Game Theory

2.1 Zero Sum Games

George Mason University, Spring 2018

Osborne and Rubinstein: "Game theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact. The basic assumptions that underlie the theory are that decision-makers pursue well-defined exogenous objectives (they are *rational*) and take into account their knowledge or expectations of *other* decision-makers' behavior (they *reason strategically*)."

Ex. Rock, Paper, & Scissors.

You are playing the game Rock, Paper, & Scissors. Paper beats rock. Rock beats scissors. Scissors beats paper. Two of the same is a draw. Which do you choose?

Ex. Stag Hunt.

You and an acquaintance have gone hunting and you must each decide whether to pursue a large stag or to pursue a small hare. Alone, you can each capture a hare. However, a stag requires two people to capture. If you cooperate and both hunt stag, then you will end up with much meat. But if you hunt stag while your acquaintance hunts hare, then you return home empty-handed. Do you hunt stag or hare?

We can explicate these informal interactive choice situations with formal mathematical models.

Def 2.1.1. A *strategic game* $\mathcal{G} = \langle \mathcal{N}, \{\mathcal{A}_i\}_{i \in \mathcal{N}}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of actions $\mathcal{A}_i$ available to each player $i \in \mathcal{N}$
- a set of outcomes $\mathcal{O}$
- a function $g : \times_{i \in \mathcal{N}} \mathcal{A}_i \to \mathcal{O}$ that maps each action profile $a \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ to an outcome $o \in \mathcal{O}$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$ where $o_1 \succcurlyeq_i o_2$ just in case o_1 is at least as preferred as o_2 by player $i \in \mathcal{N}$

Going forward, it is assumed that the preferences of each player $i \in \mathcal{N}$ are well-behaved in the sense that they are representable with an *interval utility function* $u_i : \mathcal{O} \to \mathbb{R}$.

Aside: We could model a situation where the outcomes of action profiles are affected by exogenous random variables by adding a set of states Ω.

The game $\mathcal{G}$ of Rock, Paper, & Scissors consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{A}_1 = \{\text{rock}, \text{paper}, \text{scissors}\}$
  $\mathcal{A}_2 = \{\text{rock}, \text{paper}, \text{scissors}\}$
- $\mathcal{O} = \{1 \text{ wins}, 2 \text{ wins}, \text{draw}\}$
- $g(\langle \text{paper}, \text{paper} \rangle) = \text{draw}$
  $g(\langle \text{paper}, \text{rock} \rangle) = 1 \text{ wins}$
  $g(\langle \text{paper}, \text{scissors} \rangle) = 2 \text{ wins}$
  $g(\langle \text{rock}, \text{paper} \rangle) = 2 \text{ wins}$
  $g(\langle \text{rock}, \text{rock} \rangle) = \text{draw}$
  $g(\langle \text{rock}, \text{scissors} \rangle) = 1 \text{ wins}$
  $g(\langle \text{scissors}, \text{paper} \rangle) = 1 \text{ wins}$
  $g(\langle \text{scissors}, \text{rock} \rangle) = 2 \text{ wins}$
  $g(\langle \text{scissors}, \text{scissors} \rangle) = \text{draw}$
- $1 \text{ wins} \succ_1 \text{draw} \succ_1 2 \text{ wins}$
  $2 \text{ wins} \succ_2 \text{draw} \succ_2 1 \text{ wins}$

The game $\mathcal{G}$ of Stag Hunt consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{A}_1 = \{\text{hunt stag}, \text{hunt hare}\}$
  $\mathcal{A}_2 = \{\text{hunt stag}, \text{hunt hare}\}$
- $\mathcal{O} = \{\text{both stag}, \text{both hare}, \text{1 hare only}, \text{2 hare only}\}$
- $g(\langle \text{hunt stag}, \text{hunt stag} \rangle) = \text{both stag}$
  $g(\langle \text{hunt stag}, \text{hunt hare} \rangle) = \text{2 hare only}$
  $g(\langle \text{hunt hare}, \text{hunt stag} \rangle) = \text{1 hare only}$
  $g(\langle \text{hunt hare}, \text{hunt hare} \rangle) = \text{both hare}$
- both stag $\succ_1$ both hare $\succcurlyeq_1$ 1 hare only $\succ_1$ 2 hare only
  both stag $\succ_2$ both hare $\succcurlyeq_2$ 2 hare only $\succ_2$ 1 hare only

A strategic game can be visualized in a table or in a game tree, but we'll focus on tables except when considering games with sequential play.

Generically, we refer to the Row player as Row (or Player 1) and the column player as Col (or Player 2). We can plug names in if we have them.

For example, the game $\mathcal{G}$ of Rock, Paper, & Scissors corresponds to the following *game matrix* (1 is the row player and 2 is the column player):

		Col (P2)		
		paper	rock	scissors
Row (P1)	paper	draw	1 wins	2 wins
	rock	2 wins	draw	1 wins
	scissors	1 wins	2 wins	draw

Just for reference, here's what Rock, Paper, & Scissors looks like as a game tree:

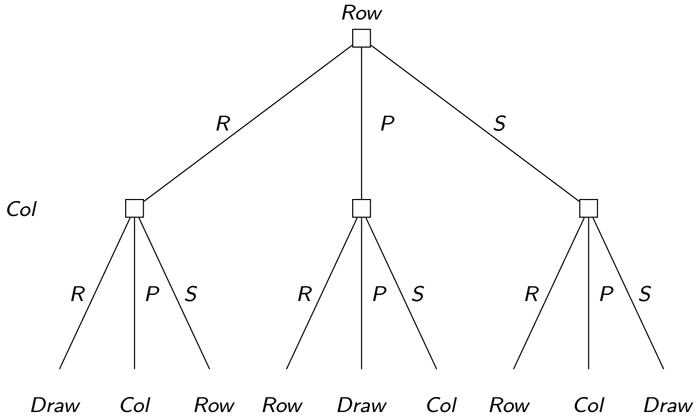

Chess: Game Tree

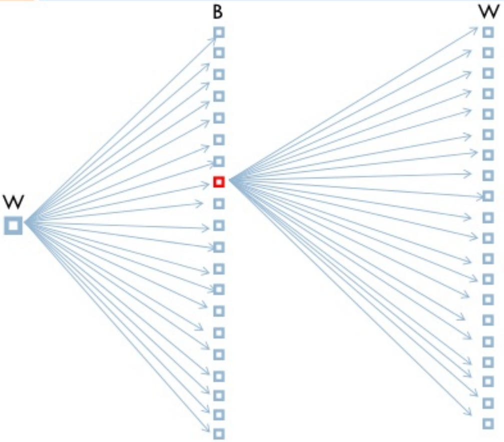

- So just TWO moves into Chess, and you already have 400 possible outcomes
- The problem is that no computer is powerful enough to calculate the solution
- It is estimated that there are 10^120 total number of possible moves in Chess.
- The estimated number of atoms in the universe are 10^80. A supercomputer would take 3x10^103 years to check out all these moves (FYI: the Sun is predicted to die by then).

Substituting the pair of utilities $\langle u_1(o), u_2(o) \rangle$ for each outcome o in the matrix gives us:

		Col		
		paper	rock	scissors
	paper	0,0	1,-1	-1,1
Row	rock	-1,1	0,0	1,-1
	scissors	1,-1	-1,1	0,0

The first utility in each ordered pair is always Row's (Player 1's), the second is Col's (Player 2's).

The game $\mathcal{G}$ of Stag Hunt corresponds to the following game matrix:

	hunt stag	hunt hare
hunt stag	both stag	2 hare only
hunt hare	1 hare only	both hare

Plugging in utilities:

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Classifying Games.

A **zero-sum** (or strictly competitive) game is one in which the winner's gain is equal to the loser's loss. A **nonzero-sum** game is one in which this criterion is not satisfied.

A **cooperative** game is one in which the players are *able* to communicate with one another and to form *binding* agreements (they may choose not to do so). **Non-cooperative** games are those that do not allow communication and binding agreements. Games with no mechanism for enforcing agreements are non-cooperative.

A **two-person** game is one with just two players (or teams). An **n-person** game is any game with more than two players (or teams).

A **simultaneous-move** game is one in which players select their strategies at the same time. A **sequential-move** game is one in which players take turns (and, so, have knowledge of the other players' previous moves).

Classifying Games (cont.).

Perfect information games are those sequential games in which all players have full information about the strategies played by other players. Think chess.

Symmetric games are those in which all players face the same strategies and outcomes.

An **iterated** game is one that is played more than one time by the same players and with the same setup. This can dramatically alter the best strategies when compared to the same game played in a *non-iterated* form.

A game is said to be **finite** iff the actions available to each player is a finite set, otherwise it is **infinite**.

More formally, we can say:

Def 2.1.2. A *strictly competitive* or *zero-sum* two-player strategic game $\mathcal{G}$ is one where for any outcomes $o_1, o_2 \in \mathcal{O}$, $o_1 \succcurlyeq_1 o_2$ if and only if $o_2 \succcurlyeq_2 o_1$.

If player 1's preferences are represented with $u_1 : \mathcal{O} \to \mathbb{R}$, then player 2's preferences are representable with $u_2 : \mathcal{O} \to \mathbb{R}$ where $u_1(o) + u_2(o) = 0$ for all $o \in \mathcal{O}$.

Rock, Paper, & Scissors is zero-sum.

Stag Hunt is not.

Zero-sum, two-player games are the most well understood by game theorists. They are the only games for which the math is pretty much complete and for which we can *always* specify a solution in terms of *equilibrium strategies*. They're also the easiest to understand.

The most famous solution concept in game theory is the **Nash equilibrium**.

Let $\langle a_j, a_{-j} \rangle$ designate the action profile where player j performs action $a_j \in \mathcal{A}_j$ and all of the other players perform $a_{-j} \in \times_{i \in \mathcal{N} \setminus \{j\}} \mathcal{A}_i$.

Def 2.1.3. A Nash equilibrium of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$g(\langle a_j^*, a_{-j}^* \rangle) \succcurlyeq_j g(\langle a_j, a_{-j}^* \rangle)$ for each $a_j \in \mathcal{A}_j$.

Note on Notation.

- a_j is an action or strategy of one of the players $j \in \mathcal{N}$, i.e., $a_j \in \mathcal{A}_j$.
- a_{-j} is an action or strategy of a *different* player.
- a_j^* is the proposed Nash equilibrium action profile for j.

Less formally, a *Nash equilibrium* is a set of strategies "such that each player's strategy maximizes his pay-off if the strategies of the others are held fixed. Thus each player's strategy is optimal against those of the others." (Nash 1950: 7)

Osborne and Rubinstein: "No player can profitably deviate, given the actions of the other players."

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Does Rock, Paper, & Scissors have any Nash equilibria?

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Does either player have incentive to deviate from $\langle \text{paper}, \text{paper} \rangle$?

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Yes. Player 1 does well to choose scissors instead of paper.

$$u_1(g(\langle \text{scissors}, \text{paper} \rangle)) > u_1(g(\langle \text{paper}, \text{paper} \rangle)).$$

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Yes. Player 2 also does well to choose scissors instead of paper.

$$u_2(g(\langle \text{paper}, \text{scissors} \rangle)) > u_2(g(\langle \text{paper}, \text{paper} \rangle)).$$

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

In fact, Rock, Paper, & Scissors has *no* Nash equilibria.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does Stag Hunt have any Nash equilibria?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does either player have incentive to deviate from $\langle \text{hunt stag}, \text{hunt stag} \rangle$?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

No. So $\langle \text{hunt stag}, \text{hunt stag} \rangle$ is a Nash equilibrium.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does either player have incentive to deviate from $\langle \text{hunt hare}, \text{hunt hare} \rangle$?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

No. So $\langle$hunt hare, hunt hare$\rangle$ is a second Nash equilibrium.

Homework for Extra Credit on Midterm (up to 5 points):

Is this a good example of a Nash equilibrium? Why or why not?

A Beautiful Mind

In a zero sum game with Players 1 and 2 (Row and Col), $u_2(o) = -u_1(o)$ so we need include only the utility of Player 1 (Row) in the game matrix. Player 2's (Col's) utilities will just be the negations of these.

	a_1	a_2	a_3
a_1	8,-8	8,-8	7,-7
a_2	0,0	-10,10	-4,4
a_3	9,-9	0,0	-1,1

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

Does this game have any Nash equilibria?

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$\langle a_1, a_3 \rangle$ is the only Nash equilibrium.

We're going to look at three ways to find equilibrium pairs for zero-sum, two player games.

- Dominance Reasoning
- Maximin Theorem
- Mixed-Strategy Solutions

Dominance Reasoning.

Dominance Reasoning (DR) is an imperfect but sometimes fruitful strategy for solving games.

To *solve* a game is to determine what rational agents would do when playing the game.

DR is imperfect because it will not give us a solution in every case and because in some cases it leads to unacceptable outcomes.

DR is sometimes fruitful because it can be used in both zero-sum and nonzero-sum games to give us a feel for the game we're dealing with.

DR makes two key assumptions:

- All players are *rational*.
- Common Knowledge of Rationality (CKR) - Each player knows that the other players are rational and knows that each of the other players knows that each player is rational, and so on.

No rational agent would play a dominated strategy.

This allows us to rule out strategies that are dominated.

Once we've ruled out a strategy, we can re-draw the matrix without that strategy. This may, in turn, rule out strategies for the other player that are dominated in teh new matrix.

Let's explore some examples.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

Does this game have any Nash equilibria?

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

Let's look at Row first. Do any strategies dominate any others?

Next, let's look at Col. (Remember that, since it's a zero-sum game, Col's outcomes are the negation of Row's.) Any dominant strategies for Col?

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

For Col, a_1 *strictly dominates* a_2 and a_3.

So, we can ignore those two strategies. Col will always play a_1.

Now, does this change things for Row? Does she have a dominant strategy now?

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

For Row, a_1 and a_2 both *strictly dominate* a_3.

So, Row will play a_1 or a_2, since she knows that Col will always play a_1.

$\langle a_1, a_1 \rangle$ and $\langle a_2, a_1 \rangle$ are both Nash equilibria of this game. (You can have more than one solution.)

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

Does this game have any Nash equilibria?

Let's look at Row first. Does any strategy dominate any other?

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

a_3 dominates the other strategies, so Row will play it and Col knows this.

Does Col have any dominant strategies in this case?

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

a_1 and a_4 are dominant for Col.

So, $\langle a_3, a_1\rangle$ and $\langle a_3, a_4\rangle$ are Nash equilibria.

Are they the only ones?

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

a_1 and a_4 are dominant for Col.

So, $\langle a_3, a_1 \rangle$ and $\langle a_3, a_4 \rangle$ are Nash equilibria.

Are they the only ones?

What about $\langle a_1, a_1 \rangle$ and $\langle a_1, a_4 \rangle$? Would either player have an incentive to deviate?

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

If we begin by looking at Col's dominant strategies, we can see this more clearly.

a_1 and a_4 dominate a_2 and a_3.

So we can rule those strategies out.

Row knows that Col will play a_1 or a_4, so Row will play a_1 or a_3.

Another example.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

$\langle a_1, a_1 \rangle$, $\langle a_1, a_4 \rangle$, $\langle a_3, a_1 \rangle$, and $\langle a_3, a_4 \rangle$ are the Nash equilbria.

Lesson.

Dominance reasoning doesn't always give us *all* of the solutions to a game.

Dominance reasoning can also fail to give us any solution at all to a game.

And dominance reasoning can sometimes lead to unacceptable outcomes.
Example. Centipede Game

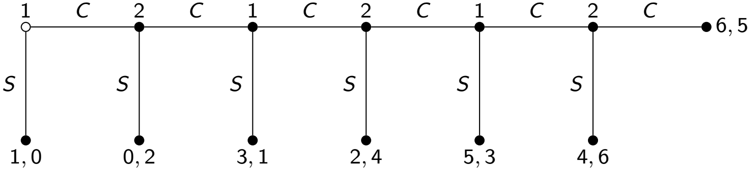

The Nash equilibria of two-player, zero-sum games have various nice properties.

- **Minimax Condition** A pair of strategies is in equilibrium if (but not only if) the outcome determined by the strategies equals the minimal value of the row and the maximal value of the column.
- These solutions can be determined using the **maximinimizer** method.
- All maximinimizer strategy pairs are Nash equilibria.
- All maximinimizer strategy pairs have the same value, and this is the *value of the game*.

See 2.1 Part 2 for Maximin Theorem and Mixed Strategies.