The Value of Cooperation in Business, Daily Life, and National Security from a Game Theory Perspective

by Dietrich Fischer <fischer@epu.ac.at>
Academic Director, European University Center for Peace Studies, Stadtschlaining, Austria
September 2006

What is Game Theory?

Even today, most economists solve problems by "optimization", which seeks the highest possible value of a goal function under given constraints. But as the Austrian economist Oskar Morgenstern (1902-1977), co-founder of game theory, pointed out, this implicitly assumes that there is one single decision-maker in control of all variables, so to speak a global dictator. This is unrealistic. In the real world, usually at least two--and often many more--independent decision-makers, with partly common and partly competing interests, determine the outcome of a problem. What we achieve depends normally not only on what we decide, but also on the decisions made by others.

In 1944, Morgenstern and the Hungarian mathematician John von Neumann (1903-1957), who were both teaching at Princeton University at that time, having escaped from Hitler, published "The Theory of Games and Economic Behavior". It is a difficult to read, highly mathematical 1200-page book and not recommended as an introduction. Since then, tens of thousands of books and articles on various aspects of game theory have been published.

In the most general and abstract sense, a game with n decision-makers (called "players") is a set of n functions

f1 (s1, s2, s3, ..., sn),
f2 (s1, s2, s3, ..., sn),
f3 (s1, s2, s3, ..., sn),
...
fn (s1, s2, s3, ..., sn),

where fi is the outcome or "payoff" to the i'th player, as a function of the strategy si chosen by i itself and the strategies chosen by all the other players. We will see a number of examples of such games. For simplicity, this introduction will be limited mostly to two-person games.

In a game of chess, which has two players, the three possible values of f1 are 1 (win), 0 (lose) or 1/2 (undecided, "remis"). f2 = 1-f1, that is, the sum of the outcomes for the two players is a constant, always equal to 1. This is an example of a constant sum game (usually called "zero-sum game"), where one player's gain is the other players loss. There exist also games in which the sum of the outcomes is variable (usually called "non-zero-sum games").

In the example of a game of chess, the strategies s1 and s2 would be descriptions of what move a player would chose as a function of all of his own and the opponent's previous moves. This is a finite set, but it is so huge that no human being could define it. When we play chess, we make up our mind from move to move, and cannot spell out in advance how we would react to every possible sequence of moves. We will consider games that are much simpler than chess.

The number of possible games of chess is also finite, because if the same board position is repeated for a third time, the game is declared undecided, and the number of possible board positions is finite. But this number is huge beyond imagination.

There is a legend that the inventor of the game of chess was called before his ruler, a wealthy maharajah, who was so delighted with the game that he offered the inventor anything in his realm that he might wish. The inventor said, "I wish one grain of rice on the first square of the chessboard, two grains on the second square, four grains on the third, eight grains on the fourth, and so on, until the last square." The maharajah exclaimed, "Why such a modest wish?! Don't you wish my daughter's hand, or one of my palaces, or one of the provinces I rule?" "No thank you," the inventor said, "all I want is the amount of rice that I explained to you." The maharajah had his treasurer calculate how much rice this would amount to. Soon he began to realize that even if the maharajah sold his entire realm with all his palaces and vast amounts of land, he would never be able to buy so much rice. This only reinforced his admiration for the inventor, and he rewarded him handsomely, but with far less than he had asked for.

An Example of a Zero-Sum Game: Paper, Scissors, Stone

Let us consider a simple game that children sometimes play, called "paper, scissors, stone". There are two players, and one of them wins with the other losing, or the game is undecided, just like in chess. But different from chess, each player has only three possible strategies to choose from: paper (showing a flat hand), scissors (showing two fingers spread apart) or stone (showing a fist). Both children must display their choice simultaneously, without knowing in advance what the other player will choose, by counting to three and making their move on three. If one child shows scissors and the other paper, scissors wins, because scissors can cut paper. If one shows stone and the other scissors, stone wins because a stone can sharpen the scissors. And if one shows paper and the other stone, paper wins because paper can wrap a stone. If both show the same symbol, the outcome is undecided. We can describe this game in the form of a "payoff matrix", where player 1 chooses a row and player 2 chooses a column in a 3x3 table, or "matrix". Such a table has 9 cells, and the figure at the bottom left of each cell indicates the payoff to player 1, whereas the figure at the top right of each cell indicates the payoff to player 2. This payoff matrix looks as follows:

player 2 player 1	paper	scissors	stone
paper	1/2 1/2	1 0	0 1
scissors	0 1	1/2 1/2	1 0
stone	1 0	0 1	1/2 1/2

If player 2 knew the strategy chosen by player 1 in advance, it would be easy to see how he can win, by choosing the one strategy that yields 1 to him and 0 to player 1. But in reality, both must play at the same time, otherwise the player who is second has an advantage. If this game is played many times, and one player always makes the same choice, the other will soon figure out how to win. Choosing one particular symbol is called a "pure strategy". To prevent that, a player can vary her or his strategy randomly. The best solution in this case turns out to be to choose each of the three possibilities with a probability of 1/3. This is an example of a "mixed strategy". This will guarantee her an average payoff of 1/2, because regardless of what player 2 chooses, player 1 will win 1/3 of the time, lose 1/3 of the time, and get a draw (undecided) 1/3 of the time. By symmetry, this is also the best strategy for player 2, with the same average outcome, a draw, valued at 1/2.

Player 1 cannot guarantee for herself any higher average outcome than 1/2. If she were to choose, say, paper with probability 1/2 and the other two choices with probability 1/4 each, the best that player 2 could do is to choose always scissors, which wins against paper, with an average payoff to player 2 of 1/2*(1) + 1/4*(0) + 1/4*(1/2) = 5/8, slightly more than 1/2. In this game, choosing each symbol with equal probability is the best mixed strategy for player 1, and also for player 2. It maximizes the minimum expected payoff, regardless of what player 2 chooses. In this case, the same strategy is best for player 2, because it minimizes the maximum payoff of player 1 (and therefore maximizes the minimum payoff to player 2), regardless of what player 1 chooses.

Already in 1928, John von Neumann provided a difficult proof for the famous "minimax theorem" for two-person constant sum games: There exists a mixed strategy for player 1 which maximizes the minimum payoff that she can expect regardless of the strategy chosen by player 2. Similarly, there exists a mixed strategy for player 2 which minimizes the maximum payoff that player 1 can expect, regardless of the strategy chosen by player 1. The two values are equal, in other words, the "minimax" is equal to the "maximin".

Most games of sports, such as football, tennis, chess etc. are of the constant-sum variety, with pure competition and no element of cooperation, but there are also some exceptions, for example frisbee, where the goal is to throw the frisbee back and forth as many time as possible without dropping it. This requires cooperation, and either both players do well together, or both play poorly. When frisbee was already popular in the United States but hardly known in the UK, an American father and son played frisbee in Hide Park in London. A British man observed them for about 15 minutes, and finally approached the father, asking, "Excuse me, please, but can you explain to me who is winning?"

Pure zero-sum games are rare in real life, usually there is some common interest and room for mutual cooperation that can help all parties achieve a better outcome, and of course also the possibility of mutual ruin, like in war.

Variable-Sum Games

Next we consider some two-person games for which the sum of outcomes is variable. Best known among them is the "prisoner's dilemma" game, but there are many more. First we will briefly examine two other variable-sum games, the survival game and "chicken". How many different games are there? Let us restrict ourselves to two-person games where each player can choose only among two strategies, and we are interested only in the ranking of outcomes, not the absolute payoffs. Assuming the 4 possible outcomes for each player are different, and taking into account that the players and their strategies are interchangeable, there are 78 such games. Prisoner's dilemma and chicken are only two of the 78. If we allow some of the outcomes to be identical (as for example in the survival game shown below), there are even more possibilities. Anyone fascinated with counting rules can calculate these numbers, but doing so here would divert us too far from the main topic.

The Survival Game

An illustration of a survival game is two people trapped in a burning house with a locked door. Neither is strong enough to break the door open alone, only if they cooperate can they escape and survive. This game has the following payoff matrix, where 1 = survival and 0 = death:

player 2 player 1	cooperate	do not cooperate
cooperate	1 1	0 0
do not cooperate	0 0	0 0

This is probably close to what we face in the world today. If humanity is to survive, we must cooperate to solve global problems that no individual, not even a superpower, can solve alone. Among such problems are nuclear disarmament, preventing global warming, preventing the destruction of the ozone layer that protects us from cancer causing ultraviolet radiation, and many more.

The first advanced civilizations emerged about 6,000 years ago in the Nile, Euphrates and Yellow River valleys when farmers faced problems that they could not solve alone. To prevent recurrent floods and droughts, it was necessary to build dams to control the flow of those rivers, requiring the organized cooperation of thousands of individuals. This gave rise to the first states, the development of written language, the codification of laws, and a flourishing of science and the arts. Today we face some problems that not even a superpower can solve by itself, which require global cooperation. At the same time, improvements in communication and transportation have made global cooperation feasible. But it is still doubtful if humanity will learn to cooperate, or destroy itself.

Many governments are still reluctant to join global agreements to deal with global problems, out of fear that they would lose part of their national sovereignty. But that fear is mistaken. No country today, for example, has sovereign control over the ozone layer. By joining a global agreement that allocates and enforces emission quotas, we do not give up control over our destiny. On the contrary, we gain added control that we do not now posses and could never achieve at the national level.

The Game of Chicken

The game of chicken is named after a game sometimes played by members of rival motor cycle gangs. Two motor cyclists drive towards each other on a straight road at high speed. If none of them moves out of the way, they collide and both die. If both move out of the way, it is a draw. But if one swerves out of the way at the last moment, and the other drives straight through, the one who did not yield is the winner, the hero, and the other is called a coward, or a "chicken". This game has the following payoff matrix, where 4, the highest outcome, denotes victory, 3 is a draw where both moved out of the way, 2 is being the chicken, and 1, the lowest outcome, is death if neither player yields.

player 2 player 1	move out of the way	drive straight
move out of the way	3 3	4 2
drive straight	2 4	1 1

The game of chicken is often pursued by heads of state who take a hard line and wait for the other to blink first. If neither blinks, it can lead to mutual destruction. The members of motor cycle gangs at least had the courage to endanger only themselves, nobody else. Heads of state put at risk whole nations by pursuing this arrogant, egotistical game. As Kenneth Boulding once said, "If you drive a schoolbus, you don't play chicken."

The Prisoner's Dilemma Game

The prisoner's dilemma can explain apparently irrational human behavior in economics, politics and daily life, such as ruinous competition, environmental pollution, arms races etc.

Before looking at the story from which this game received its name, I want to give some other typical examples of prisoner's dilemmas. Basically, it is a game where one player can gain an advantage by pursing a certain strategy (called to "defect"), as long as nobody else chooses the same strategy. If all defect, all are worse off. In a battle to defend the homeland against an aggressor, a defector may save his life as long as the others keep defending it. But if all defect, the land may be conquered by the aggressor, and the inhabitants and their families may be killed.

A harmless illustration of a multi-person prisoner's dilemma is the following: a group of people is watching a parade. If one spectator stands on tip-toe, she can see better than the others. But if all stand on tip-toe, nobody can see any better than if they would stand on their feet, they only get all tired sooner.

Competition in Extracting Oil

When petroleum was first discovered in Texas in the early 19th century, there were no rules. If someone discovered an oil well, rival oil companies quickly rushed to the scene and drilled a hole nearby to extract some of that oil before it was all taken. It even occurred that rival oil companies bombarded each other's drilling towers to get at the oil first. Soon they realized that they can never make a profit that way. Today, they appreciate that the US federal government auctions tracts of land with suspected oil reserves and grants exclusive drilling rights to the highest bidder for a parcel of territory. The oil companies pay something, but in return they enjoy the security and peace of mind that they can explore for oil without fear that someone else will come and take it away if they make a discovery.

The lawless situation that existed at the beginning with oil extraction still applies to the reserves of manganese, iron and cobalt on the deep seabed outside of any country's jurisdiction. Whoever will first develop the technology to mine these resources can reap the bulk. A similar system of orderly auctions is needed also at the global level, to prevent future wars over such resources. In addition, it would help raise some badly needed funds to help the United Nations address global problems of development, peacekeeping and protecting the global environment.

Let us now consider the situation where two oil companies, A and B, have drilled one hole each into the same underground oil well and extract oil from it. Oil should not be extracted too rapidly, otherwise the oil, a viscous fluid resting in sandy underground formations, does not have enough time to flow to the well, and gets stuck in the sand. Consider two strategies for each company, to pump oil slowly, or fast. If both pump slowly, assume that they can each extract 8 million barrels of oil from the well. If both pump fast, they can each get only 6 million barrels of oil out of the well, with the rest being wasted and lost forever. (It may be debatable whether we should extract oil now, or save it for future generations, but we certainly should not waste it.) So it is clearly in their joint interest to pump slowly, getting together a total of 16 million barrels instead of only 12 million barrels if they pump fast.

If one company pumps fast and the other slowly, they get an intermediate amount of oil, 14 million barrels together. But the one company that pumps fast gets a bigger share of the oil, let us say 10 million barrels, whereas the one that pumps slowly gets only 4 million barrels. This yields the following payoff matrix:

company B company A	pump slowly	pump fast
pump slowly	8 8	10 4
pump fast	4 10	6 6

Consider the decision-making process in company A. They do not know what company B may do, they do not talk with them, because they are rivals. If company B pumps fast, company A can also pump fast and get 6 million barrels, or it can pump slowly and get only 4 million barrels. In this case, it is better for B to pump fast. But what if company B pumps slowly? Then company A can also pump slowly and get 8 million barrels, or it can pump fast and get 10 million barrels. Again, it is better for A to pump fast. So, not knowing what company B will do, company A always prefers to pump fast, because that way it gets more oil than by pumping slowly.

By exactly the same logic, company B will also prefer to pump fast. So they end up with only 6 million barrels each, when in fact they could have gotten 8 million barrels if they cooperated. This is a typical illustration of a prisoner's dilemma.

Nash Equilibrium and Pareto Optimum

The suboptimal equilibrium where both pump fast and get 6 million barrels each is called a "Nash equilibrium", named after the mathematician John Nash, born 1928, who shared the economics Nobel prize in 1994. A Nash equilibrium is characterized by the fact that neither player can ALONE improve his or her situation, only if they move together can they jointly improve their situation. The optimal outcome, where each company gets 8 million barrels of oil, is not a Nash equilibrium, because one of the companies could unilaterally improve its position and get 10 million barrels of oil by pumping fast.

John Nash is a remarkable personality. He was a brilliant young mathematician, making creative contributions to game theory and other fields, but then got schizophrenia. By sheer willpower and intellectual brilliance, he was able to control his disease without medication. He is featured in the best-selling novel "A Beautiful Mind" by Sylvia Nasar, which was made into an Oscar winning film by the same name.

By contrast to the Nash equilibrium, a Pareto optimum is a situation where no player can improve her situation without making another player worse off. In the above payoff matrix, the situation where both companies pump slowly and get the maximum total amount of oil, 8 million barrels each, is a Pareto optimum. One player could get more, 10 million barrels, but only at the expense of another player's loss, who then receives only 4 million barrels.

A Pareto optimum has a certain similarity to a transcendent solution of a conflict, where all parties achieve together as much as is feasible.

In the early 20th century, there were two main schools of economics, the mathematical school and the historical school. Vilfredo Pareto (1848-1923) was the key representative of the mathematical school, which maintained that there are certain universally valid economic laws which can be formulated as mathematical equations. By contrast, the historical school, whose main representative Muller maintained that human behavior is unpredictable and cannot be described by mathematical equations or general laws like laws of nature. All that economists can do is to record what actually happened, like historians. At an international congress, the two debated their positions, of course without reaching any agreement, and split the audience into two camps. When the session was over at lunch time, Pareto approached Muller and proposed, "Let's have lunch together to continue our discussion, let's go to a restaurant where we get a free lunch." Muller protested, "There is no restaurant that offers a free lunch!" Pareto triumphed and said, "So after all, you do believe in some laws governing human behavior."

The Nobel-Prize winning economist Milton Friedman (born 1912) from the University of Chicago, one of the main advocates of free market economics, wrote a book "Free to Choose", which formed the basis of a public television series by the same name, in which he basically argues, "If everybody is free to choose, everybody is as well of as he can possibly be, by definition, for otherwise he would choose to do something else."

Jacob T. Schwartz, a mathematician at New York University and my thesis adviser, who has written about mathematical economics, among many other fields, pointed out that this statement confuses the Nash equilibrium with the Pareto optimum. Someone may well be free to choose what to do, but may not be able to improve his or her situation unless others cooperate. Schwartz suggested that confusing these two equilibria was as logical as if someone were to say, "If every husband has a woman as his wife, there must exist some woman who is the wife of every husband."

Some Taxes Can Improve General Welfare

Let us next consider the role of government. Friedman maintained that the best role a government can play in the economy is no role at all, to get off people's back and let them do whatever they wish. If a government practices taxation and redistribution, it takes resources away from someone to give them to someone else, keeping something for administrative overhead in between, therefore reducing the total welfare. Let us examine if that is always true.

In the above example of two oil companies, assume that the government charges a tax, the equivalent of 1 million barrels of oil for pumping slowly, and 4 million barrels of oil for pumping fast. That modifies the above payoff matrix as follows:

company B company A	pump slowly	pump fast
pump slowly	7 7	6 3
pump fast	3 6	2 2

This changes the decision-making process of the two companies. If company B pumps fast, company A can also pump fast and get 2 million barrels, or it can pump slowly and get 3 million barrels after taxes. In this case, it is better for A to pump slowly. What if company B pumps slowly? Then company A can also pump slowly and get 7 million barrels, or it can pump fast and get only 6 million barrels after taxes. Again, it is better for A to pump slowly. So company A now always prefers to pump slowly, regardless of what B chooses to do, because that way A gets more oil than by pumping fast.

The result is that each of the companies gets one million barrels of oil more, and in addition, the government collects 2 million barrels of oil in taxes, which it could use to subsidize heating oil for the poor in winter. Everyone is better off with taxes than without in this case, contrary to Friedman's assertion.

It is remarkable that contrary to a widespread popular belief, pollution taxes, which charge to polluters the full cost of the damage they cause to society, actually reduce overall taxes instead of increasing them. This is easily seen through the following thought experiment: Imagine what would happen if we did not pay for gasoline at the pump. People would use a lot more gasoline, and someone has got to pay the national annual gasoline bill anyway. That would mean that everybody's taxes would have to be increased, by much more than we now pay for gasoline, because more would be consumed. For the same reason, if do not charge those who pollute our air, water and soil, we end up paying much more, if not in terms of money, then certainly with poor health.

Why the Name "Prisoner's Dilemma"

The first illustration of this dilemma was offered in 1950 by the mathematician Albert W. Tucker (1905-1995) from Princeton University in addressing an audience of psychologists at Stanford University, where he was a visiting professor at the time. Tucker used the following story: Imagine two gangsters have robbed a bank and hidden their loot in a secret place. They are caught by the police and suspected of the robbery, but there is no proof because no witness observed the robbery. They are held in separate cells and questioned by the prosecutor. If they both confess, they each spend 5 years in prison for bank robbery. If they both deny the accusation, there is no proof, and they can only be convicted of a lesser charge, the possession of a gun without a permit, which carries a sentence of 1 year in prison for each. So if those were the only two options, it is clearly in the gangsters' interest to deny the charge.

But now comes the dilemma. If one confesses and the other denies the crime, the one who confesses is let free as a reward for assisting the prosecution, whereas the other gets 5 years in prison for bank robbery, and an additional 5 years for lying about it, seeking to cover up his crime, or a total of 10 years in prison. This is reflected in the following payoff matrix, where a minus sign in front of the years in prison indicates that less is better. 1 year in prison is preferable to 5, whereas in the earlier example, 8 million barrels of oil extracted was preferable to 6 million.

player B player A	deny	confess
deny	-1 -1	0 -10
confess	-10 0	-5 -5

Prisoner A has no idea what prisoner B said, or may say, so there are two possibilities. If prisoner B confessed the bank robbery, prisoner A better confess it, too, so he spends only 5 years in prison instead of 10 if he denies it. And if prisoner B denies the robbery, prisoner A again is better off by confessing, because then he is let go scot free. If A also denied the crime, he would spend one year in prison for illegally possessing a gun.

I find this didactically a spectacularly poor illustration, because it seems to teach that lying, denying a crime committed, is the recommended course of action, which would promote the "common good" of the two criminals, who then spend less time in prison than they deserve. In reality, cooperation between the two players is usually also in the best interest of society as a whole, as in the previous example of oil extraction.

The story would be better if the police had caught two innocent people, who did not commit the bank robbery. Because the way the penalties and rewards are set up, they would still be tempted to make a false confession, to spend less time in prison than by telling the truth. In this case, they should indeed cooperate by telling both the truth, and getting the lesser sentence, which they now deserve.

A law professor in Chicago assigned his students to reexamine the cases of people waiting for execution. In several cases they were able to prove that they were falsely convicted, on the basis of the testimony by another criminal who was bribed into making a false confession with the promise of a lesser sentence, or in one case a $100 cash payment. Such testimony is unreliable.

The "Generic" Prisoner's Dilemma

The decisions of the two players are influenced only by the rank order of the different outcomes, not the absolute magnitude. For this reason, from now on, we will use the following "generic" payoff matrix for the prisoner's dilemma, where 4 is the best outcome, 3 the second best, 2 the next to worst and 1 the worst. The payoff matrix then looks as follows:

player B player A	cooperate	defect
cooperate	3 3	4 1
defect	1 4	2 2

A Sales Transaction as a Prisoner's Dilemma

Every sale can be represented by a prisoner's dilemma. Take for example a publisher who sells journal subscriptions. From the publisher's point of view, it is preferable to sell a subscription and to send the journal than not to sell a subscription and not to send the journal. The printing costs of an additional copy of the journal must be less than the subscription price, otherwise the publisher would go bankrupt. (A son who inherited a business from his deceased father was selling below production costs. His mother expressed concern. He said, "Don't worry, mother, I will make it up in volume.") Even better for the publisher would be to collect the subscription fee without having to send the journal. The absolutely worst outcome for the publisher is to send the journal without getting paid.

For the subscriber, the rank order of the 4 outcomes is different. The best is to receive the journal without having to pay. Next (assuming she is interested in a subscription) is to receive the journal and pay for it. Somewhat worse is not receiving the journal and not paying for it. The worst possibility is to pay for the subscription and not to receive the journal. If we rank these 4 outcomes from the point of view of the two parties involved, the publisher and the subscriber, with 4 being the best and 1 the worst outcome, we get the following payoff matrix:

publisher subscriber	send the journal	do not send the journal
pay for the subscription	3 3	4 1
do not pay for the subscription	1 4	2 2

This is isomorphic to a generic prisoner's dilemma.

Arms Races as Prisoner's Dilemmas

The most dangerous arms race in history so far was the nuclear arms race between the United states and the Soviet Union. It lasted 40 years, from 1945 to 1985, until Gorbachev unilaterally ended it, and was later joined by the United States. The United States and Russia still possess thousands of nuclear warheads, but there is no race between them anymore. The United States under the Bush administration is again developing new nuclear weapons (so-called "bunker-busters" that penetrate deep into the earth before exploding), but Russia no longer competes.

For simplicity, we assume that the United States (US) and Soviet Union (SU) each had only two choices, to arm (build more nuclear weapons) or to disarm (eliminate its nuclear weapons, or at least significantly reduce their number). This yields the following two by two table:

SU US	disarm	arm
disarm	(a) both countries are safe, without the burden of heavy military spending	(b) the SU can dominate the US
arm	(c) the US can dominate the SU	(d) both countries are highly insecure, with a heavy burden of military spending paralyzing the civilian economy

If we assume that both countries make the same choice, then clearly option (a) with both disarming is far preferable to option (d) where both arm, where each country could be wiped off the face of the earth within half an hour if its opponent launched its nuclear missiles, and billions of dollars are diverted from productive investments in the civilian economy towards military purposes. But even worse from Stalin's point of view than high insecurity and an economy making sacrifices to support a nuclear arms industry is the possibility of being dominated by the US. Indeed, when the US had a monopoly on nuclear weapons in 1946, and there was a civil war in Iran with Soviet and British troops involved, supporting different sides, President Truman warned Stalin to withdraw his troops within 48 hours, or the United States would use nuclear weapons. Stalin withdrew his troops, but was even more eager to develop nuclear weapons also, so as never to be in such a weak and humiliating situation again.

The ideal outcome from Stalin's viewpoint was undoubtedly if the Soviet Union had had a monopoly on nuclear weapons, and could install "friendly" governments throughout the third world, without the United States being able to stop it, because it would have exposed itself to the risk of a nuclear attack that it had no means to deter.

From the position of a US President, being dominated by the Soviet Union must be unacceptable, the worst of the 4 outcomes.

Normal people would be satisfied with position (d) where both countries are disarmed, secure and without military spending being a burden on the civilian economy. But those who aspire for top leadership positions, and get them, often have a slightly more than average desire to be in charge, to tell others what to do, to dominate (with some great exceptions, like Mandela, Carter, Gorbachev and several others). Therefore, for a typical President, possessing a monopoly on nuclear weapons, which would make it possible to support "friendly" governments without any opposition from an adversary may look more attractive than simply to live in peace. This then gives the following ranking of the four outcomes (a) to (d) in the above table, where the preferred outcome is listed first, and the least preferred listed last:

for the Soviet Union: (b), (a), (d), (c).
for the United States: (c), (a), (d), (b).

If we assign the payoff values 4 (for the preferred), 3, 2, down to 1 for the least preferred outcome, the reader can verify that we get the generic payoff matrix of the prisoner's dilemma. Therefore the prisoner's dilemma offers an explanation for the apparently insane policy of building ever more nuclear weapons in a vain hope to achieve superiority and security, when in fact it brought only mutual insecurity.

During The early 1980s, a student asked me, "Does Reagan understand this?" I had to say, "I am afraid not." But probably Gorbachev understood it, and thus ended the nuclear arms race through unilateral restraint. He also offered a plan for complete worldwide nuclear disarmament, but unfortunately, the United States rejected it.

Through negotiations and mutual cooperation, both sides could achieve the position of mutual disarmament which brings mutual security, and frees both economies from high taxation necessary to pay for military spending.

Some people falsely argue that military spending stimulates economic activity because it creates jobs. That is like arguing, "If I wash my car, it is good for my vegetable garden, because some of the dirty water flows towards the vegetables." Why not pour all the clean water directly onto the vegetables, it will do much more good for the garden. Similarly, Caspar Weinberger ("Cap the knife"), Reagan's Secretary of defense, argued that every billion dollars spend by the Pentagon creates 30,000 jobs. He did not say that every billion spent for education creates 70,000 jobs. So in fact spending it for the military instead of education destroys 40,000 jobs, and on top of it leaves a generation poorly educated.

If some people are hired to clean up toxic waste being generated by a factory, and a new manufacturing process can eliminate the source of toxic pollution, we do not lament, "What will those poor people who were cleaning up the toxic waste now do?" There are plenty of jobs that need to be done, restoring a clean environment, providing assistance to sick and elderly people, caring for children, etc., which we currently cannot afford, but could afford if we can save elsewhere. The same applies to military spending, we should be glad if it can be reduced, because the same resources can be spent for more productive uses.

The Prisoner's Dilemma as Explanation for Environmental Pollution

Consider a concrete example, the catalytic converter, which cleans up pollutants from car motors before the exhaust is released into the environment. When it was first invented, it was an optional attachment to cars, which a buyer could add for, say, an additional $500 to the price of the car. To simplify, we consider only two decision-makers: ourselves, and everybody else. We also list only the outcome for "self". Then we have the following table of possible outcomes:

all others / self	buy a catalytic converter	do not buy it
buy a catalytic converter	(a) breathe clean and healthy air, pay $500	(b) breathe polluted, carcinogenic air, pay $500
do not buy it	(c) breathe clean and healthy air, pay nothing	(d) breathe polluted, carcinogenic air, pay nothing

I assume that the reader is a conscientious environmentalist who would buy a catalytic converter regardless of what others do, because it is the right thing to do. The "homo economicus", the kind of decision-maker that neoclassical economists assume to reflect human behavior, who thinks only of him- or herself and nobody else, is a caricature of human nature. People do care about each other to a certain extent. Kenneth Boulding (1910-1993) estimated that about 30 percent of economic transactions are not sales, but pure gifts, such as people doing volunteer work, making contributions to charitable organizations, parents paying their children's education without expecting something in return, gifts given on birthdays and religious holidays, inviting friends to dinner, etc. People are not total egoists as most economists assume in their models. There exists also a degree of altruism.

But unfortunately, altruism has limits. There are indeed many people who enjoy being a "free rider" if they can get away with it. Since a driver does not inhale the pollution emanating from the exhaust pipe behind his car, only passersby and drivers behind him, he may not be concerned about polluting the environment for others.

A selfish individual, but someone whose own health is worth to her $500, would rank the 4 outcomes above from best to worst as follows: (c), (a), (b), (d).

If we assume that everybody thinks the same way, we have again a prisoner's dilemma, and nobody buys a catalytic converter voluntarily. That is why we have so much pollution if there is no penalty for it, if people are expected to keep the environment clean on a voluntary basis.

The Importance of Public Choice

There is, however, a solution: democratic decision-making, or "public choice". Assume instead of deciding individually, people have the right to vote on the proposition, "Are you in favor of legislation making catalytic converters mandatory for all cars?" Now the only choices left are (a) and (d). Even people who would not voluntarily buy a catalytic converter may vote in favor of this proposition, because now they know that if they contribute their share of $500, everybody else has got to do the same, and they can really enjoy clean and healthy air. Such democratic decision-making may be one of the main reasons for the existence of government, by providing public services that private enterprise will not provide, including a clean environment.

Repeated Prisoner's Dilemmas

If two people play a prisoner's dilemma only once, they are tempted both to defect, and end up in the suboptimal Nash equilibrium. But if the same opponents (or partners) play a prisoner's dilemma game repeatedly, they can induce each other to cooperate for mutual benefit, rewarding cooperation by cooperating also, and punishing defection by defecting also, in retaliation.

Robert Axelrod at Michigan State University conducted an experiment about this and wrote up his findings in "The Evolution of Cooperation" (New York: Basic Books, 1984). He invited all authors of articles about the prisoner's dilemma game to describe how they would play a series of several hundred prisoner's dilemmas against a hypothetical opponent. Their goal should be to achieve the highest possible score in a tournament of everyone against everyone else (and against his or her own strategy, something impossible in a chess tournament, unless the participants have multiple personalities. A computer can very easily be told to forget what it just played). The scores were given by a generic prisoner's dilemma payoff matrix, with 4 the highest and 1 the lowest outcome. To preclude any ambiguities, those who participated had to describe their choice, cooperate or defect, as a function of all the previous moves, in the form of a computer program.

13 authors responded. The shortest and simplest program was submitted by Anatol Rapoport (born 1911), who called his strategy "TIT-FOR-TAT". He cooperates the first time, and then simply repeats what the other side did on the previous move. As long as the other side cooperates, he cooperates also on the next move. If the other side defects, he defects on the next move, but as soon as she cooperates again, he cooperates on the next move.

Others submitted more sophisticated and tricky programs. For example a Swiss mathematician Joss basically played like tit-for-tat, but randomly defected with a 10% probability, without having been provoked. He won against tit-for-tat. But he did not win the overall tournament, tit-for tat won! How is this possible? When Joss played against itself and other similarly mean-spirited strategies, very soon both defected throughout the remainder of the game, retaliating against the first random defection, and then retaliating against this retaliation, as it typically occurs in the real world in a vicious cycle of violence and mutual retaliation. This gave both strategies a low score. However when tit-for-tat played against itself and similarly nice strategies, they kept cooperating throughout the remainder of the game, and both achieved a high score, doing well by helping others to do well also.

Axelrod and his colleagues discovered that another, more forgiving strategy, would have gotten an even higher score in the tournament than tit-for-tat, but nobody submitted it. They called it "TIT-FOR-TWO-TAT". It cooperates the first two times and then defects only if the other side defects twice in a row. It did better in playing against Joss and other nasty strategies, by cooperating more often.

Axelrod wrote up the results of this first tournament, including the fact that tit-for-two-tat would have won if anyone had submitted it. They published these findings in Psychology Today, computer hobby magazines, and many other popular journals, inviting everyone who wished to participate in a second computer tournament of the prisoner's dilemma to submit a computer program describing their strategy. 73 participated in the second round. Someone did submit tit-for-two-tat in the hope of winning. Rapoport submitted again tit-for-tat, unfazed, and won again! How did that occur? Someone expected that there would be a strategy tit-for-two-tat and said to herself, "I know how to win against that one. Simply alternate between cooperation and defection, then tit-for-two-tat will never retaliate, and every second time I get the maximum score of 4 (I defect, the other cooperates), with the remainder getting 3 on every other move (both cooperating, see the payoff matrix of the generic prisoner's dilemma). This strategy, "ALTERNATE", did not win the tournament, but it had the effect that tit-for-two-tat did not win either.

Properties of Winning Strategies

Axelrod and his colleagues carefully analyzed all the strategies of the first and second tournament, to see what properties made them do well or poorly, and detected 4 properties of high-scoring strategies:

1. NICE, meaning that a strategy is never the first to defect. Of the top 15 among 73 strategies submitted for the second tournament, 14 were nice. Among the bottom 15, 14 were not nice. So it pays to be nice.

2. PROVOCABLE, meaning that if the other side defects, I defect also in retaliation, always, immediately on the next move.

3. FORGIVING, meaning that if the other side cooperates, I reward it by cooperating also, always, immediately on the next move.

4. CLEAR, meaning simple and transparent, making clear to the other what I expect, and what I will do in response to her move.

Tit-for-tat has all four of these properties.

Joss's strategy is not nice, because it may randomly defect, for the first time, without having been provoked. It is provocable (if the other side defects, it is guaranteed to defect on the next move) but it is not forgiving, because it may randomly defect even if the other side cooperated on the previous move. It is also not very clear, because it has an unpredictable random element.

Tit-for-two-tat is nice, not provocable (it does not retaliate immediately against a single defection), it is forgiving (very forgiving), and quite clear.

Alternate is not nice (it defects on every second move without having been provoked by a previous defection), it is not provocable or forgiving (it does not even examine what the other side did before choosing its move, it could be called "autistic"), but it is clear and simple.

A good article on this tournament, with additional strategies and their properties, is Douglas Hofstadter's, "Computer tournament of the Prisoner's Dilemma suggests how cooperation evolves," *Scientific American*, Vol. 248, No. 5, May 1983, pages 16-26.

These four principles emerged out of a computer tournament of the prisoner's dilemma, they were not written on stone tablets on a mountain top, nor the result of deep philosophical reflection, but they make sense nevertheless, in business, daily life and national security. The following illustrates some applications of these four principles.

Business

Take the above example of a journal publisher and a subscriber.

Being "nice" for the publisher means never to be the first to defect, that is, to keep sending the journal as long as the subscriber pays. That is good business practice. If you do not live up to contracts, the word will spread, your reputation suffers, and you will not get many new subscribers and lose some of those you already have.

Being "provocable" for the publisher means that if the subscriber fails to pay, the publisher stops sending the journal, also good business practice. If the word gets around that you do not have to pay, you receive the journal anyway, the publisher will lose income and may go bankrupt.

Being "forgiving" for the publisher means that if a subscriber has missed a payment, but later apologizes and sends a late check, explaining that he was on vacation, the publisher should accept him back. If she writes, "No way, you once missed a payment and that's enough for me. I will never accept you back as a subscriber, you unreliable misfit", she may gradually lose more and more subscribers. Being too harsh and unforgiving is also a poor business practice.

Being "clear" for the publisher means that she inform the subscribers what she expects from them, and what the consequences of their decisions will be. If she expects payment within 30 days, it is a good idea to inform customers, and not to assume that they will guess it by themselves. If the penalty for late payment is 2% surcharge for ever month delay, that should be said. She should also warn after how much delay in payment she will stop sending the journal. All of this is very sensible.

For the subscriber, being nice means to keep paying for the subscription as long as the journal is being delivered. Nonpayment has the likely consequence that the journal will no longer arrive, and as a subscriber, he presumably prefers to pay and receive the journal than not to pay and not to receive it.

Being provocable means to stop paying the subscription if the journal is not delivered, otherwise he pays for nothing.

Being forgiving means that if a journal delivery is missed, but the publisher apologizes, "We had a fire in the printing plant, but we have been able to make alternative arrangements and you will receive your journal within three weeks", have faith and do not punish the publisher by unsubscribing.

Being clear means letting the publisher know that you expect the journal in return for payment, and will stop paying if the journal is not delivered without any good reason.

Personal Life

Being nice, for example in a marriage, means never to start a quarrel.

Being provocable means to indicate how much you are prepared to do, not to give in to every unreasonable demand, not to be a "doormat" that lets everyone step over her.

Being forgiving is very important for a happy marriage. In the tournament of a prisoner's dilemma, one strategy retaliated once against the first defection, then twice if another defection occurred, then four times, then eight times, and so on, in case the other side still did not get the point. That strategy scored very poorly overall, it ended up mostly in mutual defections, which in real life corresponds to mutual violence, verbal or physical. One of the most destructive features of a marriage is if the partners keep repeating the same accusations over and over again, dozens of times. If something makes you unhappy, do not simply swallow it. Say it, but say it once, and then it is said, do not bring it up again and again. The psychologist and radio-host Dan Gottlieb once reported that the city of Houston in Texas had opened a museum on torture. He said, "They forgot one important form of torture, nagging." There is a saying, "If you want to avoid an earthquake, don't dwell on faults."

Being clear is also helpful. For example, if it is the husband's task to take out the garbage, and he forgot, his wife should not think, "He knew exactly that it is garbage collection day today and he did not take the garbage can outside on purpose, to make me angry." Maybe he simply forgot, and was preoccupied with something that worried him in his job. It is much better to say, "Darling, I am sure you remember that today is garbage day", and he will gladly take out the garbage. Say what you wish, and your wishes may be fulfilled. If you keep them secret, nobody may ever know what your would really like.

Peace and Security

Being nice means that if you want to avoid war, never start a war, something obvious.

Being provocable means to be able to defend oneself if faced with aggression. If you are totally defenseless, you may become a tempting target.

Being forgiving means, as the Koran says (8:61), "If the enemy inclines towards peace, do so also." This is very important to achieve peace, and through peace mutual security. Much has been written about the need for a strong defense and military strength as a deterrent against aggression. "Forgiveness" is not in the vocabulary of books about national security. Further below I will give and example that shows how important it is both to be able to defend oneself if attacked, but also to make clear to an opponent that as long as he does not attack, or stops the attack, he has nothing to fear.

Before World War I, the opinion prevailed that the best way to avoid war is to be prepared to win it, and the best way to win it was to deploy one's own forces more rapidly than the enemy, to invade the enemy's territory before he can invade us, so that whatever fighting there is takes place on his territory, not on ours. This led Europe to stumble rapidly into World War I. When a Serb nationalist assassinated Crown Prince Ferdinand, the heir to the Austrian throne, and his wife in Sarajevo in 1914, Austria declared war on Serbia, Serbia had a defense treaty with France, which in turn declared war on Austria, and so on, until a whole series of countries were drawn into war. They all put into effect their mobilization plans and moved their troops rapidly to their borders and beyond.

The great carnage that resulted from World War I, an estimated ten million dead, mostly soldiers, made people rethink the strategy that had precipitated this war. Many became convinced that it was necessary to be more restrained and patient, and not to retaliate all out to the first provocation. British Prime Minister Neville Chamberlain probably thought he had learned the lessons of Sarajevo when he flew to Munich to negotiate with Hitler and extracted from him a promise not to make any more territorial claims after he had annexed the Sudetenland from Czechoslovakia. Chamberlain returned to London, waving the sheet of paper that contained the agreement with Hitler's signature, and proudly announced, "I have brought peace for our time." Shortly later, Hitler invaded Poland. When someone asked him about the treaty he had just signed in Munich, he mocked, "Treaty? This is just a shred of paper."

The initial failure to resist Hitler's aggression, including the occupation of the Sarine region in 1936, may have emboldened him to seek a vast conquest. Resisting the first signs of aggression is important.

After World War II, the pendulum of strategic opinion swung back full cycle. NATO announced that it would not attack, but if attacked would not hesitate to use nuclear weapons first against larger conventional Soviet forces (the policy of "flexible response"). The Soviet Union announced that it would not use nuclear weapons first, but if attacked with nuclear weapons, would retaliate by firing all of its nuclear missiles at the United States. These strategies seem to be a strong deterrent against an attack. But not every war begins with an intentional attack.

Imagine the following hypothetical but not impossible scenario. During the Cold War, NATO helicopters regularly patrolled along the inter-German border to detect any concentration of troops and tanks as a possible prelude to a Soviet invasion of Western Europe. It could have happened, but fortunately did not, that one of these helicopters accidentally strayed into East Germany in dense fog and was shot down. The pilot, who survived, was going to be tried as a Western spy. NATO would have sent a rescue mission to save the pilot (as President Carter, at the insistence of his national security adviser Zbigniev Brzezinsky, sent a failed mission to rescue the American hostages in Teheran in 1979, and as Israel did in July 2006, invading and bombing Gaza and Lebanon allegedly to rescue three abducted Israeli soldiers). The members of the Warsaw Treaty would have brought reinforcements to defend their territory against this invasion. NATO would bring reinforcements, and, facing superior numbers of conventional Warsaw Treaty forces, might have resorted to "tactical" nuclear weapons. This could have triggered a Soviet nuclear retaliation against the United States.

The threat of excessive retaliation may deter an intentional attack, but it cannot prevent war by accident, and if that happens, it is disastrous. A better strategy is to resist an invasion of one's own territory, but at the same time to make clear to an opponent that if he does not attack us, he has nothing to fear. This should not only be announced in words, but demonstrated by the structure of armed forces, that they serve purely defensive purposes and cannot possibly be used for aggression, even if intentions were to change.

A typical example of a purely defensive measure is a tank trap, a hole in the ground covered with branches and leaves. If a tank moves over the hole, it falls down and is rendered useless. Such defenses pose no threat to another country as long as it does not attack us.

If, on the contrary, a country announces that it would pursue an aggressor into his own territory and annihilate him, this gives every incentive to the opponent to fight as long and hard as possible, because withdrawal would mean defeat and likely death. If by contrast we make clear to an opponent, "As long as you attack us, we will resist vigorously, but if you withdraw, we stop fighting, too," this gives the opponent the maximum incentive to stop fighting as soon as possible.

In other words, an effective defense must be both provocable and forgiving, not only one or the other. This corresponds to non-offensive defense.

Being clear means to tell other countries clearly that we wish to live in peace with them, that we have no aggressive intentions, but that if attacked we would defend ourselves. General Alexander Haig, Reagan's National Security Adviser, once recommended not to let our enemies know what we may or may not do, better to keep them in the dark, keep them guessing, so that they cannot sleep quietly.

Let us consider this. At the beginning of World War II, General Guisan, the commander of the Swiss army, called the German military attache to his office and told him, "If one German soldier sets foot onto Swiss soil, all the tunnels through the Alps will be blown up, they are already mined." Digging a tunnel through the Alps can take as much as 20 years, it is not something that could be done during a war. Germany used those tunnels to move empty railway cars from Italy to Germany, which then brought supplies to Italy through the Brenner tunnel in Austria. It is said that this warning, together with numerous other measures that formed part of a comprehensive military and civilian strategy called "general defense" helped persuade Germany to drop its plans for the invasion of Switzerland. If Guisan had kept this strategy as a military secret, he might have had to implement it, but that would have been too late. It is better to let an opponent know in advance what the consequences of aggression would be, to help dissuade aggression.

In other words, even in the realm of peace and security, being clear can be a useful strategy.

Frequent Meetings Encourage Cooperation

Axelrod has speculated that there tends to be less crime in small rural communities than in big cities, because in a small community everybody else knows you, and if you misbehave, you lose people's respect and their willingness to help you when you need it. But in a big city, people can practically go "underground". If they are impolite to someone on the street, chances are that they will never meet the same person again. "Atomie", the lack of a network of close acquaintances, who know what you are doing, can lead to "anomie", disrespect for social norms of behavior.

Similarly, people who play a single game of prisoner's dilemma with someone else tend to defect, but if they play repeatedly with the same partner, they can come to a tacit agreement to cooperate, for mutual benefit.

Such a phenomenon was also observed in World War I. When French and German soldiers faced each other in trenches, often for many weeks without any movement in the front-line, they discovered that if they did not shoot first, but only shot back when the other side started shooting, this persuaded the soldiers on the other side not to shoot first either. This led to long quiet periods on the front-line, without any official ceasefire agreement, simply through a silent understanding among the opposing soldiers. Officers, who were in safe positions behind the front and gave orders over the telephone, did not like this. They wanted their soldiers to fight and defeat the enemy. After some time, they discovered that if they moved the soldiers around frequently, they were not able to establish such forms of informal cooperation with their opponents.

Evolutionary Simulations

Axelrod and his colleagues also did an experiment that imitates the formation of the first forms of life in a primordial soup. They started with a mix of various strategies and let them have random encounters. Each time two strategies meet, they play a series of several hundred prisoner's dilemma games with each other, and each obtains a score. After some time, the top third of the strategies ranked according to their scores is increased by 10%, the middle third is left unchanged, and the bottom third is decreased by 10% in their numbers. Then the process of random encounters is repeated.

An interesting result emerged: If there are enough "nice" strategies initially (such as tit-for-tat and tit-for-two-tat etc.), the nice strategies score better than the nasty, aggressive strategies, with the latter gradually dying out, until there is a world populated with only nice strategies that cooperate constantly. However, if there are too few nice strategies to begin with, they are ravaged by the more aggressive, vicious strategies and have too few opportunities to get a high score in an encounter with other nice strategies. The nice strategies gradually die out, leaving a "nasty, brutish" jungle, a Hobbesian world. This would suggest that a group of nonviolent peace workers immersing itself into a war-torn society must be massive to have an impact. In a less belligerent society, although it may have some problems and conflicts, even a small number of peace workers may be able to make a difference for the better.

Axelrod also reported in his book that colonies of bacteria cooperate and protect each other, and do not devour each other. Did they understand the theory of the prisoner's dilemma and the advantage of cooperation? Of course not. In the beginning, there must also have been bacteria that eat each other and fail to support one another. But they have long died out. Those who survived until today are those who happened to cooperate.

We humans do not have the luxury to learn from such long periods of trial and error. With nuclear weapons, one error could be the final one. Fortunately, we have the capacity to learn from the failures and successes of others, even other species, like bacteria. We better learn to cooperate before it is too late.