

Lecture 8: Game Theory

Week 9

UCSB 2015

If you've seen *A Beautiful Mind*, you've heard of John Nash — one of the brightest mathematicians of our time, who revolutionized the mathematical field of game theory before being diagnosed with paranoid schizophrenia in 1959. His consequent battle with mental illness and later work — both in mathematics, where he has made massive contributions to multiple fields and as a mental health advocate — is a fantastic story, and one worth learning.

He passed away this weekend, at the age of 86 in a taxi accident. To honor his memory, it seems appropriate to spend some time discussing his work in the mathematical field of **game theory!**

1 Mathematical Games

A “game,” in the way that game theorists think about games, is not what most people mean by “games.” Rather, game theory is better thought of as the field of “strategic decision-making,” and games themselves are instances where a collection of players need to make individual decisions that interact with each other. Formally, we can define this as follows:

Definition. A **game** for n players $\{p_1, \dots, p_n\}$ consists of the following pieces of information:

- For each player p_i , we associate a list of possible **moves** or **strategies** S_i ; these are the actions available to that player.
- We also have a “payoff” function f , that given strategies $s_i \in S_i$ for all of our players, outputs the results of this collective action for each individual player. If we assume that we can represent these payoffs by real numbers, we can think of f as a function from $\prod_{i=1}^n S_i$ to \mathbb{R}^n .

A game is called **non-cooperative** if each player makes their decisions independently; i.e. while other players can certainly communicate or otherwise interact before making their decisions, there is no way for players to make outside “contracts” or other such things that restrict their decisions. Every game we consider will be assumed to be non-cooperative.

Looking at some examples might help make this more concrete:

1.1 Example Games

Game. (The stag hunt.) Two hunters have set off to get food from the forest. They have two possible game that they can hunt: either a stag, or rabbits.

To bring down a stag, both of the hunters need to cooperate, as it is too fierce to hunt alone. Rabbits, by comparison, can be hunted on one's own with no issue; however, they offer far less game than the stag.

We can describe this as a game for hunters h_1, h_2 as follows:

- Strategies: both players have the same pair of available strategies: { stag, rabbit}.
- Payoff: if both players hunt the stag, then they succeed in bringing it down; this gives them both a payoff of 2 units of food. If one player hunts the stag and the other hunts rabbits, then the stag hunter will fail while the rabbit hunter will succeed; i.e. the stag hunter gets a payoff of 0 while the rabbit hunter gets a payoff of 1. Finally, if both hunters pursue rabbits, they both receive a payoff of 1.

It is convenient to visualize such two-player games with a **payoff matrix**, where we identify the columns with player one's available moves, the rows with player two's moves, and fill entry (i, j) with the resulting pair of payoffs given those two moves:

		Hunter 1	
		Stag	Rabbit
Hunter 2	Stag	2, 2	0, 1
	Rabbit	1, 0	1, 1

Game. (Awkward social interactions.) Two people are walking down a hallway from opposite ends. The hallway is about two people wide. Both people are initially walking down the middle of the hallway; to continue walking, each must choose a side to go towards.

This too can be described in our game-theoretic language:

- Strategies: both walkers have the same strategies {left side, right side}.
- Payoffs: if both walkers choose their left, then they walkers walk past each other, and receive a payoff of one coolness point. Similarly, both players receive the same payoff if they both choose their right.

Conversely, if one chooses their left and the other chooses their right, then both walkers awkwardly stop and look at each other. This is clearly the worst thing ever and has a payoff of -1 coolness points.

Again, we can visualize this with a matrix:

		Walker 1	
		Left	Right
Walker 2	Left	1, 1	-1, -1
	Right	-1, -1	1, 1

Game. (Odds and evens.) Two schoolchildren are playing the game of “odds and evens,” which goes as follows: simultaneously, both players raise some number of fingers on their hand. If the sum is odd, then the first player wins; otherwise, if the sum is even the second player wins. We can turn this actual game into a game-theoretic game as follows:

- Strategies: both children have the same two choices, { odd, even}, if we make the observation that the number of fingers they raise is irrelevant; all we care about is the parity.
- Payoffs: Player 1 gets a payoff of 1 if the two children make different choices, while player 2 gets a payoff of 1 if the two children make the same choice.

		Player 1	
		Odd	Even
Player 2	Odd	0	1
	Even	1	0

Game. (The Prisoner’s Dilemma.) Two people have been captured by an oppressive dictatorship. Simultaneously, in different cells, both are met by their interrogator and asked to rat out the other prisoner for their alleged crimes.

If one prisoner betrays the other while the other remains silent, the silent prisoner is jailed for three years (one year for their silence, and two years for their crimes) while the betraying prisoner gets to go free. If they both betray each other, then they both get two years in prison for their crimes; finally, if they both stay silent they both get the lighter punishment of one year each (as the state has no evidence.)

We can again visualize this with a matrix, where this time the goal of our players is to minimize their payoff:

		Prisoner 1	
		Silence	Betray
Prisoner 2	Silence	1	0
	Betray	3	2

Not all games are two-player games; in these cases, we can’t draw a payoff matrix in the same way, but we can still study them:

Game. (Traffic.) Suppose you have n cars trying to travel from point A to point D on the map below:




Image from Wikipedia's [Nash Equilibrium](#) article.

Basically, we can think of the paths $A \rightarrow C$, $B \rightarrow D$, as large but indirect superhighways; no matter how much traffic goes on those two roads, it takes a car two hours to travel from start to finish. $B \rightarrow C$ is also such a highway road, though much shorter. Finally, $A \rightarrow B$, $C \rightarrow D$ can be best thought of as a pair of direct country roads; if there is no traffic on these roads, it takes a car a hour to travel from start to finish, but as traffic gets added to our system it takes longer and longer to use such a road!

We can think of this as a n -player game, where each player has one of the three choices $\{A \rightarrow B \rightarrow D, A \rightarrow C \rightarrow D, A \rightarrow B \rightarrow C \rightarrow D\}$ available to them, and the “payoff” for each player is the length of their commute; as before, players want to minimize their payoffs in this game.

1.2 Nash equilibria

In all of these games, it’s natural to wonder if there is an “optimal strategy” for our players, in some appropriate sense of “optimal.” Studying and answering this question is where Nash made his impact on the field of game theory, with the concept of **Nash equilibria**.

Intuitively speaking, suppose that you’re playing a game of the form above, and after all of the players have selected their move you are told the moves that all of the other players will make. If you would change your move, then this set of moves is not a Nash equilibria. However, if you would not change your move, and this property holds for **every** player in our game, then this set of moves is “stable:” no player can do better by unilaterally changing their strategy!

We call such a state a **Nash equilibria**. Formally, we define this as follows:

Definition. Suppose that G is a game with players p_1, \dots, p_n each with strategy sets S_i , and that we have an associated payoff function f for our game.

Suppose that $(s_1, s_2, \dots, s_n) \in \prod_{i=1}^n S_i$ is a selection of moves, one for each player, that is “stable.” That is: for any player i and any alternate move s'_i they could make, if we let f_i denote the payoff for player i , then

$$f_i(s_1, \dots, s_i, \dots, s_n) \geq f_i(s_1, \dots, s'_i, \dots, s_n).$$

We call such a state a **pure Nash equilibria**.

Several of the games we’ve studied have pure Nash equilibria:

- The stag hunt has two pure Nash equilibria: (rabbit, rabbit) and (stag, stag). If both hunters have chosen rabbit, then neither would unilaterally choose to switch to stag, as they would believe the other would stay with the rabbit and thus that they would fail in their hunt. As well, if both have chosen stag, then both would prefer to not change, as they would receive a strictly smaller payoff by switching to rabbit.

Conversely, these are the only two pure Nash equilibria for this game; if the hunters made mixed choices, then both would want to switch knowing the other hunter's choice: the stag hunter would want to switch to rabbit if they believed the other wanted to choose rabbit, while the rabbit hunter would want to choose stag if they knew they would have a partner!

- The two people walking in a hallway have (*left, left*) and (*right, right*) as their pure Nash equilibria, as in both cases neither person would want to switch knowing the other player's preferences. Again, these are the only two equilibria: if we were in either of the mixed states, both players would want to switch (thus leading to yet another conflict, and the resulting awkwardness.)

In the two games above, it seems natural to expect our players to converge onto one of these Nash equilibria. There are many conditions that (if satisfied) will insure that all of our players adopt a Nash equilibria. Here is one simple set of three conditions, that if all satisfied insure that a game with a single Nash equilibria always has its players adopt such a strategy:

- All players know the planned strategies of all of the other players, and believe that changing their strategy won't cause another player to change their strategy as well.
- All players want to optimize their payoff above all other concerns.
- All players are capable of perfect play; i.e. they don't make logical mistakes, they don't make errors in execution when performing their moves, and moreover they are all aware that the other players satisfy all of these properties and their consequences.

There are certainly instances of games in which some of these properties are not satisfied:

- Many times, the payoff function does not adequately describe the actual goals/desires of the players in the game; i.e. in the hunter scenario, perhaps the two hunters are bitter enemies, and would far prefer to not work together rather than cooperate. In this case, the value of the stag is not actually what the payoff matrix indicates, and it is entirely possible that (stag, stag) is no longer a pure Nash equilibrium.
- Other times, players are incapable of perfect play; consider playing go fish with a small child. As well, people can arrive at non-Nash equilibria if they do not believe that their opponent is playing perfectly; i.e. in the Cold War, countries often adopted strategies predicated on their opponents being irrational.
- Finally, as in the hallway scenario, it is often the case that other players are unaware of your strategy, or believe that changing their strategy will cause you to change yours as well (causing an infinite loop of shuffling left and right.)

However, in many cases these properties all hold (or approximately hold; when they do, researchers have found that in many cases populations converge on these Nash equilibria! This leads to some interesting paradoxes, which we can see if we look at our other games:

- The prisoner’s dilemma has only one Nash equilibria: both prisoners betraying each other. This is perhaps surprising, because it is clearly to our prisoner’s overall benefit for them to both stay silent.

However, if we assume our prisoners actually have the payoff functions described in their matrix (i.e. we don’t assume our prisoners have some extra secret cost associated to the guilt of betraying someone else, or anything else like this), then **if** a prisoner knew that the other person would stay silent, then they would benefit by switching their action to betrayal! So both prisoners staying silent is unstable; as well, any mixed state is clearly unstable, as the silent prisoner would change their action to betrayal if they knew that the other prisoner was betraying them as well.

The prisoner’s dilemma is arguably one of the most famous games studied in game theory, largely because of its paradoxical behavior; it seems to argue that prisoners should betray each other, while in reality there are **copious instances** of animals or people that co-operate to achieve a shared good even when defecting would lead to that individual’s own profit.

One argument here to explain the prisoner’s dilemma is that the Nash equilibria don’t quite apply; i.e. prisoners may fear retaliation in response to snitching on each other, which represents a modification to our payoff matrix that may not be there. However, even in situations where we have an accurate payoff matrix, individuals and animals seem to routinely adopt a more co-operative strategy than one might expect.

A second explanation for this situation here is that many encounters in life are **repeated**; i.e. if you are placed in such a situation with another person, you are likely to re-encounter this person later in life and thereby encounter similar situations in the future. This motivates a study of something called the **iterated prisoner’s dilemma**, which roughly goes as follows: suppose that you have two players who run through the prisoner’s dilemma game n times in a row. In theory, because you encounter the same player many times in a row, you are motivated to co-operate more often than normal, as the other player can “punish” you for your betrayal in later rounds!

If the two prisoners are aware of the value of n at the start, then (assuming optimal play / both players only care about maximizing their payoffs /etc) it is not hard to show by induction that they betray each other n times in a row:

- At the last step in this game, the two players are playing only one game of the prisoner’s dilemma, and have no later interactions with each other — i.e. they cannot punish each other for betrayal, and the only concern they have in theory is the payoff matrix. As shown above, this has exactly one Nash equilibria, in which both players betray each other.
- Both players are capable of perfect play, and therefore realize that at the last step the other will betray them. Therefore, on the second-to-last game they

again have no incentive to co-operate; so they again would betray each other with perfect play.

- The same logic then applies to the third-to-last game . . .
- . . . all the way back to the first game. So our players always betray each other.

However, if the two players are **not** aware of the number of games they will play against each other, things become much more interesting, and is an area of open research. In prisoner’s dilemma “tournaments” (in which various algorithms are written, paired up with each other, and played for a variable number of rounds, with the goal being to get the best total score over the tournament) top-scoring strategies typically had the following features:

1. Nice: they don’t defect before their opponent does.
2. Retaliatory: if their opponent betrays them, then they usually respond by betraying their opponent.
3. Forgiving: if their opponent starts to play nice in future rounds, then they too start to play nice.
4. Non-envious: their goals aren’t to outscore their specific opponent in that sequence of rounds.

A very simple strategy that (up to small modifications) meets these conditions is something called “tit-for-tat:”

1. The first move of a tit-for-tat player is to cooperate, i.e. stay silent.
2. At every subsequent move, the tit-for-tat player simply plays whatever their opponent played on the turn before.

Tit-for-tat or slightly similar strategies (“forgiving” tit-for-tat, that intermittently chooses to cooperate about $k\%$ of the time even if their opponent betrayed them on an earlier turn, for some small value of k) place very highly in pretty much all tournaments they’re entered in.

Interestingly enough, this is still an area of active research! In 2012, for example, a pair of researchers ([Press](#) and [Dyson](#)) came up with a new sort of strategies, called **zero-determinant** strategies, that players can adopt in a one-on-one match in certain scenarios to control their opponent’s score or enforce a set ratio between their own score and their opponent’s score. (These strategies apply in different scenarios than the ones that tit-for-tat operates in, i.e. they apply for someone facing a single fixed opponent who they assume is simply trying to maximize their own score, as opposed to a tournament with a wide variety of opponents that may not necessarily care about maximizing their own scores. So these results don’t invalidate the ideas above!)

If you want further reading,

https://golem.ph.utexas.edu/category/2012/07/zerodeterminant_strategies_in.html

is a really well-written discussion of Press and Dyson’s paper by a friend of mine, and

is Freeman and Dyson's own paper, which is remarkably readable (you need basically some linear algebra and patience!)

- The traffic system, as modeled above, also has a unique Nash equilibria that turns out to be collectively inefficient for all of our drivers. Suppose that we have a hundred drivers, of which 25 take the path $A \rightarrow B \rightarrow D$, another 25 take the path $A \rightarrow C \rightarrow D$, and the remaining 50 take the path $A \rightarrow B \rightarrow C \rightarrow D$.

There are 75 drivers on each of $A \rightarrow B, C \rightarrow D$, so the travel time of those two roads is 1.75, and consequently the travel time along any of our three paths is 3.75. So, in particular, no driver would benefit by switching; any $A \rightarrow B \rightarrow C \rightarrow D$ driver would not shorten their commute by taking one of the two-road paths, and any of the two-path drivers would actually increase the length of their commute.

However, this is clearly not optimal; if we simply had 50 drivers take the path $A \rightarrow B \rightarrow D$ and another 50 take $A \rightarrow C \rightarrow D$, every driver would have a commute of 3.5 hours and be home sooner! However, such a solution is not stable, as in any such commute all of our drivers would prefer to take the $A \rightarrow B \rightarrow C \rightarrow D$ route, as it would look like it would take a bit more than 3.25 hours.

This is known as Braess's paradox; as shown here, adding new connections between roads can sometimes make a commute take **longer**! This has actually been experimentally observed in real-life situations; see [this Wikipedia article](#) for a list of instances in which it's come up!

- Finally, the odds-evens game doesn't have a Nash equilibria at all! If both players have chosen the same option, player 1 will want to change; if they have chosen different options, player 2 will want to change.

It's a little frustrating that this last game doesn't have a Nash equilibrium state. We can fix this, though! Consider the following definition:

Definition. As before, let G be a game with players p_1, \dots, p_n each with strategy sets S_i , and let f be the associated payoff function for our game.

Now, instead of having each player submit a single move s_i for their strategy, we let each player submit a **probability distribution** of moves: i.e. if a player has moves $\{s_1, \dots, s_m\}$ available to them, we let them submit a probability vector (v_1, \dots, v_m) of nonnegative real values such that the sum $v_1 + \dots + v_m = 1$. Given this vector, we select a player's move at random by choosing option s_i with probability v_i . Notice that we can still simulate players deterministically choosing a fixed move by using the probability vector that is 1 on that move and 0 everywhere else.

These "randomized" strategies can be used to create a notion of a **mixed Nash equilibrium**, defined as follows:

Suppose that each player in our game has a probability vector \vec{v}_i that they use to randomly choose their strategy. Also suppose that these vectors are known to all of the other players.

If a given player, knowing everyone else's probability vector, could **change** their vector to get a greater payoff, then we consider this collection of strategies to be unstable. However, if no player can increase their payoff by changing their own vector while keeping the others constant, then we consider this state to be stable, and call it a **mixed Nash equilibria**.

Arguably Nash's greatest result in game theory was the following:

Theorem. Every game with a finite number of players and strategies has a mixed Nash equilibria.

We give a pair of examples of such equilibria here:

- In the people walking in a hallway problem, beyond the two pure Nash equilibria of (left, left) and (right, right), there is a mixed-Nash equilibria where each walker chooses randomly to go left or right with 50% probability for either choice; i.e. $((.5, .5), (.5, .5))$ is a mixed Nash equilibria. This is not hard to see: if we are in such a state and one of our walkers unilaterally changes their probability vector, they haven't changed their odds of avoiding a collision! So they do not benefit from this change, and presumably would stay the same.
- In the odds-evens game, that had no pure Nash equilibria beforehand, there is a mixed Nash equilibria: have both players play $((.5, .5), (.5, .5))$; i.e. have both players randomly choose odd or even! If they do this, then by the same reasoning as above no player would want to change their strategy upon knowing the other's, because they wouldn't improve their odds – no matter what you do, if your opponent is playing at random, your odds of success are 50%.

Cool!