CSCI699: Topics in Learning and Game Theory

Lecture 3

Lecturer: Shaddin Dughmi *Scribes: Brendan Avent, Cheng Cheng*

1 Introduction to Game Theory

Game theory is the mathematical study of interaction among rational decision makers. The goal of game theory is to predict how agents behave in a game. For instance, poker, chess, and rock-paper-scissors are all forms of widely studied games.

To formally define the concepts in game theory, we use *Bayesian Decision Theory*. Explicitly:

- Ω is the set of future states. For example, in rock-paper-scissors, the future states can be 0 for a tie, 1 for a win, and -1 for a loss.
- A is the set of possible actions. For example, the hand form of rock, paper, scissors in the game of rock-paper-scissors.
- For each $a \in A$, there is a distribution $x(a) \in \Omega$ for which an agent believes he will receive $\omega \sim x(a)$ if he takes action a.
- A rational agent will choose an action according to Expected Utility theory; that is, each agent has their own utility function $u : \Omega \to R$, and chooses an action $a^* \in A$ that maximizes the expected utility .
  - Formally, $a^* \in \arg\max_{a \in A} \mathop{\mathbb{E}}_{\omega \sim x(a)} [u(\omega)]$.
  - If there are multiple actions that yield the same maximized expected utility, the agent may randomly choose among them.

2 Games of Complete Information

2.1 Normal Form Games

In games of complete information, players act simultaneously and each player's utility is determined by his actions as well as other players' actions. The payoff structure of the game (i.e., the map from action profiles to utility vectors) is common knowledge to all players in the game.

For instance, in the game of rock-paper-scissors, the payoff structure is a 3 by 3 matrix. The rows are indexed by the three possible actions (rock, paper, and scissors) of player 1, and the columns are indexed by the three possible actions of player 2. Each cell of the matrix gives the utility for both players as a result of their joint actions; e.g., if player 1 chooses the action of scissors and player 2 chooses the action of rock, than the result in given cell of the matrix may be -1 for player 1 and +1 for player 2.

The following defines the typical mathematical representation of such games known as the *normal form*:

- A set of players $N = \{1, \dots, n\}$
- For each player i, a set of actions A_i
  - Let $A = A_1 \times A_2 \times \dots \times A_n$ denote the set of action profiles.
- For each player i, a utility function $u_i : A \to \mathbb{R}$.
  - If players play $a_1, \dots, a_n$, then $u_i(a_1, \dots, a_n)$ is the utility of player i. This is typically thought of as an n-dimensional matrix, indexed by $a = (a_1, \dots, a_n) \in A$, with entry $(u_1(a), \dots, u_n(a))$.

2.2 Strategies in Normal Form Games

A strategy of a player is the way he choose an action.

- A *pure strategy* of player i is simply an action $a_i \in A_i$.
- A *mixed strategy* of player i is a distribution s_i supported on A_i. The player randomly draws an action $a_i \sim s_i$. An example is a player uniformly randomly choosing to play either rock, paper, or scissors.
- A *strategy profile* is (s1, . . . , sn). When each player chooses mixed strategy, correlation matters.
  - In most cases we discuss: independent. (Nash equilibrium)
  - Sometimes, there are ways of correlating (coordinating) their random choice (Correlated equilibrium).

- When we refer to mixed strategy profiles, we mean independent randomization unless stated otherwise.
- We can extend utility functions to mixed strategy profiles. Using $s(a)$ as shorthand for the probability of action a in strategy s,

$$u_i(s_1, \ldots, s_n) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j).$$

- Best Response: A strategy of a player that gives the best utility given the strategy of other players.
  - A mixed s_i of player i is a best response to a strategy profile s_{-i} of the other players if $u_i(s) \geq u_i(s'_i, s_{-i})$ for every other mixed strategy s'_i.
  - Note: There is always a pure best response.
  - The set of mixed best responses is the randomization over pure best responses.
  - Examples
    * Prisoner's Dilemma (slides p8): If prisoner 2 chooses to cooperate, the best response for prisoner 1 it to defect since the utility is larger compared to not $(0 > -1)$. If prisoner 2 chooses to defect, the best response for player 1 is still to defect since the utility is larger compared to not $(-2 > -3)$
    * The Battle of the Sexes (slides p9): If the woman chooses football, the best response for the man is to choose football since the utility is larger than not $(2 > 0)$. If the woman chooses movie, the best response for the man is to choose movie since the utility is larger than not $(1 > 0)$.

2.3 Sequential Games

These won't be examined much in class, but it's good to know they exist. In sequential games, the game is played over time (i.e., not simultaneously); e.g., chess, video game, poker.

- More naturally modeled using the extensive form tree representation rather than normal form.
- Each non-leaf node is a step in the game and associated with a player.

- Outgoing edges represent the actions available at that step.
- Leaf nodes are labeled with utility of each player.
- A pure strategy here is a choice of action for *each* contingency (i.e., for each non-leaf node).
- In terms of equilibrium (as defined in next section) for such games, subgame-perfect equilibrium is the most pertinent. Sits between DSE and Nash.

2.4 Equilibrium

An equilibrium concept identifies, for every game, one or more distributions over action profiles (the equilibria). Game theory predicts that the outcome of the game is distributed as one of the equilibria.

The following four types of equilibria are defined from the most restricted to the most relaxed form.

2.4.1 Dominant Strategy Equilibrium

A strategy s_i of player i is a dominant strategy if it is a best response to every strategy profile s_{-i} of the other players. Formally, for all profiles s_{-i} of players other than i, we should have that $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for any other strategy s'_i of player i.

- A *dominant-strategy equilibrium* (DSE) is a strategy profile where each player plays a dominant strategy.
- DSE is the most restricted form of equilibrium. If one exists, i doesn't need to know what others are doing in order to best respond.
- This is the best kind of equilibrium, since it requires minimal knowledge assumptions.
- If there is a mixed dominant strategy, there is also a pure one (since mixed DS is just randomization over pure DS).
- It may be pure or mixed (independent randomization). Mixed DSE are not of much interest since a pure DSE will always exists when mixed D.
- Every DSE is also a Nash Equilibrium.

- In the example of Prisoner's dilemma, an DSE is when both players defect; if either chooses to cooperate, regardless of what the opponent chooses, his utility can be increased by defecting. Therefore, both players will choose to defect no matter what the other side chooses.

2.4.2 Nash Equilibrium

A Nash equilibrium is a strategy profile $(s_1, ..., s_n)$ such that, for each player i, s_i is a best response to s_{-i}.

- If each s_i is pure, we call it a pure Nash equilibrium; otherwise, we call it a mixed Nash equilibrium.
- All players are optimally responding to each other, simultaneously.
- Unlike DSE, there may exist a mixed Nash Equilibrium but no pure Nash equilibrium (e.g., rock-paper-scissors).
- Pure Nash equilibria and dominant strategy equilibria do not always exist. However, mixed Nash equilibrium always exists when there is a finite number of players and actions!
- Every Nash equilibrium is a correlated equilibrium.
- For the Battle of the Sexes game, both pure and mixed Nash equilibria exists.
  - Pure: two cases, when both the man and the woman choose football or when both the man and the woman choose movie. Since for both sexes, if they know their partner's option, their best response is to do the same thing in order to maximize utility. For instance, if the man realizes the woman chose football, he will gain more utility for choosing football than for choosing movie (2 vs. 0). He also realizes that if the woman chooses movie, he will gain more utility for choosing movie (1 vs. 0).
  - Mixed: The man chooses football w.p. $\frac{2}{3}$ and movie w.p. $\frac{1}{3}$, and W chooses movie $\frac{2}{3}$ and football w.p. $\frac{1}{3}$. If both the man and woman know the other's strategy, the expected utility for both is the same:

$$Man: \quad 2 \times \frac{1}{3} + 0 \times \frac{2}{3} = \frac{2}{3}$$

$$Woman: 0 \times \frac{2}{3} + 2 \times \frac{1}{3} = \frac{2}{3}$$

By randomizing, the man makes the woman indifferent and vice-versa, so that a mixed Nash Equilibrium is achieved.

2.4.3 Correlated Equilibrium

A correlated equilibrium can be thought as having an extra player in the game that gives "fair" and randomized recommendations to the remaining players.

A correlated equilibrium is distribution x over action profiles such that, for each player i and action $a_i^\star \in A_i$, we have

$$\mathop{\mathbb{E}}_{a\sim x}[u_i(a)|a_i = a_i^*] \geq \mathop{\mathbb{E}}_{a\sim x}[(a_i', a_{-i})|a_i = a_i^*]$$

for all $a_i' \in A_i$.

- In other words: if all players other than i follow their recommendations, then when i is recommended $a_i^\star$, his posterior payoff is maximized by following his recommendation as well.
- In Battle of the Sexes, the correlated equilibrium is to randomize between (Football, Football) and (Movie, Movie). For both the man and woman, think of it as if there is a mediator that gives recommendations based on the result of a fair coin toss. For example, if the coin shows heads, both the man and woman are told to choose football; if the coin shows tails, both the man and woman are told to choose movie. Since both players know that the other side is told the same thing, they both don't have incentive to select another choice which may lower their utility. In this case, their expected utility is still the same:

$$Man: \quad 2 \times \frac{1}{2} + 1 \times \frac{1}{2} = \frac{3}{2}$$

$$Woman: 1 \times \frac{1}{2} + 2 \times \frac{1}{2} = \frac{3}{2}$$

- In the Chicken Game, a correlated equilibrium can be randomizing among the Nash equilibria (STOP, GO) and (GO, STOP) as in Battle of the Sexes. Another would be uniformly randomizing between (STOP,GO), (GO,STOP), and (STOP,STOP).
  - To prove the latter case, we can assume that a fair mediator is will give recommendation to the players with a probablity of $\frac{1}{3}$ for each of the cases, but only given what he is supposed to do.

– If the recommendation to Player 2 is GO, he will have no incentive to disobey since he knows that this is the only option given that Player 1 has to choose STOP. If the recommendation to Player 2 is STOP, he knows the only two possible result if he choose STOP is (STOP, STOP) and (GO, STOP), the expected utility OS:

$$STOP : 0 \times \frac{1}{2} + -2 \times \frac{1}{2} = 0.5$$

But if he does not follow the recommendation and choose GO, the result is:

$$GO : -2 \times \frac{1}{2} + -10 \times \frac{1}{2} = -6$$

He will gain more utility if he selects to cooperate ($0.5 > -6$).

2.4.4 Coarse-correlated Equilibrium

The most relaxed form.

3 Games of Incomplete Information

A more realistic setting for games is one where players don't have perfect knowledge of the payoff structure of the game or the private information of others. Without this information though, it is generally impossible for a player to play a Nash equilibrium, since it's unclear exactly what the game is!

As an example, consider a first-price auction for a car. Recall that in a first-price auction, the highest bidder receives the item and pays an amount equal to their bid; all other players pay nothing and receive nothing. We model the utility for the winner of the car as simply the difference in their value for the car and the amount they paid (and 0 for everyone else). Player 1's value for the car is $v_1 = \$3,000$, and player 2's value v_2 is either $\$1,000$ or $\$2,000$. If both players know the other player's exact valuation, then a Nash equilibrium is simply a bid of $b_1 = v_2 + \epsilon$ from Player 1 and a bid of $b_2 = v_2$ for player 2, where ϵ is the smallest positive biddable unit. But this requires Player 1 to know exactly Player 2's valuation – it is unclear what the equilibrium bids are, or even how they are defined, when their mutual valuations aren't perfectly known. To tackle these issues, *games of incomplete information* were defined.

3.1 Prior-free Games

Prior-free games are those in which players only knows what the possible values of everyone else's private data are, without knowing or assuming anything about the likelihoods of these values and, in turn, their strategies. To formalize, we define that a *game of strict incomplete information* is given by

- A set of players $N = \{1, \ldots, n\}$.
- For each player i, a set of *actions* A_i.
  - The set of *action profiles* are $A = A_1 \times \ldots \times A_n$.
- For each player i, a set of *types* T_i.
  - The set of *type profiles* are $T = T_1 \times \ldots \times T_n$.
- For each player i, a utility function $u_i : T \times A \to \mathbb{R}$.

For this course, we're interested in a slightly more restricted definition, where the utility functions of each player aren't directly dependent on the *types* of the other players, but just of the player's own type in conjunction with the *actions* of the other players. That is to say, each player only cares about the other players insofar as the actions they take, not as the type they are (since that merely influences what actions they will take). Formally, this means that the utility function is modeled as $u_i : T_i \times A \to \mathbb{R}$. Concretely $u_i(t_i, \vec{a})$ is the utility of i when i has type t_i and players play $\vec{a} = (a_1, \ldots, a_n)$. Notice that when each player only has one type, this reduces to a complete information game.

As with complete information games, we have two types of strategies: pure and mixed. Here though, they have more generalized definitions.

- A *pure strategy* for player i is $s_i : T_i \to A_i$ – a choice of action $a_i \in A_i$ for every type $t_i \in T_i$.
  - Example: Truthful bidding (i.e., bidding your valuation).
  - Example: Bidding half your valuation.
- A *mixed strategy* for player i is a distribution $s_i(t_i)$ over actions A_i for each type $t_i \in T_i$.
  - Example: Bidding b_i uniformly in $[0, v_i]$.

Notice that here, we have strategies defined as only depending on a player's own private info (i.e., their type), NOT on other players' private info or actions. A *strategy profile* $(s_1, \ldots, s_n)$ describes what would happen in each "state of the world" (i.e., type profile) $(t_1, \ldots, t_n)$, regardless of the relative frequency of various states. Concretely, in state $(t_1, \ldots, t_n)$, players play $(s_1(t_1), \ldots, s_n(t_n))$.

For prior-free games, the main equilibrium concept is the *dominant strategy equilibrium* (DSE), a generalization from the complete information context. Here, we say that a strategy $s_i : T_i \to \Delta(A_i)$ is a *dominant strategy* (DS) for player i if $\forall\ t_i \in T_i$ and $a_{-i} \in A_{-i}$ and $a'_i \in A_i$, we have that

$$u_i(t_i, (s_i(t_i), a_{-i})) \geq u_i(t_i, (a'_i, a_{-i})).$$

That is, $s_i(t_i)$ is a best response to $a_{-i} \sim s_{-i}(t_{-i})$ for all t_{-i}, s_{-i}, and realization a_{-i}. To choose the best response, player i only needs to know t_i, not any of the other players' types or their strategies. Similar to complete information games, if there is a mixed dominant strategy, then there is a pure dominant strategy.

A *dominant strategy equilibrium* here is still simply a strategy profile where each player plays a dominant strategy. Without information about the relative frequency of type profiles (e.g., prior probabilities over players' types), a DSE is the only reasonable equilibrium concept for strictly incomplete information games. The pros and cons of this DSE definition are essentially the same as for the DSE definition for complete information games; reference those for more info.

As an example of a DSE in this setting, consider a Vickrey (i.e., second-price) auction for a car. Recall that a second-price auction works by having the player with the highest bid winning the item, but only paying an amount equal to what the second highest bidder had bid; all other players pay nothing and receive nothing. Again, we model the utility for the winner of the car as simply the difference in their value for the car and the amount they paid (and 0 for everyone else). We claim that each player truthfully bidding exactly their own valuations is a DSE. While we leave the details of this as an exercise to prove on your own, we will start you out: with player i's valuation of the car as v_i and their submitted bid as b_i, consider the following cases: $b_i > v_i$, $b_i = v_i$, and $b_i < v_i$.

3.2 Bayesian Games

Prior-free information games, while having the positive of requiring minimal assumptions, comes with the consequence of being highly restrictive – in reality, some knowledge is likely known about the private data of other players. To remedy this, we consider a different way to model uncertainty: using a Bayesian common prior. Specifically, in this setting, we assume that players' private data is drawn from a distribution that is common knowledge. Players explicitly know their own private data, but unlike with the prior-free approach, they also know what distributions the other players' data are coming from.

To formalize, an n-player *Bayesian game of incomplete information* is given by

- A set of players $N = \{1, \ldots, n\}$.
- For each player i, a set of *actions* A_i.
  - The set of *action profiles* are $A = A_1 \times \ldots \times A_n$.
- For each player i, a set of *types* T_i.
  - The set of *type profiles* are $T = T_1 \times \ldots \times T_n$.
- For each player i, a utility function $u_i : T \times A \to \mathbb{R}$.
- A *common prior*, which is simply a distribution $\mathcal{D}$ over T.

Again, we're going to focus on the more restricted independent private values model, where the utility function for each player i is modeled as $u_i : T_i \times A \to \mathbb{R}$.

As an example of a Bayesian game, we'll look at a more advanced version of the previously-discussed first-price auction for a car. Here, for each player i, we define the set of actions as $A_i = [0, 1]$ and the set of types as $T_i = [0, 1]$; that is, player i will have a type (i.e., a valuation for the car) between 0 and 1, and will also be able to bid any amount between 0 and 1. We'll define utility the same as before, with $u_i(v_i, b) = v_i - b_i$ if $b_i > b_{-i}$ and 0 if $b_i < b_{-i}$ (ignoring ties). Our common prior here will be $\mathcal{D}$, which has each v_i being drawn uniformly from T_i (and independently across all players).

Before we continue with this game, we'll discuss the main equilibrium concept in this setting: the Bayes-Nash Equilibrium (BNE). As before, a pure strategy s_i for player i is a map from T_i to A_i. A pure *Bayes-Nash Equilibrium* of a Bayesian game of incomplete information is a set of strategies $\{s_1, \ldots, s_n\}$ where $s_i : T_i \to A_i$ such that for all $i, t_i \in T_i, a_i' \in A_i$ we have

$$\mathop{\mathbb{E}}_{t_{-i} \sim D | t_i} u_i(t_i, s(t)) \geq \mathop{\mathbb{E}}_{t_{-i} \sim D | t_i} u_i(t_i, (a_i', s_{-i}(t_{-i})))$$

where the expectation is over t_{-i} drawn from $\mathcal{D}$ after conditioning on t_i.

We can define a mixed BNE analogously, by allowing $s_i : T_i \to \Delta(A_i)$ and including this additional randomness in the expectations accordingly. It should be noted that every DSE is also a BNE; unlike the DSE though, a mixed BNE is always guaranteed to exist.

Jumping back to the first-price car aution now, assume that there are only 2 players ($n = 2$) and that for the case of a tie in bids, $u_i(v_i, b) = \frac{1}{2}(v_i - b_i)$. We leave it as an exercise here to show that in this game, the strategies $b_i = v_i/2$ for both players is a pure BNE.