percent; a tariff this high would surely arouse louder consumer protests.

thereby be difficult to isolate. This was so in our example of the exercise of leadership in proposing higher taxes. In such a case, it is far harder to infer or allege cheating. While positive action is there for all to see, there are numerous excuses for inaction: greater urgency of other issues, time needed to consolidate forces, and so on.

3. Punishment of Cheaters

Behind every good scheme to encourage cooperation is usually some mechanism to punish cheaters. A prisoner who confesses and implicates his collaborators may become the target of revenge by the others' friends. The prospect of getting out of prison more quickly may look less alluring given the knowledge of what waits outside. Police have been known to scare drug dealers into confessing by threatening to release them. The threat is that if they are released, their suppliers will assume they have squealed.

In the example of the Texas A&M classroom experiment, if the students could detect who had reneged on the conspiracy for all of them to write 1, they could ostracize the cheaters for the rest of the semester. Few students would risk that for the sake of fifty cents. In OPEC, because of the social and political cohesion of the Arab states in the 1970s, a country thinking of cheating may have been deterred by a fear of ostracism. These are examples of punishments that are added on to the original game, in order to reduce the incentive to cheat.

Other kinds of punishments arise within the structure of the game. Usually this happens because the game is repeated, and the gain from cheating in one play will lead to a loss in other plays. We illustrate this using the crude oil example with Iran and Iraq.

The possibility of punishment arises because the two countries are involved in this game day after day. Suppose they

start on a basis of trust, each producing 2 million barrels a day and helping keep the price high. Each will constantly be tempted to sneak in a defection. Look again at the table of daily profits. A successful day of cheating while Iraq stays honest will raise Iran's profit from $46 million to $52 million, a gain of $6 million.

Table of Profits (Iran, Iraq)

Iran's Output	Iraq's Output: 2	Iraq's Output: 4
2	42 46	44 26
4	22 52	24 32

The question is what happens when Iraq recognizes what has gone on. A plausible scenario is that the mutual trust will collapse, and the two will settle down to a regime of high outputs and low prices from that day onward. Relative to the continuation of trust, this gets Iran $14 million a day (46 32) less profit. The short-term gain from cheating seems small in comparison with the subsequent cost: if it takes Iraq a month to detect Iran's cheating and respond, the month's extra profit to Iran ($180 million) would be wiped out if the period of collapsed trust lasts just 13 days. Of course time is money, and higher profits today are worth more than an equal reduction of profit in the future; but still this calculation looks distinctly unfavorable. For Iraq, breaking the cartel is even worse; the daily gain while its cheating goes undetected and unpunished is $2 million, whereas the daily cost once trust collapses is $18 million. It appears that in this instance, even a slight fear of

the collapse of their mutual trust should be enough to keep the two competitors abiding by the agreement.

Trust can break down for all sorts of reasons. For example, the war between Iran and Iraq made it difficult for OPEC to impose production quotas on either country. Trust in maintaining cartel quotas is based on the subsequent ability to punish those who violate the agreement. But what additional punishments could be imposed on two countries already punishing each other with explosives and "human wave" assaults? With the war ended, there is once again a potential for cooperation because there is a potential for punishment.

To sum up, there is no solution that achieves reciprocal cooperation in a one-time game. Only in an ongoing relationship is there an ability to punish, and thus a stick to motivate cooperation. A collapse of cooperation carries an automatic cost in the form of a loss of future profits. If this cost is large enough, cheating will be deterred and cooperation sustained.

There are some caveats to this general principle. The first arises when the relationship has some natural end, such as the end of a term in an elected office. In these situations, the game is repeated only a fixed number of times. Using the principle of looking ahead and reasoning back, we see that cooperation must end when there is no longer any time left to punish. Yet neither wants to be left cooperating while the other cheats. If ever someone cooperates, then someone must get stuck in the end. Since neither is willing to play the fool, cooperation never gets started. This is true no matter how long the game is, provided the end is known.

Let us look at this argument a little more carefully. Right from the start, both players should look ahead to predict the last play. On this last play, there will be no future to consider, and the dominant strategy is to cheat. The outcome of the last play is a foregone conclusion. Since there is no way to affect the last play of the game, the penultimate play effectively becomes the last one to consider.

Once again, cheating is a dominant strategy. The reason is that the play in the next-to-last period has no effect on the strategies chosen in the final period. Thus the penultimate period can be considered in isolation. For any period in isolation, cheating is a dominant strategy.

Now the play of the final two periods can be taken as given. Cooperating early on won't help, as both players are resigned to cheat in the final two periods. Hence, the third-to-last period is effectively the last one to consider. The same argument applies and cheating is a dominant strategy. This argument unwinds all the way back, so that there is no cooperation even in the first play.

The logic of this argument is impeccable, and yet in the real world we find episodes of successful cooperation. There are various ways to explain this. One is that all actual games of this kind are repeated only a finite number of times, but that number is unknown. Since there is no fixed last time, there is always the possibility that the relationship will go on. Then the players have some incentive to sustain the cooperation for the sake of such future contingencies; if this incentive is large enough, the cooperation will persist.

Another explanation is that the world contains some "nice" people who will cooperate no matter what the material advantages of cheating may be. Now suppose you are not so nice. If you behaved true to your type in a finitely repeated game of prisoners' dilemma, you would start cheating right away. That would reveal your nature to the other player. To hide the truth (at least for a while) you have to behave nicely. Why would you want to do that? Suppose you started by acting nicely. Then the other player would think it possible that you are one of the few nice people around. There are real gains to be had by cooperating for a while, and the other player would plan to reciprocate your niceness to achieve these gains. That helps you, too. Of course you are planning to sneak in a defection near the end of the game, just as the other player is.

But you two can still have an initial phase of mutually beneficial cooperation. Thus while each side is waiting to take advantage of the other, both are benefiting from this mutual deception.

A third qualification to the emergence of trust in a repeated prisoners' dilemma is that the gains from cheating take place before the costs of the breakdown of cooperation occur. Therefore the relative importance of the two depends on the relative importance of the present versus the future. In business contexts, current and future profits are compared using an appropriate interest rate to discount the future. In politics, the judgment of present versus future is more subjective, but it seems that time beyond the next election counts for very little. This makes cooperation hard to achieve. Even in business, when times are bad, the whole industry is on the verge of collapse, and the management feels that there is no tomorrow, competition may become more fierce than in normal times. Similarly, the needs of war made current profits more important to Iran and Iraq, and contributed to the difficulties of OPEC.