Unit 3
GAME THEORY

Lesson 29

Learning Objective:

- On completion of this lesson you will be familiar with the key economic models dealing with interactive competition amongst firms when the behavior of rivals must be accommodated.
- You will learn to apply creative approaches to game theory.
- This course concentrates on the thinking process itself and the application of this process to decisions under conflict. You will apply new theories and creative approaches on a challenge from your work.

Hello students,

Let us begin with a very important practical game of the concept under our study.

The Prisoners' Dilemma

Tucker's invention of the Prisoners' Dilemma example was very important. This example, which can be set out in one page, could be the most influential one page in the social sciences.

This remarkable innovation did not come out in a research paper, but in a classroom.

While addressing an audience of psychologists at Stanford University, where he was a visiting professor, Mr. Tucker created the Prisoners' Dilemma to illustrate the difficulty of analyzing" certain kinds of games. "Mr. Tucker's simple explanation has since given rise to a vast body of literature in subjects as diverse as philosophy, ethics, biology, sociology, political science, economics, and, of course, game theory."

The Game

Tucker began with a little story, like this:

- Two burglars, Bob and Al, are captured near the scene of a burglary and are given the "third degree" separately by the police.
- Each has to choose whether or not to confess and implicate the other.
- If neither man confesses, then both will serve one year on a charge of carrying a concealed weapon.
- If each confesses and implicates the other, both will go to prison for 10 years.
- However, if one burglar confesses and implicates the other, and the other burglar does not confess, the one who has collaborated with the police will go free, while the other burglar will go to prison for 20 years on the maximum charge.

The strategies in this case are: confess or don't confess. The payoffs (penalties, actually) are the sentences served.

We can express all this compactly in a "payoff table" of a kind that has become pretty standard in game theory.

Just look at the payoff table for the Prisoners' Dilemma game:

Table 1

		Al	
		Confess	Don’t
Bob	Confess	10,10	0,20
	Don’t	20,0	1,1

The table is read like this:

- Each prisoner chooses one of the two strategies. In effect, Al chooses a column and Bob chooses a row.
- The two numbers in each cell tell the outcomes for the two prisoners when the corresponding pair of strategies is chosen.
- The number to the left of the comma tells the payoff to the person who chooses the rows (Bob) while the number to the right of the column tells the payoff to the

person who chooses the columns (Al). Thus (reading down the first column) if they both confess, each gets 10 years, but if Al confesses and Bob does not, Bob gets 20 and Al goes free.

So:

How to solve this game?

What strategies are "rational" if both men want to minimize the time they spend in jail?

Al might reason as:

"Two things can happen: Bob can confess or Bob can keep quiet. Suppose Bob confesses. Then I get 20 years if I don't confess, 10 years if I do, so in that case it's best to confess. On the other hand, if Bob doesn't confess, and I don't either, I get a year; but in that case, if I confess I can go free. Either way, it's best if I confess. Therefore, I'll confess."

But Bob can and presumably will reason in the same way -- so that they both confess and go to prison for 10 years each. Yet, if they had acted "irrationally," and kept quiet, they each could have gotten off with one year each.

Dominant Strategies

What has happened here is that the two prisoners have fallen into something called a "dominant strategy equilibrium."

DEFINITIONS

What is Dominant Strategy?

Let an individual player in a game evaluate separately each of the strategy combinations he may face, and, for each combination, choose from his own strategies the one that gives the best payoff. If the same strategy is chosen for each of the different combinations of strategies the player might face, that strategy is called a "dominant strategy" for that player in that game.

Now,

What is Dominant Strategy Equilibrium?

If, in a game, each player has a dominant strategy, and each player plays the dominant strategy, then that combination of (dominant) strategies and the corresponding payoffs are said to constitute the dominant strategy equilibrium for that game.

In the Prisoners' Dilemma game, to confess is a dominant strategy, and when both prisoners confess, that is a dominant strategy equilibrium.

Now, let us see some issues regarding Prisoners' Dilemma

Issues With Respect to the Prisoners' Dilemma

This remarkable result -- that individually rational action results in both persons being made worse off in terms of their own self-interested purposes -- is what has made the wide impact in modern social science. For there are many interactions in the modern world that seem very much like that, from arms races through road congestion and pollution to the depletion of fisheries and the overexploitation of some subsurface water resources. These are all quite different interactions in detail, but are interactions in which (we suppose) individually rational action leads to inferior results for each person, and the Prisoners' Dilemma suggests something of what is going on in each of them. That is the source of its power.

Having said that, we must also admit candidly that the Prisoners' Dilemma is a very simplified and abstract -- if you will, "unrealistic" -- conception of many of these interactions. A number of critical issues can be raised with the Prisoners' Dilemma, and each of these issues has been the basis of a large scholarly literature:

- The Prisoners' Dilemma is a two-person game, but many of the applications of the idea are really many-person interactions.
- We have assumed that there is no communication between the two prisoners. If they could communicate and commit themselves to coordinated strategies, we would expect a quite different outcome.
- In the Prisoners' Dilemma, the two prisoners interact only once. Repetition of the interactions might lead to quite different results.
- Compelling as the reasoning that leads to the dominant strategy equilibrium may be, it is not the only way this problem might be reasoned out. Perhaps it is not really the most rational answer after all.

We will consider some of these points in what follows.

An Information Technology Example

Game theory provides a promising approach to understanding strategic problems of all sorts, and the simplicity and power of the Prisoners' Dilemma and similar examples make them a natural starting point. But there will often be complications we must consider in a more complex and realistic application.

Let's see how we might move from a simpler to a more realistic game model in a real-world example of strategic thinking: choosing an information system.

For this example,

The players will be a company considering the choice of a new internal e-mail or intranet system, and a supplier who is considering producing it.

The two choices are to install a technically advanced or a more proven system with less functionality.

Assume that the more advanced system really does supply a lot more functionality, so that the payoffs to the two players, net of the user's payment to the supplier, are as shown in Table 2.

Table 2

		User	
		Advanced	Proven
Supplier	Advanced	20,20	0,0
	Proven	0,0	5,5

We see that both players can be better off, on net, if an advanced system is installed. (We are not claiming that that's always the case! We're just assuming it is in this particular decision). But the worst that can happen is for one player to commit to an advance system while the other player stays with the proven one. In that case there is no deal, and no payoffs for anyone. The problem is that the supplier and the user must have a **compatible standard,** in order to work together, and since the choice of a standard is a strategic choice, their strategies have to mesh.

Although it looks a lot like the Prisoners' Dilemma at first glance, this is a more complicated game. We'll take several complications in turn:

Looking at it carefully, we see that there this game has no dominated strategies. The best strategy for each participant depends on the strategy chosen by the other participant. Thus, we need a new concept of game-equilibrium, that will allow for that complication.

When there are no dominant strategies, we often use an equilibrium conception called the Nash Equilibrium**,** named after Nobel Memorial Laureate John Nash. The Nash Equilibrium is a pretty simple idea: we have a Nash Equilibrium if each participant chooses the best strategy, given the strategy chosen by the other participant.

In the example,

If the user opts for the advanced system, then it is best for the supplier to do that too. So (Advanced, Advanced) is a Nash-equilibrium.

But, hold on here!

If the user chooses the proven system, it's best for the supplier to do that too.

There are two Nash Equilibria! Which one will be chosen?

It may seem easy enough to opt for the advanced system which is better all around, but if each participant believes that the other will stick with the proven system -- being a bit of a stick in the mud, perhaps -- then it will be best for each player to choose the proven system -- and each will be right in assuming that the other one is a stick in the mud! This is a danger typical of a class of games called coordination games -- and what we have learned is that the choice of compatible standards is a coordination game.

- We have assumed that the payoffs are known and certain. In the real world, every strategic decision is risky -- and a decision for the advanced system is likely to be riskier than a decision for the proven system. Thus, we would have to take into account the players' subjective attitudes toward risk, their **risk aversion,** to make the example fully realistic. We won't attempt to do that in this example, but we must keep it in mind.

- The example assumes that payoffs are measured in money. Thus, we are not only leaving risk aversion out of the picture, but also any other subjective rewards and penalties that cannot be measured in money. Economists have ways of measuring subjective rewards in money terms -- and sometimes they work -- but, again, we are going to skip over that problem and assume that all rewards and penalties are measured in money and are transferable from the user to the supplier and vice versa.

- Real choices of information systems are likely to involve more than two players, at least in the long run -- the user may choose among several suppliers, and suppliers may have many customers. That makes the coordination problem harder to solve. Suppose, for example, that "beta" is the advanced system and "VHS" is the proven system, and suppose that about 90% of the market uses "VHS." Then "VHS" may take over the market from "beta" even though "beta" is the better system. Many economists, game theorists and others believe this is a main reason why certain technical standards gain dominance. (This is being written on a Macintosh computer. Can you think of any other possible examples like the beta vs. VHS example?)

- On the other hand, the user and the supplier don't have to just sit back and wait to see what the other person does -- they can sit down and talk it out, and commit themselves to a contract. In fact, they have to do so, because the amount of payment from the user to the supplier -- a strategic decision we have ignored until now -- also has to be agreed upon. In other words, unlike the Prisoners' Dilemma, this is a **cooperative game,** not a **non-cooperative game.** On the one hand, that will make the problem of coordinating standards easier, at least in the short run. On the other hand, Cooperative games call for a different approach to solution.

So let us recapitulate Zero-Sum Games

By the time Tucker invented the Prisoners' Dilemma, Game Theory was already a going concern. But most of the earlier work had focused on a special class of games: zero-sum games.

In his earliest work, von Neumann made a striking discovery. He found that if poker players maximize their rewards, they do so by bluffing; and, more generally, that in many games it pays to be unpredictable. This was not qualitatively new, of course -- baseball pitchers were throwing change-up pitches before von Neumann wrote about mixed strategies. But von Neumann's discovery was a bit more than just that. He discovered a unique and unequivocal answer to the question:

"How can I maximize my rewards in this sort of game?" without any markets, prices, property rights, or other institutions in the picture.

It was a very major extension of the concept of absolute rationality in neoclassical economics. But von Neumann had bought his discovery at a price. The price was a strong simplifying assumption: von Neumann's discovery applied only to zero-sum games.

For example,

Consider the children's game of "**Matching Pennies**."

In this game, the two players agree that one will be "even" and the other will be "odd." Each one then shows a penny. The pennies are shown simultaneously, and each player may show either a head or a tail. If both show the same side, then "even" wins the penny from "odd;" or if they show different sides, "odd" wins the penny from "even". Here is the payoff table for the game.

Table 3

		Odd	
		Head	Tail
Even	Head	1, -1	-1,1
	Tail	-1,1	1, -1

If we add up the payoffs in each cell, we find 1-1=0. This is a "zero-sum game."

Zero-Sum game: If we add up the wins and losses in a game, treating losses as negatives, and we find that the sum is zero for each set of strategies chosen, then the game is a "zero-sum game."

In less formal terms, a zero-sum game is a game in which one player's winnings equal the other player's losses. Do notice that the definition requires a zero sum for every set of strategies. If there is even one strategy set for which the sum differs from zero, then the game is not zero sum.

Another Example

Here is another example of a zero-sum game. It is a very simplified model of price competition. Like Augustin Cournot (writing in the 1840's) we will think of two companies that sell mineral water. Each company has a fixed cost of $5000 per period, regardless whether they sell anything or not. We will call the companies Perrier and Apollinaris, just to take two names at random.

The two companies are competing for the same market and each firm must choose a high price ($2 per bottle) or a low price ($1 per bottle). Here are the rules of the game:

1) At a price of $2, 5000 bottles can be sold for a total revenue of $10000.

2) At a price of $1, 10000 bottles can be sold for a total revenue of $10000.

3) If both companies charge the same price, they split the sales evenly between them.

4) If one company charges a higher price, the company with the lower price sells the whole amount and the company with the higher price sells nothing.

5) Payoffs are profits -- revenue minus the $5000 fixed cost.

Here is the payoff table for these two companies

Table 4

		Perrier	
		Price=$1	Price=$2
Apollinaris	Price=$1	0,0	5000, -5000
	Price=$2	-5000, 5000	0,0

(Verify for yourself that this is a zero-sum game.)

For two-person zero-sum games, there is a clear concept of a solution. The solution to the game is the maximin criterion -- that is, each player chooses the strategy that maximizes her minimum payoff.
In this game, Appolinaris' minimum payoff at a price of $1 is zero, and at a price of $2 it is -5000, so the $1 price maximizes the minimum payoff. The same reasoning applies to Perrier, so both will choose the $1 price.

Here is the reasoning behind the maximin solution: Apollinaris knows that whatever she loses, Perrier gains; so whatever strategy she chooses, Perrier will choose the strategy that gives the minimum payoff for that row. Again, Perrier reasons conversely.

SOLUTION: Maximin criterion For a two-person, zero sum game it is rational for each player to choose the strategy that maximizes the minimum payoff, and the pair of strategies and payoffs such that each player maximizes her minimum payoff is the "solution to the game."

Mixed Strategies

Now let's look back at the game of matching pennies.

> It appears that this game does not have a unique solution. The minimum payoff for each of the two strategies is the same: -1. But this is not the whole story. This game can have more than two strategies. In addition to the two obvious strategies, head and tail, a player can "randomize" her strategy by offering either a head or a tail, at random, with specific probabilities. Such a randomized strategy is called a "mixed strategy." The obvious two strategies, heads and tails, are called "pure strategies." There are infinitely many mixed strategies corresponding to the infinitely many ways probabilities can be assigned to the two pure strategies.

DEFINITION

Mixed strategy If a player in a game chooses among two or more strategies at random according to specific probabilities, this choice is called a "mixed strategy."

The game of matching pennies has a solution in mixed strategies, and it is to offer heads or tails at random with probabilities 0.5 each way.

Here is the reasoning:

- If odd offers heads with any probability greater than 0.5, then even can have better than even odds of winning by offering heads with probability 1.
- On the other hand, if odd offers heads with any probability less than 0.5, then even can have better than even odds of winning by offering tails with probability 1.
- The only way odd can get even odds of winning is to choose a randomized strategy with probability 0.5, and there is no way odd can get better than even odds.
- The 0.5 probability maximizes the minimum payoff over all pure *or mixed* strategies.
- And even can reason the same way (reversing heads and tails) and come to the same conclusion, so both players choose 0.5.

Von Neumann's Discovery

We can now say more exactly what von Neumann's discovery was.

Von Neumann showed that every two-person zero sum game had a maximin solution, in mixed if not in pure strategies. This was an important insight, but it probably seemed more important at the time than it does now. In limiting his analysis to two-person zero sum games, von Neumann had made a strong simplifying assumption. Von Neumann was a mathematician, and he had used the mathematician's approach: take a simple example, solve it, and then try to extend the solution to the more complex cases. But the mathematician's approach did not work as well in game theory as it does in some other cases. Von Neumann's solution applies unequivocally only to "games" that share this zero-sum property. Because of this assumption, von Neumann's brilliant solution was and is only applicable to a small proportion of all "games," serious and non-serious. Arms races, for example, are not zero-sum games. Both participants can and often do lose. The Prisoners' Dilemma is not a zero-sum game, and that is the source of a major part of its interest. Economic competition is not a zero-sum game. It is often possible for most players to win, and in principle, economics is a win-win game. Environmental pollution and the overexploitation of resources, again, tend to be lose-lose games: it is hard to find a winner in the destruction of most of the world's ocean fisheries in the past generation. Thus, von Neumann's solution does not -- without further work -- apply to these serious interactions.

The serious interactions are instances of "non-constant sum games," since the winnings and losses may add up differently depending on the strategies the participants choose. It is possible, for example, for rival nations to choose mutual disarmament, save the cost of weapons, and both be better off as a result -- so the sum of the winnings is greater in that case. In economic competition, increasing division of labor, specialization, investment, and improved coordination can increase "the size of the pie," leading to "that universal opulence which extends itself to the lowest ranks of the people," in the words of Adam Smith. In cases of environmental pollution, the benefits to each individual from the polluting activity is so swamped by others' losses from polluting activity that all can lose -- as we have often observed.

Poker and baseball are zero-sum games. It begins to seem that the only zero-sum games are literal games that human beings have invented -- and made them zero-sum -- for our own amusement. "Games" that are in some sense natural are non-constant sum games. And even poker and baseball are somewhat unclear cases.

A "friendly" poker game is zero-sum, but in a casino game, the house takes a proportion of the pot, so the sum of the winnings is less the more the players bet. And even in the friendly game, we are considering only the money payoffs -- not the thrill of gambling and the pleasure of the social event, without which presumably the players would not play. When we take those rewards into account, even gambling games are not really zero-sum.

Von Neumann and Morgenstern hoped to extend their analysis to non-constant sum games with many participants, and they proposed an analysis of these games. However, the problem was much more difficult, and while a number of solutions have been proposed, there is no one generally accepted mathematical solution of non-constant sum games.

To put it a little differently, there seems to be no clear answer to the question, "Just what is rational in a non-constant sum game?"

So, now let us summarize today's discussion:

Summary

We have discussed about:

- Several Examples
- The Prisoner's Dilemma
- Information Technology example
- The game of matching pennies