Introduction

The context for this unit

In earlier units, we modeled individuals making decisions which were constrained by technological knowledge and productivity of land, environmental considerations, and economic data. Now we will look at how other people and firms act as a constraint on behavior; specifically, we will model the constraining effect on behavior of other peoples anticipated actions. To do so, we introduce game theory.

This unit 1

In the introduction to this unit, the problem of climate change was discussed to highlight the issues presented by social dilemmas. Social dilemmas—like climate change—occur when people do not take adequate account of the effects of their decisions on others, whether these are positive or negative.

It is not always the case that what's best for an individual is best for society.

This unit 2

Who is cleaning the kitchen today?

If its never you and always your roommate, you are a free rider. Of course, I am sure that none of you would are trying to free ride through the group assignments required in this course.

The presence of common property creates a social dilemma. The tragedy of the commons is an example of how decisions which do not consider the effects on others can lead to socially sub-optimal outcomes.

The final topic we consider in this lecture is how we can resolve social dilemmas and why we sometimes fail.

Game Theory: Key concepts

Social and strategic actions

Devising policies which benefit society requires understanding if the incentives encourage individuals to behave in a way that promotes general well-being or lead to undesirable results.

As we go about our daily life, we often find ourselves in a situation where we interact with other people and making decision which affect them. Social interaction is a part of our daily life.

If we are aware of how our actions affect others, then we are engaged in a strategic interaction. And the action we take when we are aware of how are actions affect others is referred to as our strategy.

We will use games to model these strategic interactions.

Game

A game is a description of a social interaction. To model such an interaction, start by describing who is playing: who are the agents that are interacting with each other. Next, we define what course of actions they may pursue: what are the feasible strategies. The third step is to state what information the players have when making decisions. The final aspect is to describe the outcome of those actions.

In the next example, two farmers must decide which crops to sow.

Example: Crop choice

This is a one-shot game: the players will interact once. The single decision about which course of action to pursue is the only time the two players interact with each other and we assume that no one else is involved or affected. Each farmer must decide to grow rice or cassava without knowing what the other farmer is deciding to do: they make the decision simultaneously without any knowledge of what the other is doing.

The payoffs of their decisions are presented in the 2x2 matrix.

Anil's payoff is presented in the lower left of each quadrant, while Bala's payoff is presented in the upper right. The payoff each receives depends not only on which course of action they choose, but on what the other does as well.

Reading across the first row we can see the payouts if Anil decides to grow rice. Anil will receive a payout of 1 if Bala also grows rice and a payout of 2 if Bala grows cassava. Still in the first row, which presents the payouts when Anil grows rice, we see that Bala will receive a payout of 3 for choosing rice and 2 for cassava.

The matrix describes the 4 possible outcomes and what happens in each case. To determine the outcome (that is to find the equilibrium), we need to first introduce another concept: the best response.

Optimal decision making

To find the best response, we simply look for the strategy that yields the greatest payout. Recall that each players payout depends on what the other player does. And this must be taken into consideration.

Bala's best response if Anil grows rice is to also grow rice. We found this by considering Bala's payouts presented in the first row.

Bala's best response if Anil grows cassava, now looking at Bala's payout in the second row, is to grow rice.

Anil's best response if Bala grows rice, reading down the first column, is to grow cassava.

Anil's best response if Bala grows cassava, now reading down the second column, is to grow rice.

Crop choice example

In this game, each player has a dominant strategy. Regardless of what the other player does, there is a single course of action which yields the greatest payout for each player. And when both players have the dominant strategy, we have a dominant strategy equilibrium.

Regardless of what Anil does, Bala should grow rice. And regardless of what Bala does, Amil should grow cassava.

Nash equilibrium

Same game, different payouts.

Using the dot and circle approach, find the best response.

If Bala grows rice, what is Anil's best response. Read down column 1 and put a dot in the cell that represents the course of action which yields the highest payout. Repeat for when Bala grows cassava.

If Anil grows rice, what is Bala's best response? Read across row 1 and put a circle in the cell that yields the highest payout. Repeat for when Anil grows cassava.

Note that unlike the previous game, there is no dominant strategy.

But what we do have is two Nash equilibriums. There is a set of strategies which is a best response given the course of action by the other player.

If Anil grows rice, Bala should grow cassava. And, if Anil grows cassava, Bala should grow rice. The converse is also true.

Resolving social dilemmas

The prisoner's dilemma

Our two farmers now face a different problem, how to deal with pest insects.

There are two strategies: use the cheaper chemical which kills every insect but also contaminates the water both farmers use or release beneficial insects which is more expensive but doesn't contaminate the water.

If only 1 farmer choses the chemical, the damage is minimal. However, if both farmers chose the chemical, the problem becomes much graver.

Find the best response the same way as before. Verify that the dominant strategy for both players is to use the chemical.

In this game, the dominant strategy leads to a socially sub-optimal outcome. The total payoff of the dominant strategy is lower than the alternative strategy.

How might the game change if they could collaborate before making the decision simultaneously?

Check the video on the next slide to see one possible way to play this game.

Why did we predict this outcome?

The prisoner's dilemma contrasts with the invisible had game introduced earlier in the lecture. In one instance self-interest led to an optimal outcome while in the other it led to sub-optimal outcome.

Here you can see the three aspects of the interaction that led to the sub-optimal outcome and how they can be addressed.

Social preferences: Altruism

As we have seen in an earlier game, social dilemmas may arise when self-interest remains unchecked. However, most people do care about others.

To model altruistic preferences, we return to the concepts of indifference curves and the feasible set.

The feasible frontier here represents the possible combinations of the 10,000-rupee prize money shared between Anil and Bala. The preferences of Anil are represented by the indifference curves; the indifference curves show combinations of amounts for Anil and Bala for which Anil is indifferent.

If Anil were purely motivated by self interest and had no concern for Bala, he would keep the all the prize money. The indifference curves would be a vertical line. Just like before, he prefers curves further to the right. The best option is A.

However, if Anil cares about Bala, his indifference curves will be downward sloping. The more Anil cares about Bala, the flatter the curves will be. When Anil is somewhat altruistic, the best option is point B.

Resolving the prisoner's dilemma

Back to the prisoner's dilemma, a suboptimal outcome that arose from a failure to account for the costs the agent's actions inflict on others. So how does the game change when we introduce altruistic preferences.

Each players payoff is represented on an axis. The altruistic preferences cause the indifference curve to be downward sloping which contrasts with the vertical, completely selfish indifference curve. Each of the possible outcomes are represented by dots.

When completely selfish, Anil chooses T. This is exactly what we modeled earlier.

With altruistic preferences, Anil will choose I. (I,I) is preferred to (T,I)and (I,T) is preferred to (T,T).

For Anil, when selfish, the dominant strategy is the chemical pesticide and when altruistic, the dominant strategy is the environmentally friendly approach.

Caring about each other makes it easier to resolve social dilemmas.

Social preferences: Other types

Altruistic preferences are not the only type of social preferences. Here you can see a few others.

Lack of concern for others is only one reason that we arrived at less-than-optimal outcome in the prisoners' dilemma. Let's consider another reason.

Repeated games

The previous games were a one-shot game, the implication is that I can cause you harm and there are no repercussions. If we change the rules so that it is a repeating game, that is so that we interact consecutively, then social norms, reciprocity, and punishment may alter the dominant strategy.

Public goods game: Farming example

The public goods game describes a situation where agents decide to contribute which causes them to incur a personal cost but provides a benefit to everyone.

In this example, a group of 4 farmers are contributing to an irrigation project. If a farmer contributes $10, it yields a benefit of $8 to everyone. We are looking at the payoff for Kim which depends not only on if she contributes, but also on how many other farmers contribute.

If Kim is the only farmer to contribute, her payoff is -$ 2. If no one contributes, including Kim, she will receive a payoff of 0 which is preferred to -$2.

If Kim and one other farmer contribute, everyone receives a benefit of 2*$8=$16, but accounting for her contribution of $10 reduces her benefit to $6. If one other farmer contributes and Kim freerides, her benefit goes up to $8.

It becomes apparent, that while everyone is better off if everyone contributes, the dominant strategy is to freeride. Whatever, the other farmers decide, Kim's payoff is greatest if she does not contribute.

Elinor Ostrom, the first female to win the Nobel prize in economics and she was not even an economist, has contributed greatly to our understanding of the public goods problem. A history of trust and greater equality are two characteristics of communities which have been more successful in sustaining cooperation.

Reciprocity and social norms

A game much like the one we just described with Kim and the other farmers was recreated as an experiment all over the world. Initially, contributions were high. Altruistic preferences seemed apparent. But the tragedy becomes apparent with the downward trend observable in every country. It is likely that reciprocity was also at play here. People become disappointed that others are not contributing and thus reciprocate with lower contributions themselves.

What would happen if the free riders could be punished…and after very recently completing the first group assignment, I bet some your ears just perked up…what we can punish free riders?

Peer punishment

Punishing free riders increases contributions. If the punishment imposes a cost to the punisher, it is a form of altruism. The punisher is willing to incur a cost to deter free riding behavior which is detrimental to overall well-being of society.

Maybe y'all should go back and revise the team policies considering this newfound knowledge.

The ultimatum game

Here we have a two-person, one-shot, sequential game (as opposed to the simultaneous games earlier). The first mover is the proposer who has received $100 and must propose a split of the form x for proposer, y for responder where x+y=$100. Any split is fine, but it is a take it or leave it offer. If the responder rejects the offer, both the responder and proposer get nothing.

We are modeling an interaction from which economic rents arise. The proposer receives the pie provisionally. If the two parties can successfully divide the pie, if the proposer makes an offer the responder accepts, they both receive a rent, a slice of the pie. The reservation option or next best alternative is nothing.

The diagram presented on this slide is the game tree. It shows the actions, orders of actions, and results of those actions. As the proposer you must consider what the responder will do which makes this a strategic interaction.

The secret to this game as the proposer is to offer the minimum acceptable offer. The value of this offer is equivalent to the amount of pleasure the responder would get from punishing the proposer for violating social norms for what is considered fair.

Another approach for the proposer who cares only about their own payoffs is try and maximize the expected payoff. The expected payoff is calculated by multiplying the payoff you would receive if the offer is accepted by the probability it will be accepted.

Example: Kenyan farmers and US students

Again, like the public goods game, the ultimatum game has been recreated in experiments the world over. The height of each bar represents the share of respondents who accepted the offer which is presented as a share of the pie on the horizontal axis. The darker shade on the upper part of the bars represents the probability of the offer being rejected.

Apparently, US students do not care nearly as much about fairness as Kenyan farmers. Attitudes and social norms most definitely vary across country and other social characteristics.

The rules of the game matter

What if we change the rules of the game? Now suppose there are two responders; all it takes is one for the responders to accept the offer and split the pie between them and the proposer. If both responders accept, one is chosen at random to receive the payout.

When there is competition between the responders, they act much more like a purely self-interested individual. The share of offers rejected, presented on the vertical axis, declines with two responders, especially at the lowest offers.

The primary point is that the rules of the game matter!

Coordination issues

We have seen a variety of games, the invisible hand, prisoners' dilemma, and public goods game where the highest payoff did not depend on the other players strategy: these games had a dominant strategy.

Often in life, there is no dominant strategy. Which side of the road you drive on is an easy example— it depends what side of the road everyone else is driving on. We saw a game like this earlier when Anil and Bala were choosing crops.

Use the dot and circle method to confirm which outcomes are the two Nash equilibriums. When we confront a dilemma that has more than a single Nash equilibrium we must consider, which outcome do we expect to observe and do any of the outcomes create conflict because one outcome is preferable to some but not all.

The question confronting policy makers remains how we can coordinate action to yield the outcome which is socially optimal. This question remains elusive as we will see in the next slide as we conclude out lecture with a discussion on climate change.

Example: Climate change

In this game, the two players are the US and China are considering how to respond to global carbon emissions. Each country has two strategies they can adopt: reduce emissions or continue with business as usual. The payoff matrix presented on the slide is a bit different than the matrices we saw before; the payoffs are ranked best through worst. We refer to this ranking as an ordinal scale. The ordinal scale conveys which is better, but not by how much.

Find the dominant strategy. If China restricts, the US will receive the highest payout by free riding with business as usual approach. If china continues with business as usual, then the US receives the highest payout by also choosing business as usual. This is the prisoners' dilemma; society ends up at the socially sub-optimal outcome.

Let's change the rules and assume that our players are not purely motivated by self-interest but value inequality aversion and reciprocity. To model this, we change the payouts. Now the dominant strategy for both players is to restrict and society arrives at the socially optimal outcome. Only if it were this easy in real life. Why do you think its so difficult for countries to negotiate a binding treaty to reduce emissions?