1 Introduction

Game theory is the study of *strategic decision making*, what people should do when faced with common and conflicting interests.

It models not our usual games of Chess or Checkers, but instead the kinds of games that take place between people—rational, thinking beings who have their own interests. Developed around 1928 by John von Neumann and initially applied to economics, it has found applications to fields as wide apart as evolutionary biology and political science. In fact, the study of the difference between the optimal strategies prescribed by game theory and what people actually do is an interesting problem in psychology.

Game theory shattered many of the intuitive facts that people held for granted about economics, and in fact decision-making in general. Without game theory, you might think the following.

1. If in a certain "game" that you knew everything about, and you are faced with a certain situation, then you should always make the same move... right? After all, you can calculate what would happen if you took different moves, and one of them would be better than the other. Take Tic-Tac-Toe for instance. If you were to analyze every possible move, then some positions would lead to wins or draws and others would lead to losses, so you would just pick the ones that are good for you.

2. Capitalism is good, right? I mean, it's worked the best. The basic idea is that with everyone furthering their own interests, people will create more value for the world—whether it is that better computer or bestselling book.

Of course the above two ideas are gross oversimplifications, but game theory shatters both of these conceptions in spectacular ways. It's hard to imagine what economics would be like today without game theory. By learning game theory, you will learn the *language* to talk about cooperation and competition in the real world, and learn the *method* to mathematically describe and then solve these kinds of problems.

1.1 Objectives

In this lesson you will learn to

1. *recognize* situations that can be modeled by game theory.

2. *describe* in plain language how game theory models situations in the real world such as
   - arms races
   - goats fighting for mates
   - responses to global warming
   - competing corporations

3. *classify* situations into common categories such as prisoners' dilemmas, games of chicken, dollar bill auctions, and so forth, and in so doing, understand the common traps that associated with these situations

4. given a two-person zero-sum game where each person has two choices, *calculate* the optimal strategy for both players.

2 Coin Poker

Game 1 (Coin poker)**:** Consider the following game. There are two players, Player I and Player II.

1. First, Player I puts down \$30 on the table and player II puts down \$20 on the table.
2. Next, Player I flips a coin, and keeps the result hidden from Player II. Then Player I either *raises* by putting down \$30 more on the table, or *passes* by doing nothing.
3. If Player I raised, then Player II must either *see* by also putting in more money—\$20 more—or *fold*, that is, give up. If Player II folds, then (s)he is out of the game, and Player I takes all the money on the table.
4. Otherwise, players look at the coin. If heads, Player I wins all the money, and if tails, player II wins all the money.

(In this version of the game, Player II is allowed to look at the coin even if (s)he folds.)

Which player do you want to be? Why? Play this game a few times; can you come up with a strategy? In the long run, does Player I or Player II win money?

2.1 Discussion

What are some possible strategies for each player? Let's consider for the moment *pure* strategies: strategies that tell you to do the same thing when faced with the same circumstances. For instance, one pure strategy for player I would be to pass if you flip tails (because if your bet and your opponent sees, you'd lose money) and bet if you flip heads (because you'd be in a position to win more money).

> **Problem 1:** What are two pure strategies that it might make sense for Player I to follow? What about Player II?

Player I might potentially do the following.

1. Pass on tails and raise on heads.
2. Raise on tails and raise on heads. ("Bluff" when you have tails, in the hope that your opponent will think you had heads and fold.)

Note that the other two strategies—raise on tails and pass on heads, or pass on both of them—are worse than these strategies no matter what the outcome, because you can always do at least as well if you raise on heads, compared to when you pass on heads. We say that these strategies are **dominated** by the strategies (1) and (2), and do not consider them.

Player II might do the following.

1. Fold when Player I raises.
2. See when Player I raises.

(When Player I passes, Player II has no options.)

To compare the strategies, it's helpful to compute the **payoff matrix**, a table where we quantify how good each strategy is.

I \ II	Fold	See
Pass/Raise		
Raise/Raise		

To quantify how good a strategy is, we look at the *amount of money you're expected to win* if you follow that strategy.

> ▶ **Definition:** The **expected value** is the average value of a variable (for instance, the amount of money you win), weighted by the probability for each value.
> For instance, if you gain \$10 with probabilty .8 and lose \$20 with probability .2, then your expected winnings is
> $$.8(\$10) + .2(-\$20) = \$8 - \$4 = \$4.$$

Say you are Player I. Suppose you pass when you get tails and raise when you get heads, and your opponent folds whenever you raise. Then

- If you flip tails, you pass, and you lose \$30.

- If you flip heads, you raise, your opponent folds, and you win \$20.

Thus the average, or *expected* value that you'd earn is $\frac{\$30+(-\$20)}{2} = \$5$. Similarly, you can calculate the other values in this table.

I \ II	Fold	See
Pass/Raise	-5	5
Raise/Raise	20	-10

Which strategy should you pick as Player I? If you pick pass/raise, and Player II figures out your strategy, then Player II should fold—so you'd lose \$5 on average. If you raise/raise, and Player II figures out your strategy (i.e. you're bluffing half the time), then Player II should always see—so you'd lose \$10 on average. Either way it seems like you're in trouble.

However, there's no requirement that you follow a *pure* strategy. You could pass some of the time when you get a 1 and raise some of the time when you get a 1. This way, when you raise, your opponent wouldn't know whether you have a 1 or 2, and this could be to your advantage.

To model this, suppose you Pass/Raise with probability p, and Raise/Raise with probability $1-p$. Note that when p is small, the worst thing happens when your opponent folds, and when p is large, the worst thing happens when your opponent sees. We want to find the *maximum* of the *least* you can expect to earn no matter what your opponent does.

When p gets larger from 0, when your opponent folds you do better, and when p gets smaller from 1, when your opponent sees you do better. Thus, at some point, you'll do equally well whether your opponent folds or sees, and this will be the best we can do. (Draw a graph to see this.) We hence try to solve

$$\text{Expected winnings when opponent folds}\text{Expected winnings when opponent sees}$$
$$-5p + 20(1-p) = 5p - 10(1-p)$$
$$-25p + 20 = 15p - 10$$
$$30 = 40p$$
$$\frac{3}{4} = p.$$

This means that your optimal strategy is to raise always when you get heads, and

- With probability $\frac{3}{4}$, pass when you get tails.
- With probability $\frac{1}{1}4$, raise when you get tails.

In this case, your expected winning per game is

$$-5p + 20(1-p) = -25p + 20 = -25 \cdot \frac{3}{4} + 20 = 1.25 \text{ dollars}.$$

We say that our strategy is a mixed strategy because it doesn't tells us absolutely what to do in situation—we do different things depending on chance.

► **Definition:** A **pure strategy** is one where you do the same thing every time the same situation comes up.
A **mixed strategy** is one where you might do different things when the same situation comes up, depending on chance—think of it as following different pure strategies with different probabilities.

Problem 2: What is the optimal mixed strategy for player II to follow? Remember that player II is trying to *minimize* the amount that player I wins.

Using similar calculations, letting q be the probability player II follows the "fold" strategy, we solve

$$-5q + 5(1-q) = 20q - 10(1-q)$$

and find $q = \frac{3}{8}$: if player I raises, player 2 should

- Fold with probability $\frac{3}{4}$.
- Pass with probability $\frac{1}{4}$.

This will ensure that the second player loses no more than $1.25 on average.

2.2 Summary

In the previous section, we saw that by modeling the problem mathematically, we were able to find the best strategy for both players. We had to expand our notion of "strategy" to include strategies where we could do different things with different probabilities—essentially we are "flipping a coin" to figure out what move we should make.

Often the best strategy is a *mixed* strategy, one where you make different moves with different probabilities.

You can remember this saying something like: sometimes the best way to make a decision is with a coin toss! (Of course, you have to mathematically analyze it to see if this is the case.)

Why were we able to analyze this game mathematically so easily? This kind of game is especially nice to analyze because of one key property: it is zero-sum.

▶ **Definition:** A game is **zero-sum** if the total amount of money (or value) earned by both players is 0: whatever one player loses, the other player gains.

When we put ourselves in the shoes of player 1, we *assumed the opponent wants to do the worst thing for us*; this is the assumption behind our equations for p, q. In fact, we have the following general theorem.

Theorem 1 (von Neumann Minimax Theorem)**:** Given a two-person *zero-sum* game, there exists an optimal strategy for both players, in the following sense.

1. Following his optimal strategy, player I can be guaranteed earn at least V on average.
2. Following her optimal strategy, player II can prevent player I from earning more than V on average.

This quantity V is called the **value** of the game.

We found the value of "coin poker" to be $1.25 for player 1—following the optimal strategy player 1 can on average win $1.25 on average, and following her optimal strategy player 2 can make sure player 1 doesn't win more than $1.25 on average.

A (2-player) zero-sum game is especially easy to analyze because the fact that the opponent always wants the opposite of what we want makes the problem simpler. We can alway find optimal strategies, as follows. (This is somewhat simplified.)

1. First, model the problem mathematically by writing the payoff matrix for the problem.
2. Delete all *dominated* strategies—strategies that do worse than another strategy no matter what. (For instance, we deleted the Raise/Pass and Pass/Pass strategies.)
3. Now assign a probability to the options that player I has, and find the probabilities such that player I would expect to earn the same no matter what player II does. Plug in to see how much player I can expect to earn.
4. Do the same for player II.

! Don't be fooled into thinking that all games have nice solutions like zero-sum problems! Non-zero-sum games often do not have optimal strategies, and the problem becomes much more psychological rather than mathematical.

We'll spend the rest of the lesson looking at non-zero-sum games.

3 The Dollar Bill Auction

Take your pretend money, and let's play a game. Here are the rules.

Game 2 (Dollar bill auction)**:** I have a dollar bill up for auction, with a minimum bid of one cent. Think of it: you can earn a dollar by paying less than a dollar! The rules of the auction are the following.

1. The minimum bid is 1 cent.
2. You may raise any bid, but by at most 10 cents.
3. If no one responds for a while, I will say, "Going once, going twice, sold!" and the highest bidder will get the dollar bill, after paying me what (s)he bid.
4. No communication is allowed between the players.

However, there is one extra rule.

5. Both the highest bidder and the second highest bidder have to pay what they bid.

Play this a few times. What happens?

Here are tables for you to follow the excitement (if you aren't caught up in the excitement of bidding yourself).

Name	Bid

Name	Bid

Name	Bid

Name	Bid

3.1 Discussion

Martin Shubik introduced the dollar bill auction and published an account of it in 1971. He auctioned off many dollar bills at social gatherings, and on average, the dollar bill sold for $3.40, and he earned more than $5 per dollar bill sold.

The game has been studied in psychological experiments, all with similar results. In a study of 40 groups of college students, in every group the bid went over $1—that is, people (who are intelligent beings) bought a dollar for more than $1! What is this insanity?

Let's analyze the game. Most likely something like the following happens.

1. *Someone* will most likely begin the game. After all, you can start with a bid of 1¢, and who wouldn't want to get $1 with 1¢?

   The game then keeps going forwards. If someone else bids 1¢, then why shouldn't you bid a little more, 2¢? It would be ridiculous to have that first person get $1 for 1¢. The game goes on: suppose you bid 30¢and someone bids 31¢. Then you'd think: it's better to pay 32¢ for a dollar than *lose* 31¢! So you bid 32¢.

2. The game passes the 50¢ mark. What's special about this mark? When two players have bid 50¢ or more, then the auctioneer would now make a profit! Maybe this enters the players' minds, but still, *they'd* still be in a position to make a profit, if they bid more. Paying 52¢ for a dollar isn't too bad.

3. Someone offers 100¢ for the dollar. At this point, anyone who bids more would be paying more than $1 for the dollar, so why would the bidding continue? Suppose you made the previous bid of 99¢. If you stop, then your opponent pays $1 for a dollar (neither earning nor losing) and you lose 99¢! So it would be better to bid 101¢—because then all you'd lose is 1¢.

   But the same reasoning applies to your opponent—and the bidding spirals upwards.

In all experiments, bidding eventually reduced to a battle between two people—each of whom is already too heavily invested in the bidding to quit (better to pay $3.50 for a dollar than $3.40 and get no dollar...). Researchers noticed increase in stress—players sweated, shouted, and showed characteristics of extreme tension "similar to what parachutists experience just before jumping out of an airplane." Moreover, the battle stopped being just about the money, but more about pride.

When asked about the bidding afterwards, subjects tended to say, "My opponent went crazy!" and almost never blamed themselves. Moreover, players made the same mistakes even after multiple rounds of the game.

How did your game compare? Knowing this, how would you play differently?

4 The Prisoner's Dilemma

Game 3 (Prisoner's Dilemma)**:** You and your "friend" have committed a serious crime worth 10 years in jail, but the police do not have the evidence to implicate you on that count. However, the prosecutor does know that you and your "friend" are guilty of tax evasion, worth 1 year in jail.

The prosecutor would very much like to close the case, so he makes the following offer separately to you and your "friend," who are placed in different cells.

"Here is the plea bargain.

1. If you *confess* to your crime, then I will set you **free**, and we'll forgive your case of tax evasion. Your accomplice will be in jail for 10 years, and the case will be closed forever.
2. However, this offer is only valid if your friend doesn't confess. If both of you confess, then your confession isn't useful to us, and we'll jail you both for **5 years**.
3. If neither of you confess, we won't be able to implicate either of you—but you'll still be in jail for **1 year** for tax evasion.
4. Finally, if your friend confesses but you do not, you'll be the one in jail for **10 years**!

Remember that I have made the exact same offer to your accomplice. I await your answer tomorrow morning at 9. Just think—you can be free at 10!"

Note the goal of this game is to get the best outcome for yourself. Assume you don't care at all about the fate of your "friend," or any silly concepts like "honor."

Problem 3:

1. Organize the above information into a table. Both you and your opponent have 2 options, so there are 4 possibilities in total. What happens in each case?
2. Would you choose to confess or keep silent, and why? Remember that your opponent is mulling over the same problem! Do you have a "dilemma"?

You \ Opponent	Keep silent	Confess
Keep silent		
Confess		

Now play the game. What happened? How did your fellow pairs of prisoners in the class fare?

Pairs who...	Number
both kept silent	
kept silent/confessed	
both confessed	

4.1 Discussion: The Dilemma

Remember that we can keep track of a game using a *payoff matrix*: in each entry of the table we write a number representing how good the outcome is. Since you *don't* want to spend more years in jail, we'll write the number of years with a negative number. You'd prefer to spend 0 years in jail over spending 10 years in jail, which corresponds to the fact that

$$0 > -10.$$

If you keep silent and your friend also keeps silent, then you just get 1 year in jail. If you keep quiet and your friend confesses, then you're in big trouble—you get 10 years in jail.

You \ Opponent	Keep silent	Confess
Keep silent	$\mathbf{-1}$	$\mathbf{-10}$
Confess		

If you confess and your friend keeps silent, then you go free; if you confess and your friend also confesses, then you get 5 years in jail.

You \ Opponent	Keep silent	Confess
Keep silent	-1	-10
Confess	$\mathbf{0}$	$\mathbf{-5}$

So it seems like you would do better to confess in either case:

1. If your friend keeps silent, you would go free rather than spend a year in jail: $0 > -1$.
2. If your friend confesses, you would spend 5 rather than 10 years in jail: $-5 > -10$.

Using the terminology from the previous section, we might say that the "confess" strategy is dominating.

So your best strategy is to confess. Since your opponent is in the same position, he would also choose to confess.

4.2 The paradox

Let's write the payoff matrix again. Note that because the game is not zero-sum, we really should keep track of the outcomes for your opponent too, and we'll do that by writing both the outcomes for you and your opponent in the table.

You \ Opponent	Keep silent	Confess
Keep silent	$(-1, -1)$	$(-10, 0)$
Confess	$(0, -10)$	$(-5, -5)$

As we said, "intelligent" prisoners would both confess, because they'd both be better off no matter what their opponent did.

You \ Opponent	Keep silent	**Confess**
Keep silent	$(-1, -1)$	$(-10, 0)$
Confess	$(0, -10)$	$(\mathbf{-5}, \mathbf{-5})$

This means that following their "optimal" strategies, both players get 5 years in jail. But they could have done better if they had both kept silent: only 1 year in jail! What went wrong?

You \ Opponent	Keep silent	**Confess**
Keep silent	$(\mathbf{-1}, \mathbf{-1})$	$(-10, 0)$
Confess	$(0, -10)$	$(\mathbf{-5}, \mathbf{-5})$

4.3 The general prisoner's dilemma

The prisoner's dilemma comes up so often in applications that we want to have some unifying terminology for it.

1. We call the "keep silent" option cooperating. You're cooperating with your accomplice.
2. Call the "confess" option **competing**.

Using this new terminology, we see that the best outcome is when you compete and your friend cooperates (then your friend is really screwed). Second best is when you both cooperate, and worse is when you both compete.

4.4 The iterated prisoner's dilemma

Now what would happen if you played the prisoner's dilemma over and over?

Game 4 (Iterated Prisoner's Dilemma)**:** Play with a friend. Each turn check either "cooperate" or "compete," without showing your opponent. Then you and your friend receive the amount of dollars given in the table below.

You \ Opponent	Cooperate	Compete
Cooperate	$(3,3)$	$(0,5)$
Compete	$(5,0)$	$(1,1)$

Repeat.

First, make sure you can see why this game is a "prisoner's dilemma."

Play a bunch of games in a row (at least 20). What strategy did you use? Now break into 2 groups to discuss strategies, and try again. Did you do better the second time?

Here are tables you can use.

I \ II	Cooperate	Compete
Cooperate	$(3, 3)$	$(0, 5)$
Compete	$(5, 0)$	$(1, 1)$

I coop	I comp	II coop	II comp	$ earned

I coop	I comp	II coop	II comp	$ earned

4.5 Discussion: Iterated Prisoner's Dilemma

Can cooperation theoretically develop in a world where everyone is governed by their own interests?

Political scientist Robert Axelrob proposed a competition for the iterated prisoner's dilemma: write a computer program to play the game.

The winning program? Tit-for-tat, just two rules:

1. Cooperate in the first round.
2. In all subsequent rounds, do whatever the opponent did in the previous round.

5 Games for thought

Classify and comment on the following situations.

1. **Arms race**
2.