NU Econ 101 Lecture 14: Game Theory

Lecturer: Shane Auerbach

Thursday 30/06/16

Where are we now?

What is game theory?

Game theory

The analysis of behavior in strategic environments.

- A **strategic** environment is one in which the outcome of your choices depends on the choices of other agents.

Who uses game theory?

- GT is omnipresent in contemporary economic theory.
- Also used commonly in biology, computer science, political science, military theory, engineering and sociology.

What is a game theoretic model?

A game theoretic model consists of two elements:

1. A description of a strategic situation (a *game*).
2. A set of behavioral/epistemic assumptions.
    - Usually implicit in the choice of a solution concept.

Game theory's noumenal role

Noumenal: an object as it is in itself independent of the mind.

- Noumenal analysis focuses on the theory in a vacuum, disconnected from the world.

Noumenal game theory:

- GT provides a set of tools to investigate the meaning and implications of rationality in strategic settings.
- A game should provide a complete (i.e. not simplified) description of the strategic situation faced by all players.
- Different solution concepts are used to explore the implications of different assumptions on players rationality as well as their properties.

Game theory's phenomenal role

Phenomenal: perceptible by the senses or through experience.

- Phenomenal analysis focuses on using theory to make predictions in the world.

Phenomenal game theory:

- GT tries to explain why things happen. But it is often very difficult to model real world situations as a game.
- Will Google profit from making their own smartphones?
  - GT does not directly answer the question. It tells you how to go from a given game to a prediction (for a given solution concept). It doesn't tell you what the game is.
  - But understanding abstract strategic concepts may help you identify important features of an environment.

What is a game?

In a **static** game, players make choices simultaneously, only once.

A static game G can be described as this collection of objects:

$$G = \langle I, (A_i, u_i)_{i \in I} \rangle$$

- $I = \{1, 2, \ldots, n\}$ is the set of players.
  - We'll often refer to the i^{th}, j^{th} and even $(-i)^{\text{th}}$ players.
- A_i is the set of actions (i.e. choices) of the i^{th} player.
- $u_i : A_1 \times A_2 \times \ldots \times A_n \to \mathbb{R}$ is utility function of player i.

An example: coordination game (1)

Suppose we play the following two-person game:

You are player 1 (P1) and I am player 2 (P2).

- Each of us chooses to place a coin with heads up or tails up.
- If we match, we are both given one dollar.
- If we do not match, we get nothing.

An example: coordination game (2)

We can represent this game in the **normal form**:

	H	T
H	1,1	0,0
T	0,0	1,1

- P1 chooses the row. P2 chooses the column.
- The first number in each cell is P1's utility from that outcome. The second is P2's.
  - For example: $u_1(H, H) = u_2(H, H) = 1$.

Where are we now?

What will people play in the coordination game?

Looking at the coordination game, what would be our prediction of how people might play?

	H	T
H	1,1	0,0
T	0,0	1,1

In a coordination game, incentives are aligned.

We might predict they either play (H, H) or (T, T) so that each gets one util (instead of zero).

- Note that game theory is not about 'winning' or 'losing'. What matters is how much utility you get, not how much you get relative to the other player.

Best response

To define Nash equilibrium, we first define a best response:

Best response

An action a_i^* is a **best response** to an opponent's action a_{-i} if,

$$\forall a_i \neq a_i^* \quad u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$$

If this is the case, we say $a_i^* \in r_i(a_{-i})$, or a_i^*.

Nash equilibrium

To get predictions from a game, we need a **solution concept**.

Most famous solution concept: **Nash equilibrium (NE)**.

Nash equilibrium

An action profile $a^* = (a_i^*)_{i \in I}$ is a pure Nash equilibrium if, for every i, $a_i^* \in r_i(a_{-i}^*)$.

In words: An action profile is a Nash equilibrium if each agent is best responding to all other agents.

Finding pure Nash equilibria in a normal form

Defining pure Nash equilibrium formally might a little technical, but finding them in simple games is very easy!

For each player, and for each profile of opponent actions, underline the utility resulting from the best response to that profile.

- In a two-player game, a profile of opponent action is just an opponent's action.

Let's do it for the coordination game . . .

Finding pure Nash equilibria in coordination game

First, for player 1, underline the utility corresponding to the best response to each of P2's actions:

	H	T
H	$\underline{1}$,1	0,0
T	0,0	$\underline{1}$,1

Then, for player 2, underline the utility corresponding to the best response to each of P1's actions:

	H	T
H	$\underline{1}$,$\underline{1}$	0,0
T	0,0	$\underline{1}$,$\underline{1}$

Every cell with both numbers underlined is a NE.

- (H,H) and (T,T) are NE.

Does the coordination game have phenomenal relevence?

The coordination game seems crazy simple.

Is it relevant to any real world situations?

Yes, any situation in which coordination is critical:

- What side of the street should de drive on?
  - u(all on left)=u(all on right) = Good. Otherwise bad!
- With which hand should we shake hands?
  - u(all use left)=u(all use right) = Good. Otherwise bad!
- In both situations above, **real world outcomes match NE**.

A game of conflict: matching pennies

We have seen a simple game of **coordination**. What about a simple game of **conflict**?

Matching pennies:

- Each of us chooses to place a coin with heads up or tails up.
- If we match, you win a dollar (I do not).
- If we do not match, I win a dollar (you do not).
- You are P1. I am P2.

	H	T
H	1,0	0,1
T	0,1	1,0

Finding pure Nash equilibria in matching pennies

First, for player 1, underline the utility corresponding to the best response to each of P2's actions:

	H	T
H	$\underline{1}$,0	0,1
T	0,1	$\underline{1}$,0

Then, for player 2, underline the utility corresponding to the best response to each of P1's actions:

	H	T
H	$\underline{1}$,0	0,$\underline{1}$
T	0,$\underline{1}$	$\underline{1}$,0

There are no pure Nash equilibria for matching pennies.

No pure Nash equilibria?

Some games, like matching pennies, have no pure Nash equilibria.

But there are also **mixed Nash equilibria**.

- In a mixed NE, players randomize their action.
- Learn about these in a game theory course!

Every game has at least one Nash equilibrium.

Makes sense to have no pure NE here as you do not want to play predictably in matching pennies. Better to randomize!

Where are we now?

Prisoner's dilemma (1)

Suppose you and I have committed a crime. We are being interrogated in separate rooms. We can cooperate with eachother (i.e. not nark the other one out) or we can defect (i.e. nark).

- If we both cooperate with eachother (i.e. we do not nark), then they send us to jail for a year each on lesser charges.
- If I cooperate and you defect, they send me to jail for three years and you get off free.
- If I defect and you cooperate, they send you to jail for three years and I get off free.
- If we both defect, we both go to jail for two years.

Prisoner's dilemma (2)

Suppose that for each of us, our utility is:

$$u(x \text{ years you spend in prison}) = 3 - x$$

- Note that this implies that you only care how long you spend in prison, not about your friend. If you did care, we would represent it in the utility function.

We can represent the game as follows:

	C	D
C	2,2	0,3
D	3,0	1,1

Nash equilibrium of Prisoner's dilemma (1)

Underlining best responses to find pure NE:

	C	D
C	2,2	0,<u>3</u>
D	<u>3</u>,0	<u>1</u>,<u>1</u>

The Nash equilibrium is (D,D).

In fact, D actually **strictly dominates** C for both players. Whether your opponent is playing C or D, D gives you strictly more utility.

- You'll define dominance more formally in a GT class!

Nash equilibrium of Prisoner's dilemma (2)

	C	D
C	2,2	0,<u>3</u>
D	<u>3</u>,0	<u>1</u>,<u>1</u>

Why is the Nash equilibrium troubling/surprising?

- The Nash equilibrium is *bad*. If we cooperate with eachother (didn't nark), we would get two utils. But Nash equilibrium predicts we will both defect, ending up with one util.
- (C,C) is a Pareto improvement over (D,D)
    - (D,D) is not Pareto efficient.

Sometimes strategic interaction leads to bad outcomes!

Prisoner's dilemma and the phenomenal

Are there prisoner's dilemmas in your life?

- Doing the dishes.
  - You and your roommate don't do the dishes. Would be better for both if each did own dishes immediately after using them.
- Working on a project with a partner.
  - You can work hard or be lazy. Better for both if you work hard, but individual incentives to be lazy.

Prisoner's dilemma and cartels

We saw yesterday, in studying cartels, that each member of a cartel had an incentive to cheat on the cartel (and overproduce). If C is producing half of the monopoly quantity and D is defecting to maximize own profits, we get the following game of profits:

	C	D
C	High, High	Low, <u>Very high</u>
D	<u>Very high</u>, Low	<u>Medium</u> , <u>Medium</u>

The prediction of Nash equilibrium on this game would be that cartels will not exist. But cartels do exists. Is NE wrong?

- No. More likely, we have modeled the strategic situation faced by members of a cartel (the game) incorrectly.

Where are we now?

Repeated games

We may have erred in modeling cartels as a static game.

- A static game means players interact just once.
- In a cartel, firms may interact over a long period of time.

To model a strategic situation with repeated interactions, we use repeated games . . .

Finitely repeated games

Suppose we are going to play a prisoner's dilemma N times. Then the game is the sequence of encounters. We call what happens in each period **the stage game**:

	C	D
C	2,2	0,$\underline{3}$
D	$\underline{3}$,0	$\underline{1}$,$\underline{1}$

- Since the stage game is just the prisoner's dilemma, the NE of the stage-game is still (D,D).
- **But can we sustain (C,C) in a Nash equilibrium of the repeated game?**

Backward induction

To solve for NE in finitely repeated games, we use a technique called **backward induction**.

- Essentially, we solve a game backwards, starting with the N^{th} stage and working our way back to the first stage.

The reason we don't solve forwards is that behavior in the first stage might be affected by what players believe about the consequences (i.e. punishment) of the first stage outcome on future behavior.

- Must understand stage $n+1$ before analyzing stage n.

What happens in stage N?

In stage N, no matter what happened before, we play (D,D).

- Once we're at stage N, the game we face is identical to the static game.
- Maybe we have played (C,C) in the past, but might as well cheat on the cartel in the last period as there is no subsequent period for punishment.

What happens in earlier stages?

In stage N − 1, no matter what happened, we play (D,D).

- We know that we will play (D,D) in stage N no matter what we do here. Given that, whether you are playing C or D today, it is better for me to play D.

But this same logic means you should play D in period $N-2$, which in turn means you should play D in $N-3$ and this continues all the way back to the first stage!

The unique Nash equilibrium of the finitely repeated prisoner's dilemma is playing (D,D) in every stage!

Intuition of the result

The intuition of our result is very simple:

- We both know the other will play D in the last period no matter what.
- Given that, we both know each of us will play D in the second-to-last period no matter what.
- Given that, we both know each of us will play D in the third-to-last period no matter what.
- ...
- Given that, we both know we will play D in the first period.

Where are we now?

Infinitely repeated games (1)

But, again, we know cartels exist. What are we doing wrong?

What screwed us was the existence of a last stage.

- Everybody maximizes self-interest in the last period.
- But that starts the unraveling!

To get around this last period problem, let's consider modeling this as an infinitely repeated game!

Infinitely repeated games (2)

Do firms interact over an infinite time frame?

- Of course not!

But, using a model of infinitely repeated games is also valid for instances when it is uncertain when an interaction ends.

- Any stage could be the last, but players do not know whether the current stage is or is not.
- We can get them to behave because they fear consequences in possible future interactions (which may or may not occur).

Discount rate

Let's suppose that in any given stage, there is a probability of $\delta \in (0, 1)$ that there will be another stage.

- We will call δ the **discount rate**.

We call δ the discount rate because it tells us how much we should value a potential payoff in a future stage.

- A potential payoff of one util tomorrow is worth $\delta \cdot 1 = \delta$ to us today as δ is the probability that there will be a tomorrow.
- A potential payoff of one util in two stages is worth δ^2.

Valuing infinite payoff strings

How much is a potentially infinite string of payoffs worth?

Suppose we're trying to value a string of payoffs of one util going infinitely out into the future. Let's say it has value S.

$$S = 1 + \delta + \delta^2 + \delta^3 + \delta^4 + \ldots$$
$$S = 1 + \delta \left(1 + \delta + \delta^2 + \delta^3 + \delta^4 + \ldots\right)$$

But wait, that junk in the parentheses is just S!

$$S = 1 + \delta \cdot S$$
$$S = \frac{1}{1 - \delta}$$

Strategies in infinitely repeated games

A Nash equilibrium is a profile of strategies.

- In the static game, an action (like C, or D) is a strategy.
- In games where more than one decision is made, a strategy must specify an action for each possible decision-making scenario.
  - In an infinitely repeated game, that's infinitely many possible decision-making scenarios that we must specify in a strategy.

A profile of strategies is one strategy for each player.

Grim trigger strategy profile

Consider a strategy profile where . . .

- Both players play C in the first stage.
- In all subsequent stages, both players play D if D has ever previously been played by anybody. Otherwise, both play C.

This is a strategy profile as it suggests what each player should do in any possible scenario that can reached.

This is called a **grim-trigger strategy profile**.

- Note that, if followed, (C,C) will occur in each stage!
- But is it a Nash equilibrium?

Profitable deviations

Remember, in a Nash equilibrium, each player is best responding to the other.

- To check if a strategy is a best response, we can ask if a player has a profitable deviation. If not, it is a best response.
- The possible profitable deviation would be that a player would deviate to playing D instead of C.
  - If that increases his expected payoff, the strategies we have proposed are not best responses to eachother.

Does a player want to deviate to D? (1)

Suppose we're in stage one. Is a deviation to D profitable?

	C	D
C	2,2	0,3
D	3,0	1,1

- Expected payoff from not deviating (sticking with C):

$$2 + 2\delta + 2\delta^2 + 2\delta^3 + \ldots = \frac{2}{1-\delta}$$

- Expected payoff from deviating to D:

$$3 + \delta + \delta^2 + \delta^3 + \ldots = 3 + \delta \frac{1}{1-\delta}$$

Does a player want to deviate to D? (2)

A player **does not** want to deviate if:

$$\frac{2}{1-\delta} \geq 3 + \delta\frac{1}{1-\delta}$$
$$2 \geq 3(1-\delta) + \delta$$
$$2 \geq 3 - 2\delta$$
$$\delta \geq \frac{1}{2}$$

So a player does not want to deviate in the first stage so long as there is at least a 50% chance of another stage.

- By symmetry, this will be the same for the other player.

One stage down, infinitely more to check . . .

Just kidding. We're actually done. Why?

- We enter every stage and the future is identical to the future we perceived in the stage before (it's recursive!).
  - We're still supposed to play C.
  - There's still a probability δ of another stage.

Conclusions:

- The grim-trigger strategy profile is a Nash equilibrium of the infinitely repeated game.
- We finally have a model that predicts that cartels can sustain so long as it is sufficiently likely that they continue to interact (and there is no known end-date to their interaction).

Where are we now?

Useful problems

- Chapter 14: 5, 6