Lecture 13: Game Theory I: Static Games with Complete Information

Alexander Wolitzky

MIT

6.207/14.15: Networks, Spring 2022

Game Theory

Second half of the course considers social/economic networks where individuals make decisions **strategically**.

- This means that each individual's optimal choice depends on what others are doing.
- We already saw some examples: behavioral SIR model, Morris contagion model.
- More examples**:** choosing a route, adopting a new technology, pricing/bargaining/bidding, sharing or learning from a piece of information.

Such situations are formally modeled as **games**: that is, multi-person decision problems.

The analysis of games is called **game theory**.

Today: Static Games of Complete Information

Today's lecture provides an introduction to game theory, focused on the simplest type of games: **static games of complete information**.

- "Static" means game is played all at once, not over time.
- "Complete information" means there is no uncertainty or private information in the game.

Coming lectures cover network applications of this simple class of games.

We consider more complex games (dynamic, incomplete information) later in the course.

Games

A **game** consists of

1. A set $N = \{1, \ldots, n\}$ of **players**.
   - Who is playing the game.
   - In network models, often each player is a node in the network.
2. A **strategy set** S_i for each player $i \in N$.
   - The set of options available to each player.
   - For example, $S_i = \{0, 1\}$ means each player can choose action 0 or action 1.
3. A **payoff function** $u_i : S_1 \times \ldots \times S_n \to \mathbb{R}$ for each player $i \in N$.
   - Each player cares about what she does *and what everyone else does*.

Strategy Profiles

- We write $S = S_1 \times \ldots \times S_n$.
- An element $s = (s_1, \ldots, s_n) \in S$ is called a **strategy profile**.
  - Each player's payoff $u_i(s)$ depends on the entire strategy profile.
  - Ex. your payoff from using technology 1 depends on who else uses technology 1.
- We write $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ for a vector of strategies for everyone except player i, and we write $(s_i, s_{-i}) = (s_1, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n)$.
  - Since player i cannot control the other players' strategies, she takes s_{-i} as given, and chooses her own strategy s_i to maximize her payoff $u_i(s_i, s_{-i})$.
- The set of players other than player i is called **player i's opponents**.
  - This does not mean they try to "beat" player i. Every player tries to maximize her own payoff. This is not very restictive, because i's payoff function should already account for any altruism or spite that i feels towards j.

Example: The Prisoner's Dilemma

- Two criminals have been arrested and are held in separate cells.
- Each criminal can either *cooperate* with the other criminal by keeping quiet, or *defect* by helping the police.
- If both cooperate, they each get 1 year in prison.
- If both defect, they each get 2 years in prison.
- If one cooperates and the other defects, the one who cooperates gets 3 years in prison, and the one who defects gets 0 years.

The Prisoner's Dilemma (cntd.)

With the normalization

$$\text{payoff } = 3 - (\text{number of years in prison}),$$

the game is given by

- $N = \{1, 2\}$
- $S_1 = S_2 = \{C, D\}$
- $u_1(C, C) = u_2(C, C) = 2$
- $u_1(D, D) = u_2(D, D) = 1$
- $u_1(C, D) = u_2(D, C) = 0$
- $u_1(D, C) = u_2(C, D) = 3$

Payoff Matrix Representation

In 2-player games with a small number of strategies for each player, we can represent the game as a **payoff matrix** with player 1's strategies in the rows, player 2's strategies in the columns, and the payoff pairs (u_1, u_2) filling out the matrix.

For the prisoner's dilemma, the payoff matrix is

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

More Examples

Two children in a house must go *Upstairs* or *Downstairs*

- If they like being together, we have a **coordination game,** with payoff matrix

	U	D
U	1,1	0,0
D	0,0	1,1

- If they like being apart, we have an **anti-coordination game**, with payoff matrix

	U	D
U	0,0	1,1
D	1,1	0,0

- If player 1 likes being with player 2 but player 2 likes avoiding player 1, we have a **hide-and-seek game** (also called **matching pennies**), with payoff matrix

	U	D
U	1,0	0,1
D	0,1	1,0

Solution Concepts for Games

What do we expect to happen when a game is played?

- Related question: how should it be played?

An answer to these questions is a prediction/prescription about how games will/should be played.

- This is called a **solution concept**.
- There are different solution concepts. We will focus on the important and most widely applied ones. A game theory course would spend more time discussing alternatives.

Dominant/Dominated Strategies

Almost all solution concepts predict that, if a player has a strategy that's *always* uniquely optimal (no matter what her opponents play), she will play that strategy.

- Similarly, if a player has a strategy that's never optimal, she will not play that strategy.

A strategy $s_i \in S_i$ for player i is **strictly dominant** if, for every alternative strategy $s_i' \in S_i$, we have $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$.

- A player should take a strictly dominant strategy regardless of what others do.

A strategy $s_i \in S_i$ for player i is **strictly dominated** if there exists an alternative strategy $s_i' \in S_i$ such that $u_i(s_i, s_{-i}) < u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$.

- A player should never take a strictly dominated strategy regardless of what others do.
- If there exists a strictly dominant strategy, then every other strategy is strictly dominated by it.

Example: Prisoner's Dilemma

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

Is there a strictly dominant strategy? What is it?

- For each player, D is the unique optimal action regardless of what other do.
- Thus, game theory (or really just individual optimization) predicts that the outcome of the game will be (D, D).

- **Note:** if the players could jointly agree to play (C, C), they would both be better off.
- But (C, C) is not the game-theoretic prediction, because it is not a stable outcome: an agreement to play (C, C) would not be credible, because each player has an incentive deviate to D.
- The prisoner's dilemma thus demonstrates a conflict between individual optimization and social optimality.

Possible Objections to this Prediction

- What if the players dislike hurting each other, e.g. they are altruistic?
- Remember that a player's payoff is, by definition, what she maximizes.
- The entries in the payoff matrix must already take into account motivations such as altruism. They are not only "material payoffs."
- In real-world applications of game theory, we must sometimes be careful in thinking about how to account for non-material payoffs.

- Couldn't the players support the play of (C, C) via some threatened punishment if there is a deviation to D?
- We will see later in the course that this is indeed possible in richer settings, such as if the game is played repeatedly.
- But if we specify that the whole game is just the (static) prisoner's dilemma, such punishments are not possible.

Example: Coordination Game

	U	D
U	1,1	0,0
D	0,0	1,1

Is there a strictly dominant strategy? What is it?

- The prisoner's dilemma is unusual in having strictly dominant strategies.
- The more typical case (the one where game theory is really needed, as opposed to just individual optimization) is when a player's optimal action does depend on what everyone else does.
- What is our prediction/prescription for how the game will/should be played when no strategy is strictly dominant?

Nash Equilibrium

The simplest and most widely applied prediction is that the players will play a **Nash equilibrium**.

- This simply means that each player plays optimally, *taking as given* what everyone else is doing.

Formally, a **pure-strategy Nash Equilibrium (PSNE)** is a strategy profile $s^* \in S$ such that

$$u_i\left(s_i^*, s_{-i}^*\right) \geq u_i\left(s_i, s_{-i}^*\right) \text{ for all } s_i \in S_i, i \in N.$$

A PSNE is a stable outcome: no player has an incentive to unilaterally change her strategy (or "deviate").

- The words "pure-strategy" distinguish this definition from a more general version of NE that we'll see later in this lecture.
- If any player i has a strictly dominant strategy, she must play that strategy in every PSNE.

Best Responses

One more useful piece of terminology: given the opponents' strategies s_{-i}, a "best response" for player i is an optimal action against s_{-i}.

- Formally, s_i is a **best response** to s_{-i} if $u_i\left(s_i, s_{-i}\right) \geq u_i\left(s_i', s_{-i}\right)$ for all $s_i' \in S_i$.

Note that a PSNE is simply a strategy profile where each player is playing a best response.

- To check whether a strategy profile is a PSNE: for each player separately, hold fix the other players' strategies and see if that player is playing a best response.

Prisoner's Dilemma Again (cntd.)

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

We can also check that (D, D) in a PSNE directly from the definition: for each player, if the opponent is playing D, then D is a best response.

- To check this, we must check that for each player $u(D, D) \geq u(C, D)$.
- The fact that $u(D, C) > u(C, C)$ is relevant for showing that D is a strictly dominant strategy, but it is not relevant for showing that (D, D) is a PSNE.
  (Since for Nash equilibrium we only consider a player's incentives *holding everyone else's strategy fixed*.)

Coordination Game Again

What are the PSNE in the coordination game?

	U	D
U	1,1	0,0
D	0,0	1,1

- (U, U) is a PSNE: if 1 is playing U, 2's best response is U (and vice versa)
- (D, D) is also a PSNE
- (U, D) is not a PSNE: if 1 is playing U, 2's best response is U
- (D, U) is also not a PSNE

Comment: Morris Contagion Model

The Morris contagion model we studied last week is simply the coordination game played on a network, where each player wants to coordinate with as many of her neighbors as possible.

The "equilibrium" concept we studied was precisely the PSNE of this game.

Exercise: Go back to the lecture notes on the Morris model and make sure you understand this connection.

Hide-and-Seek

	U	D
U	1,0	0,1
D	0,1	1,0

What are the PSNE?

- (U, U) is not a PSNE: 2 will deviate to D
- (U, D) is not a PSNE: 1 will deviate to D
- (D, D) is not a PSNE: 2 will deviate to U
- (D, U) is not a PSNE: 1 will deviate to U

Uh-oh.

What Now?

	U	D
U	1,0	0,1
D	0,1	1,0

What do we expect to happen in hide-and-seek?

- The prediction that the players will play any one of the four pure-strategy profiles 100% of the time is not a good prediction, because none of these are stable outcomes.

What Now? (cntd.)

	U	D
U	1,0	0,1
D	0,1	1,0

Instead, let's try to predict what **fraction** of the time each player will play each action.

- A prediction that player 1 will play U strictly more than $1/2$ the time will not come true.
- If player 1 plays U more than $1/2$ the time, player 2 will play D 100% of the time.
- But then player 1 should play D 100% of the time, and never play U!
- Similarly, a prediction that player 1 will play D strictly more than $1/2$ the time will not come true either, and similarly for player 2.
- The unique stable prediction is that each player will play each action exactly 50% of the time.

What Now? (cntd.)

	U	D
U	1,0	0,1
D	0,1	1,0

Instead, let's try to predict what **fraction** of the time each player will play each action.

- A prediction that player 1 will play U strictly more than $1/2$ the time will not come true.
- If player 1 plays U more than $1/2$ the time, player 2 will play D 100% of the time.
- But then player 1 should play D 100% of the time, and never play U!
- Similarly, a prediction that player 1 will play D strictly more than $1/2$ the time will not come true either, and similarly for player 2.
- The unique stable prediction is that each player will play each action exactly 50% of the time.
- This is the unique "mixed-strategy Nash equilibrium".

Mixed Strategies

A **mixed strategy** σ_i for player i is a probability distribution over S_i.

- Simplest interpretation: player i flips a coin or spins a roulette wheel before choosing her action.
- There are also other, more realistic ways of interpreting mixed strategies. We will mention them later.

Let Σ_i denote i's set of mixed strategies, with $\Sigma = \Sigma_1 \times \ldots \times \Sigma_n$.

- It is assumed that any players who use mixed strategies must randomize independently.

A player i's payoff at a mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is simply her expected payoff under independent randomization:

$$\begin{aligned} u_i(\sigma) &= \sum_{s \in S} \Pr^{\sigma}(s)\, u_i(s) \\ &= \sum_{(s_1,\ldots,s_n) \in S} \left(\prod_{j=1}^{n} \sigma_j(s_j) \right) u_i(s_1, \ldots, s_n). \end{aligned}$$

Hide-and-Seek Again

	U	D
U	1,0	0,1
D	0,1	1,0

What are the payoffs at mixed strategy profile $\left(\sigma_1 = \left(\frac{1}{2}U, \frac{1}{2}D\right), \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$?

What are the payoffs at mixed strategy profile $\left(\sigma_1 = U, \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$?

What are the payoffs at mixed strategy profile $\left(\sigma_1 = \left(\frac{2}{3}U, \frac{1}{3}D\right), \sigma_2 = \left(\frac{1}{3}U, \frac{2}{3}D\right)\right)$?

Mixed-Strategy Nash Equilibrium

Recall that a **pure-strategy Nash Equilibrium (PSNE)** is a strategy profile $s^* \in S$ such that

$$u_i\left(s_i^*, s_{-i}^*\right) \geq u_i\left(s_i, s_{-i}^*\right) \text{ for all } s_i \in S_i, i \in N.$$

- A PSNE is a pure-strategy profile where every player is playing a best response.

Similarly, a **mixed-strategy Nash Equilibrium (NE)** is a strategy profile $\sigma^* \in \Sigma$ such that

$$u_i\left(\sigma_i^*, \sigma_{-i}^*\right) \geq u_i\left(\sigma_i, \sigma_{-i}^*\right) \text{ for all } \sigma_i \in \Sigma_i, i \in N.$$

- A NE is a (possibly mixed) strategy profile where every player is playing a best response.

Verifying Mixed NE

Even with only two pure strategies, there are infinitely many mixed strategies.

How can we check that none of them does better than σ_i^*?

Fortunately, it suffices to check that there is no **pure** strategy that does better than σ_i^*.

- If some mixed strategy σ_i is strictly better than σ_i^*, then at least one of the pure strategies in its support must also do strictly better than σ_i^*.
  (Otherwise, the expected payoff from σ_i could not be higher than that from σ_i^*.)

Theorem

σ^* is a NE if and only if

$$u_i\left(\sigma_i^*, \sigma_{-i}^*\right) \geq u_i\left(s_i, \sigma_{-i}^*\right) \text{ for all } s_i \in S_i, i \in N.$$

Hide-and-Seek Again

	U	D
U	1,0	0,1
D	0,1	1,0

Find **all** the NE (pure and mixed).

We know there aren't any pure-strategy NE.

Let's prove that $\left(\sigma_1 = \left(\frac{1}{2}U, \frac{1}{2}D\right), \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$ is a mixed-strategy NE.

- When one player plays $\left(\frac{1}{2}U, \frac{1}{2}D\right)$, the other player's expected payoff is $\frac{1}{2}$ for **any** strategy she might play.
- In particular, for each player, $\left(\frac{1}{2}U, \frac{1}{2}D\right)$ is a best response to $\left(\frac{1}{2}U, \frac{1}{2}D\right)$.

Hide-and-Seek (cntd.)

	U	D
U	1,0	0,1
D	0,1	1,0

Let's prove that $\left(\sigma_1 = \left(\frac{1}{2}U, \frac{1}{2}D\right), \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$ is the **only** mixed-strategy NE.

Suppose σ is a NE.
Let $\sigma_i(a)$ be the probability that player i takes action a.

- If $\sigma_1(U) > \sigma_1(D)$, then 2's unique best response is D, so (since σ is a NE) player 2 must be playing D.
  But then 1 would deviate to D.
- If $\sigma_1(U) < \sigma_1(D)$, then 2's unique best response is U, so player 2 must be playing U. But then 1 would deviate to U.
- Hence, in any NE we must have $\sigma_1(U) = \sigma_1(D)$, and symmetrically $\sigma_2(U) = \sigma_2(D)$.
- In other words, only $\left(\sigma_1 = \left(\frac{1}{2}U, \frac{1}{2}D\right), \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$ can be a NE.

Coordination Game Again

	U	D
U	$1, 1$	$0, 0$
D	$0, 0$	$1, 1$

We saw this game has two PSNE: (U, U) and (D, D).

It also has a mixed NE: $\left(\frac{1}{2}U + \frac{1}{2}D, \frac{1}{2}U + \frac{1}{2}D\right)$.

- This is a more compact way to write $\left(\sigma_1 = \left(\frac{1}{2}U, \frac{1}{2}D\right), \sigma_2 = \left(\frac{1}{2}U, \frac{1}{2}D\right)\right)$.

Exercise: Prove that these three are the only NE in the coordination game.

Challenging Exercise: Can there be mixed NE in the Morris model (i.e. coordination game on a network)? What do they look like?

A Theorem

A game is **finite** if each player's strategy set S_i is finite.

Theorem

Every finite game has a NE (allowing mixed NE).

This is called the **Nash existence theorem**.
(The theorem that won John Nash the Nobel Prize in Economics.)

- Proof is not too hard but beyond our scope.
- This is a very important theorem, as it ensures that NE is a usable solution concept for most models we consider.
- The theorem extends to some classes of infinite games but not all infinite games: for example, consider the game where each player names an integer, and the player who names the greatest integer wins \$1.

Summary

- We represent multi-person decision problems as games.
- The basic prediction/"solution" of a game is Nash equilibrium.
- A NE is a stable outcome of the game, where no player has a profitable deviation.
- Some games have only pure-strategy NE, some have only mixed-strategy NE, and some have both kinds.
- Every finite game has at least one NE.

14.15 / 6.207 Networks
Spring 2022