

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)

		Iraq's Output	
		2	4
Iran's Output	2	42 44	46 26
	4	22 24	52 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	Passive
Republicans	Active	2 2	1 4
	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)

		Iraq's Output	
		2	4
Iran's Output	2	42	44
	4	22	24

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.

4. THE PUNISHMENT IS GUARANTEED

The neatest trick is enforcing price collusion through a punishment guarantee, all in the name of competition. Here we turn to New York City and its stereo wars. Crazy Eddie has made his trademark "We cannot be undersold. We will not be undersold. Our prices are the lowest — guaranteed. Our prices are insane." His main competitor, Newmark & Lewis, is no less ambitious. With any purchase, you get the store's "Lifetime low-price guarantee." It promises to rebate double the difference if you can find a lower price elsewhere.

"If, after your purchase, you find the same model advertised or available for sale for less (confirmed printed proof required) by any other local stocking merchant, in this marketing area, during the life-

time of your purchase, we, Newmark & Lewis, will gladly refund (by check) 100% of the difference, plus an additional 25% of the difference, or if you prefer, Newmark & Lewis will give you a 200% gift certificate refund (100% of the difference plus an additional 100% of the difference, in gift certificates)."

— from Newmark & Lewis's Lifetime Low-Price Guarantee

Yet, although they sound competitive, these promises to beat the rival's price can enforce discipline in a price-setting cartel. How can this happen?

Suppose each VCR costs \$150 wholesale, and for the moment both Crazy Eddie and Newmark & Lewis are selling it for \$300. Crazy Eddie is contemplating a sneaky cut to \$275. Without the beat-the-rival promise, Crazy Eddie would hope that his lower price would attract some of the customers who would otherwise have gone to his rival — say, because they lived nearer to a Newmark & Lewis outlet, or had bought from them before. Unfortunately for Crazy Eddie, his price cut has the reverse effect. With the Newmark & Lewis price guarantee, these people are now tempted just to walk over to Newmark & Lewis and buy the VCR for \$300 and then claim a \$50 rebate. This is just as if Newmark & Lewis had reduced its price to \$250, automatically undercutting Crazy Eddie. But of course Newmark & Lewis would prefer not to give away the fifty dollars. Its response will be to lower the price to \$275. In any event, Crazy Eddie is worse off than where he started. So why bother? The price stays at \$300.

Although cartels are illegal in the United States, Crazy Eddie and Newmark & Lewis have the makings of one. You can see how their implicit cartel works in terms of the requirements of enforcement we mentioned before: detection of cheating, and punishment of cheaters. Newmark & Lewis can more easily detect Crazy Eddie's cheating. The customers who bring them the news of Crazy Eddie's lower price, and ask them to beat that, are acting as unwitting enforcement agents for the cartel. The punishment comes in the form of the collapse of

the pricing agreement and consequently lower profits. The "beat-the-competition" ads also set the punishment in motion, automatically and quickly.

A celebrated antitrust case before the Federal Trade Commission concerned the use of a similar device that appears to make competition more fierce, but can in fact serve as a cartel enforcement mechanism. E. I. Du Pont, Ethyl, and other manufacturers of antiknock gasoline additives were charged with using a "most-favored-customer" clause. This clause says that the seller will offer to these favored customers the best price they offer to anyone. Taken at face value, it seems that the manufacturers are looking out for their favored customers. But let's look deeper. The clause means that the manufacturer cannot compete by offering selective discounts to attract new customers away from his rival, while charging the old higher price to his established clientele. They must make general price cuts, which are more costly, because they reduce the profit margin on all sales. You can see the advantage of this clause to a cartel: the gain from cheating is less, and the cartel is more likely to hold.

In evaluating most-favored-customer clauses, the Federal Trade Commission ruled that there was an anticompetitive effect, and forbade the companies from using such clauses in their contracts with customers.* How would you rule if such a case were brought against Crazy Eddie and Newmark & Lewis? One yardstick by which to judge the fierceness of competition is the level of markups. Many "discount" stereo stores charge almost a hundred-percent markup over the wholesale cost of their components. It is hard to say what part of the markup is due to the costs of carrying inventory and adver-

* This ruling was not without some controversy. The Commission's chairman, James Miller, dissented. He wrote that the clauses "arguably reduce buyers' search costs and facilitate their ability to find the best price-value among buyers." For more information, see "In the matter of Ethyl Corporation et al." FTC Docket 9128, *FTC Decisions*, pp. 425-686.

tising, but there is at least a prima facie case that there is method to Crazy Eddie's madness.

5. A CHOICE OF PUNISHMENT

When several alternative punishments could deter cheating and sustain cooperation, how should one choose among them? Several criteria have a role.

Perhaps most important are simplicity and clarity, so that a player thinking of cheating can easily and accurately calculate its consequences. A criterion that infers someone has cheated if your discounted mean of profits from the last seventeen months is 10 percent less than the average real rate of return to industrial capital over the same period, for example, is too complicated for most firms to figure out, and therefore not a good deterrent.

Next comes certainty. Players should have confidence that defection will be punished and cooperation rewarded. This is a major problem for the European countries looking to enforce the General Agreement on Tariffs and Trade. When one country complains that another has cheated on the trade agreement, GATT initiates an administrative process that drags on for months or years. The facts of the case have little bearing on the judgment, which usually depends more on dictates of international politics and diplomacy. Such enforcement procedures are unlikely to be effective.

Next we ask how severe a punishment should be. Most people's instinctive feeling is that it should "fit the crime." But that may not be big enough to deter cheating. The surest way to deter cheating is to make the punishment as big as possible. Since the punishment threat succeeds in sustaining cooperation, it should not matter how dire it is. The fear keeps everyone from defecting, hence the breakdown never actually occurs and its cost is irrelevant.

The problem with this approach is that it ignores the risk of mistakes. The detection process may go wrong, indicating cheating by a member of the cartel when the real cause of low prices is an innocent one such as low demand. If punishments are as big as possible, then mistakes will be very costly. To reduce the cost of mistakes, the punishment should be the smallest size that suffices to deter cheating. Minimal deterrence accomplishes its purpose without imposing any extra costs when the inevitable mistakes occur.

6. TIT-FOR-TAT

This list of the desirable properties of a punishment mechanism looks quite demanding. But University of Michigan political scientist Robert Axelrod claims that the rule of *tit-for-tat* does very well in meeting these demands.² *Tit-for-tat* is a variation of the “eye for an eye” rule of behavior: do unto others as they have done onto you.* More precisely, the strategy cooperates in the first period and from then on mimics the rival’s action from the previous period.

Axelrod argues that *tit-for-tat* embodies four principles that should be evident in any effective strategy: clarity, niceness,

* In Exodus (21:22), we are told, “If men who are fighting hit a pregnant woman and she gives birth prematurely but there is no serious injury, the offender must be fined whatever the woman’s husband demands. But if there is a serious injury, you are to take life for a life, eye for eye, tooth for tooth, hand for hand, burn for burn, wound for wound, bruise for bruise.” The New Testament suggests more cooperative behavior. In Matthew (5:38) we have, “You have heard that it was said, ‘Eye for Eye, and Tooth for Tooth.’ But I tell you, do not resist an evil person. If someone strikes you on the right cheek, turn to him the other also.” We move from “Do unto others as they have done onto you” to the golden rule, “Do unto others as you would have them do unto you” (Luke 6:31). If people were to follow the golden rule, there would be no prisoners’ dilemma. And if we think in the larger perspective, although cooperation might lower your payoffs in any particular game, the potential reward in an afterlife makes this a rational strategy even for a selfish individual.

provocability, and forgivingness. *Tit-for-tat* is as *clear* and simple as you can get. It is *nice* in that it never initiates cheating. It is *provocable*, that is, it never lets cheating go unpunished. And it is *forgiving*, because it does not hold a grudge for too long and is willing to restore cooperation.

Axelrod confirmed the power of *tit-for-tat* through experiment, not just theory. He staged a tournament of two-person prisoners’-dilemma games. Game theorists from around the world submitted their strategies in the form of computer programs. The programs were matched against each other in pairs to play a prisoners’-dilemma game repeated 150 times. Contestants were then ranked by the sum of their scores.

The winner was Anatol Rapoport, a mathematics professor at the University of Toronto. His winning strategy was *tit-for-tat*. Axelrod was surprised by this. He repeated the tournament with an enlarged set of contestants. Once again Anatol Rapoport submitted *tit-for-tat* and beat the competition.

One of the impressive features about *tit-for-tat* is that it did so well overall even though it did not (nor could it) beat any one of its rivals in a head-on competition. At best, *tit-for-tat* ties its rival. Hence if Axelrod had scored the competition as a winner-take-all contest, *tit-for-tat* would have scored below .500 and so could not have won.

But Axelrod did not score the pairwise plays as winner-take-all: close counted. The big advantage of *tit-for-tat* is that it always comes close. At worst, *tit-for-tat* ends up getting beaten by one defection; i.e., it gets taken advantage of once and then ties from then on. The reason *tit-for-tat* won the tournament is that it usually managed to encourage cooperation whenever possible while avoiding exploitation. The other entries were either too trusting and open to exploitation or too aggressive and knocked one another out.

In spite of all this, we believe that *tit-for-tat* is a flawed strategy. The slightest possibility of misperceptions results in a complete breakdown in the success of *tit-for-tat*. This flaw

was not apparent in the artificial setting of a computer tournament, because misperceptions did not arise. But when tit-for-tat is applied to real-world problems, misperceptions cannot be avoided and the result can be disastrous.

For instance, in 1987 the United States responded to the Soviet spying and wiretapping of the U.S. embassy in Moscow by reducing the number of Soviet diplomats permitted to work in the United States. The Soviets responded by withdrawing the native support staff employed at the U.S. Moscow embassy and placed tighter limits on the size of the American delegation. As a result, both sides found it more difficult to carry out their diplomatic functions. Another series of tit-for-tat retaliations occurred in 1988, when the Canadians discovered spying on the part of the visiting Soviet diplomats. They reduced the size of the Soviet delegation and the Soviets reduced the Canadian representation in the Soviet Union. In the end, both countries were bitter, and future diplomatic cooperation was more difficult.

The problem with tit-for-tat is that any mistake “echoes” back and forth. One side punishes the other for a defection, and this sets off a chain reaction. The rival responds to the punishment by hitting back. This response calls for a second punishment. At no point does the strategy accept a punishment without hitting back. The Israelis punish the Palestinians for an attack. The Palestinians refuse to accept the punishment and retaliate. The circle is complete and the punishments and reprisals become self-perpetuating.

The long-standing feuds between the Hatfields and the McCoys or Mark Twain’s Grangerfords and Shepherdsons offer more examples of how tit-for-tat behavior leads to mutual loss. Feudists on either side are not willing to end the feud until they consider themselves even. But in a continuing attempt to get even, they end up knocking each other further and further down. Eventually they’ll end up dead even. Rarely is there any hope of going back and solving the dispute at its origin, for

once begun, it takes on a life of its own. When Huck Finn tries to understand the origins of the Grangerfords-Shepherdsons feud, he runs into the chicken-or-egg problem:

“What was the trouble about, Buck? — Land?”

“I reckon maybe — I don’t know.”

“Well, who done the shooting? Was it a Grangerford or a Shepherdson?”

“Laws, how do I know? It was so long ago.”

“Don’t anyone know?”


“Oh yes, pa knows, I reckon, and some of the other old people, but they don’t know now what the row was about in the first place.”

What tit-for-tat lacks is a way of saying “Enough is enough.” It is dangerous to apply this simple rule in situations in which misperceptions are endemic. Tit-for-tat is *too easily provoked*. You should be more forgiving when a defection seems to be a mistake rather than the rule. Even if the defection was intentional, after a long-enough cycle of punishments it may still be time to call it quits and try reestablishing cooperation. At the same time, you don’t want to be too forgiving and risk exploitation. How do you make this trade-off?

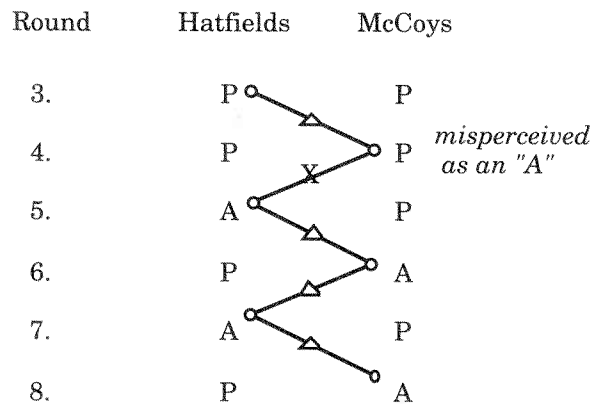
A useful way to evaluate a strategy is to measure how well it performs against itself. If one thinks in terms of evolution, the “fittest strategies” will become dominant in the population. As a result, they will encounter each other often. Unless a strategy performs well against itself, any initial success will eventually become self-defeating.

At first glance, tit-for-tat does very well against itself. Two tit-for-tatters will start off cooperating, and since each is responding in kind, this cooperation seems destined to go on forever. The pair of strategies appears to completely avoid the problem of the prisoners’ dilemma.

But what happens if there is a chance that one side misperceives the other’s move? To find out, we follow two families, the Hatfields and the McCoys, as they use tit-for-tat in their neighborly relations. They begin peacefully (*P*).



Suppose that in round 4 a Hatfield misinterprets a McCoy. Although the McCoys were truly peaceful, the Hatfields mistakenly saw an act of aggression (A).




The single misunderstanding "echoes" back and forth. In round 5, the imagined McCoy aggression becomes real in the Hatfield response. Now the two tit-for-tat clans are trapped in a situation in which they alternate retaliating against the other for the previous retaliation. In round 6, the McCoys punish the Hatfields for their aggression in round 5, which leads the Hatfields to retaliate once more in round 7. And so it goes. Trying to get even for being down one just doesn't work.

The situation continues like this until a second misinterpretation arises. Two developments are possible. The Hatfields could misinterpret peace for aggression or they could

misinterpret aggression as peace.* If aggression is misinterpreted as peace, the feud is ended (at least until the next misperception).

If the second misperception is peace as aggression, both sides will resort to continual retaliation. This is illustrated below in round 9. Here the single helixlike twisting strand of peace is misinterpreted as aggression. Consequently, the Hatfields respond by retaliating in round 11. Until another misperception occurs, both sides continue to punish the other for the other's previous punishments. Although tit-for-tatters can give it, they can't take it.



What can we conclude about the performance of tit-for-tat? When misperceptions are possible, *in the long run tit-for-tat will spend half the time cooperating and half of it defecting.* The reason is that once misperceptions arise, they are just as likely to get compounded as they are to get cleared up. Hence, tit-for-tat will do no better than a strategy based on a coin toss that cooperates and defects with equal probability.

In this discussion, we seem to have left out an important ingredient: the probability that a misperception occurs. In fact, our conclusion does not depend on this probability! No

* Alternatively, these misunderstandings could also arise on the part of the McCoys, and the effect would be the same.

matter how unlikely a misperception is (even if it is one in a trillion), in the long run tit-for-tat will spend half of its time cooperating and half defecting, just as a random strategy does. When the probability of a misperception is small, it will take a lot longer for the trouble to arise. But then once a mistake happens, it will also take a lot longer to clear it up.

The possibility of misperceptions means that you have to be more forgiving, but not forgetting, than simple tit-for-tat. This is true when there is a presumption that the chance of a misperception is small, say five percent. But what strategy would you adopt in a prisoners' dilemma in which there is a fifty percent chance that the other side will misinterpret (reverse) your actions? How forgiving should you be?

Once the probability of misunderstanding reaches fifty percent there is *no* hope for achieving any cooperation in the prisoners' dilemma. You should always defect. Why? Consider two extremes. Imagine that you always cooperate. Your opponent will misinterpret your moves half the time. As a result, he will believe that you have defected half the time and cooperated half the time. What if you always defect? Again, you will be misinterpreted half the time. Now this is to your benefit, as the opponent believes that you spend half your time cooperating.

No matter what strategy you choose, you cannot have any effect on what your partner sees. It is as if your partner flips a coin to determine what he thinks you did. There is simply no connection with reality once the probability of a mistake reaches fifty percent. Since you have no hope of influencing your partner's subsequent choices, you might as well defect. Each period you will gather a higher payoff and it won't hurt you in the future.

The moral is that it pays to be more forgiving up to a point. Once the probability of mistakes gets too high, the possibility of maintaining cooperation in a prisoners' dilemma breaks down. It is just too easy to be taken advantage of. The large

chance of misunderstandings makes it impossible to send clear messages through your actions. Without an ability to communicate through deeds, any hope for cooperation disappears.

A 50 percent chance of a misperception is the worst possible case. If misperceptions were certain to occur, you would interpret every message as its opposite, and there would be no misunderstandings. A stock forecaster whose advice is always dead wrong is as good a predictor as one who is always right. You just have to know how to decode the forecasts.

With this in mind, we look for a way out of the dilemma when there is a chance of misperception, but not too big of a chance.

7. AN ALTERNATIVE TO TIT-FOR-TAT

The basic properties of clarity, niceness, provocability, and forgivingness seem likely to be true of any good rule of behavior for extricating oneself from a prisoners' dilemma. But tit-for-tat is too quick to punish someone who has a history of cooperating. We need to find a strategy that is more discriminating: it should be more forgiving when a defection appears to be an exception, and it should punish when defection appears to be the rule.

You can consider the following guidelines as a step in that direction. (1) Begin cooperating. (2) Continue cooperating. (3) Keep count of how many times the other side appears to have defected while you have cooperated. (4) If this percentage becomes unacceptable, revert to tit-for-tat. Note that unlike before, tit-for-tat is not used as a reward for good behavior; instead, tit-for-tat is the *punishment* if it appears that the other side is trying to take advantage of you.

To determine what is an unacceptable percentage of defections, you need to know both a short-, medium-, and long-term history of the other side's actions. The long run is not enough.

Just because someone has been cooperating for a long time does not mean that he now won't take advantage of you while he runs down his reputation. You also need to know "What have you done for me lately?"

Here is an example of one such strategy. It is nicer, more forgiving, not quite as provokable, and a little more complicated than tit-for-tat. Start cooperating and continue to do so until one of the four tests below fails.

- First impression: A defection on the first move is unacceptable. Revert to tit-for-tat.
- Short term: Two defections in any three turns is unacceptable. Revert to tit-for-tat.
- Medium term: Three defections out of the last twenty periods is unacceptable. Revert to tit-for-tat.
- Long term: Five defections out of the last one hundred periods is unacceptable. Revert to tit-for-tat.

The punishment of tit-for-tat need not last forever. Keep track of how often the other side has violated any of these four tests. On the first violation, return to cooperation after twenty periods of the tit-for-tat "echo" of alternating defections. But put the other side on probation. Reduce the number of defections allowed in the medium- and long-term tests by one. If the other side does not violate the probation for fifty periods, then strike the record clean and return to the original standards. If the other side violates the probation, resort to tit-for-tat forever.

The exact rules for first, short-term, medium-term, and long-term impressions will depend on the probabilities of error or misperception, the importance you place on future gains and current losses, and so on. But this type of strategy is likely to outperform tit-for-tat in the imperfect real world.

The important principle to remember is that when misperceptions are possible, you shouldn't punish every defection you see. You have to make a guess as to whether a misper-

ception has occurred, either by you or by your partner. This extra forgiveness allows others to cheat a little on you. But if they cheat, they use up their goodwill. When the eventual misperceptions arise you will no longer be inclined to let the incident pass. Opportunism on the part of your opponent will be self-defeating.

8. CASE STUDY #4: *CONGRESS V. FEDERAL RESERVE*

The United States Congress and the Federal Reserve often clash over economic policy. To explain why the conflict arises and where it leads, we present Princeton economist Alan Blinder's game-theoretical analysis of the conflict.³ The two institutions have separate and largely independent powers in making economic policy. Fiscal policy (taxation and expenditures) is the responsibility of the Congress, and monetary policy (money supply and interest rates) that of the Federal Reserve. Each can deploy its policies in an expansionary mode or a contractionary mode. Expansionary fiscal policy means high expenditures and low taxes; this reduces unemployment but carries a risk of inflation. Expansionary monetary policy means low interest rates and therefore easier borrowing conditions, but again at the risk of inflation.

The two branches have also developed separate preferences about economic outcomes. Voters like the benefits they get from government spending, as in cheaper mortgages, and dislike paying taxes. Congress responds to this by favoring expansionary policies, unless inflation is imminent and serious. In contrast, the Fed takes a longer viewpoint and thinks inflation the greater problem; therefore it favors contractionary policies.

In 1981–82, Congress no longer regarded inflation as a suf-

ficiently great risk. It felt that the economy could afford an expansionary fiscal policy and wanted the Fed to accommodate by pursuing an expansionary monetary policy. But the Fed under Paul Volcker was afraid that this would just rekindle the fires of inflation. The Fed's first preference was for both fiscal and monetary policies to be contractionary. What seemed best for the Congress was worst for the Fed and vice versa.

The interests of the Congress and the Fed were not entirely opposed. In search of a compromise, the two sides debated the relative merits of combining one expansionary and one contractionary policy. Either way the policies were mixed would have similar effects on general employment and inflation, but differed in other important respects. Fiscal expansion and monetary contraction would lead to a large budget deficit and high interest rates as the need to finance this deficit ran up against the tight money. The high interest rates would hurt such important sectors as autos and construction especially hard. Foreign capital would flow in, attracted by the high U.S. interest rates. The dollar would rise and our international competitiveness would suffer.

Fiscal contraction and monetary expansion would have just the opposite effects — low interest rates and a low dollar — favoring our auto and construction industries, and making our traded goods more competitive. Both Congress and the Fed preferred this second combination of policies to the first.

What would you predict in this situation? How would you judge the outcome? What reforms in the policy-making process would you prescribe?

Case Discussion

This is a prisoners' dilemma. (Otherwise the case wouldn't be in this chapter, would it?) Let the Congress and the Fed rank the four possible policy combinations, 1 being the best and 4 the worst in each case. Then we have this table.

Rankings of Outcomes for (Fed, Congress)

		Congress's Choice	
		High Expenditures	Low Expenditures
Fed's Choice	Easy Money	1 4	2 2
	Tight Money	3 3	4 1

High expenditures is a dominant strategy for the Congress; tight money, for the Fed. When the two think in this way and each selects its preferred strategy, the result is a budget deficit and tight money. This is exactly what happened in the early 1980s. But there is a better outcome for both, namely a budget surplus and looser money.

What prevents them from reaching an outcome both prefer? The answer lies, once again, in the interdependence of decisions. The jointly preferred outcome arises when each chooses its *individually worse* strategy. Congress must restrict spending to achieve a balanced budget. Having done so, how can it be sure that the Fed will not respond with a tight money supply? It knows that the Fed has a temptation to sneak a switch to a tight money supply to achieve its ideal outcome, which would result in the worst possible outcome for the Congress. Congress does not trust the Fed to refrain from this temptation. It is their inability to make credible promises to each other that locks the adversaries into an outcome they could jointly improve upon.

Can we suggest a way out of this dilemma? The two have an ongoing relationship, and cooperation might emerge in the repeated game. However, that only happens if the players put

sufficient weight on future benefits; Congressmen who must run for reelection every two years find it hard to act with such forethought.

Let us try a different avenue. The Federal Reserve is itself a creation of Congress. In most other countries, the government (the Treasury Department) exercises much more control over the central bank. If the same were true in the United States, the Congress could impose an expansionary monetary policy on the Fed and achieve its most preferred outcome. Of course those who share the Fed's concern for inflation would find this regrettable.

This seems a no-win situation: coordination of fiscal and monetary policies is tantamount to a triumph of the short-sighted political objectives of the Congress, but the checks and balances supplied by an independent Federal Reserve lead to a prisoners' dilemma. Perhaps a solution is to let the Fed choose expenditures and taxes, and let the Congress set the money supply?

5

Strategic Moves

"We must organize a merciless fight. The enemy must not lay hands on a single loaf of bread, on a single liter of fuel. Collective farmers must drive their livestock away and remove their grain. What cannot be removed must be destroyed. Bridges and roads must be dynamited. Forests and depots must be burned down. Intolerable conditions must be created for the enemy." — Joseph Stalin, proclaiming the Soviets' "scorched earth" defense against the Nazis, July 3, 1941.

Today Stalin's campaign lives on in the battlefields of corporate control. When Western Pacific attempted to "annex" the publishing company Houghton Mifflin, the publishing house responded by threatening to empty its stable of authors. John Kenneth Galbraith, Archibald MacLeish, Arthur Schlesinger Jr., and many profitable textbook authors threatened to find new publishers if Houghton Mifflin were acquired. "When Western Pacific Chairman Howard (Mickey) Newman got the first few letters from authors, he thought it was a big laugh, and called it a 'put-up job.' When he began getting more letters, he began to realize, 'I'm going to buy this company and I ain't going to have nothing.'"¹ Western Pacific withdrew its bid, and Houghton Mifflin remained independent.

This strategy doesn't always work. When Rupert Murdoch was interested in acquiring *New York* magazine, the incumbent management attempted to fight him off. Many of the magazine's best-known writers threatened to quit if Murdoch attained control. Murdoch was not deterred. He acquired *New*