

CS 440/ECE448 Lecture 30: Game Theory

Mark Hasegawa-Johnson

These slides are in the public domain.

Prisoner A \ Prisoner B	Prisoner B stays silent (<i>cooperates</i>)	Prisoner B betrays (<i>defects</i>)
Prisoner A stays silent (<i>cooperates</i>)	Each serves 1 year	Prisoner A: 3 years Prisoner B: goes free
Prisoner A betrays (<i>defects</i>)	Prisoner A: goes free Prisoner B: 3 years	Each serves 2 years

https://en.wikipedia.org/wiki/Prisoner's_dilemma

Today: Games with Simultaneous Moves

Assume:

- two-player game, deterministic environment (not necessary, but simplifies the problem),
- rational players (each player tries to maximize their own reward),
- not zero-sum (game can have 0, 1, or 2 winners),
- simultaneous moves.

Some surprising results:


1. The rational course of action changes may depend on your belief about what the other player will do (Nash equilibrium).
2. There are different ways to define “optimum” (Pareto optimal outcomes).
3. There may be a Pareto optimal outcome that a rational player is forced to reject (Dominant strategy).

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
 - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
 - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
 - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Paparazzi Game (Antisymmetric Games)
 - No pure-strategy Nash equilibrium

Notation: sequential games


- Players take turns acting (e.g., dog moves first, then cat)
- Each node represents the action of one player (e.g., each animal can go either L or R)
- Terminal node is marked with the value for each player



Notation: simultaneous games

The payoff matrix shows:

- Each column is a different move for player 1.
- Each row is a different move for player 2.
- Each square is labeled with the rewards earned by each player in that square.



		Player 2 (Cat)	
		L	R
Player 1 (Dog)	L	1, 2	7, 4
	R	5, 1	5, 4

Payoff matrix

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
 - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
 - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
 - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Paparazzi Game (Antisymmetric Games)
 - No pure-strategy Nash equilibrium

Stag hunt



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

		Alice	
		Defect	Cooperate
Bob	Defect	10 / 10	10 / 0
	Cooperate	0 / 10	100 / 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

Apparently first described by Jean-Jacques Rousseau:

- If both hunters (Bob and Alice) cooperate in hunting for the stag → each gets to take home half a stag (100lbs)
- If one hunts for the stag, while the other wanders off and bags a hare → the defector gets a hare (10lbs), the cooperator gets nothing.
- If both hunters defect → each gets to take home a hare.

Nash Equilibrium



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

		Alice	
		Defect	Cooperate
Bob	Defect	10 / 10	0 / 10
	Cooperate	0 / 10	100 / 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

A Nash Equilibrium is a game outcome such that each player, knowing what the other player will do, responds rationally.

Respond rationally = the player behaves in a manner that maximizes their reward. If all rewards are the same, then all actions are rational.

Nash Equilibrium



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

		Alice	
		Defect	Cooperate
Bob	Defect	10, 10	0, 10
	Cooperate	0, 10	100, 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

Example: (Defect,Defect) is a Nash equilibrium.

- Alice knows that Bob will defect, so she defects.
- Bob knows that Alice will defect, so he defects.
- Neither player can **rationaly** change his or her move, unless the other player also changes.

Nash Equilibrium



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

		Alice	
		Defect	Cooperate
Bob	Defect	10, 10	0, 10
	Cooperate	0, 10	100, 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

(Cooperate, Cooperate) is also a Nash equilibrium!

- Alice knows that Bob will cooperate, so she cooperates!
- Bob knows that Alice will cooperate, so she cooperates!
- Neither player can **rationaly** change his or her move, unless the other player also changes.

Surprising result #1: Nash equilibrium depends on belief



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

Bob

Defect
Cooperate

		Alice	
		Defect	Cooperate
Bob	Defect	10, 10	10, 0
	Cooperate	0, 10	100, 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

Surprising result: The rational course of action depends on what you believe the other player will do.

How is “belief” formed?

- Many real-world cases: You can watch each other acting, and you have plenty of time to change your own action to match
- Other cases: You must trust the other player!

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
 - Nash Equilibrium: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Games)
 - Pareto Optimal outcome: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
 - Dominant Strategy: an action that is rational regardless of what the other player does
- Example: Paparazzi Game (Antisymmetric Games)
 - No pure-strategy Nash equilibrium

Mending Fences

I let my neighbor know beyond the hill;
And on a day we meet to walk the line
And set the wall between us once again.
We keep the wall between us as we go.

...
There where it is we do not need the wall:
He is all pine and I am apple orchard.
My apple trees will never get across
And eat the cones under his pines, I tell him.

- from *Mending Wall*, Robert Frost



https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

Asymmetric Games



https://commons.wikimedia.org/wiki/File:Boblar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

		Alice	
		Wheat	Cattle
Bob	Wheat	20, 10	10, 10
	Cattle	10, 10	10, 20



https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

Alice has lots of cattle pasture, and little wheat field.

Bob has lots of wheat field, and little cattle pasture.

- Mend the wheat fields: Bob earns \$20, Alice earns \$10
- Mend the cattle pastures: Alice earns \$20, Bob earns \$10
- If they don't cooperate, they each earn \$10

Asymmetric Games



https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield_fence_on_Bakerton_Road.jpg

		Alice	
		Wheat	Cattle
Bob	Wheat	20, 10	10, 10
	Cattle	10, 10	10, 20



https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

These 3 outcomes are Nash equilibria.

- (W,W): Bob has no reason to change, Alice has no reason to change
- (W,C): Bob has no reason to change, Alice has no reason to change
- (C,C): Bob has no reason to change, Alice has no reason to change

What happens if they trust one another?



https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

		Alice	
		Wheat	Cattle
Bob	Wheat	20 / 10	10 / 10
	Cattle	10 / 10	10 / 20



https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

What happens if they discuss their actions, and make promises, and trust one another? It depends on whose needs are more urgent.

- If Bob's needs are more urgent, they will work together to mend the wheat fields
- If Alice's needs are more urgent, they will work together to mend the cattle pasture

Pareto optimal outcomes



https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Baker_Road.jpg

		Alice	
		Wheat	Cattle
Bob	Wheat	20, 10	10, 10
	Cattle	10, 10	10, 20



https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

An outcome is Pareto optimal if there is no way to increase value for one player except by decreasing value for the other.

- (W,C) and (C,W) are **not** Pareto-optimal: either player, by changing their action, can increase reward for the other player **without decreasing reward for themselves**
- (C,C) and (W,W) are **both** Pareto-optimal: the only way to increase reward for Bob is by decreasing reward for Alice, or vice versa

Stag hunt



Photo by Scott Bauer, Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=245466>

		Alice	
		Defect	Cooperate
Bob	Defect	10 / 10	10 / 0
	Cooperate	0 / 10	100 / 100



By Ancheta Wis, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=68432449>

An outcome is Pareto optimal if there is no way to increase value for one player except by decreasing value for the other.


For the stag hunt, only (C,C) is Pareto-optimal.

Pareto optimality in engineering

- In engineering, Pareto optimality is a way to discuss tradeoffs between two different optimality criteria
- Suppose you're designing a system with two optimality criteria: A and B
- The "Pareto frontier" is the set of solutions for which A cannot be improved without degrading B, and vice versa

Pareto optimality in engineering


- Example:
 - A = True positive rate (TPR)
 - B = False positive rate (FPR)
- The receiver operating curve (ROC) is the curve showing TPR versus FPR
- Each classifier has:
 - A hyperparameter
 - Adjusting the hyperparameter changes the balance of TPR/FPR
 - Each classifier has a different ROC



CC-SA 3.0, <https://commons.wikimedia.org/wiki/File:Roccurves.png>

Pareto optimality in engineering

- The “Pareto frontier” is the set of solutions for which TPR cannot be improved without degrading FPR
- In this example, each of the 3 classifiers is on the Pareto frontier for some values of its hyperparameter



CC-SA 3.0, <https://commons.wikimedia.org/wiki/File:Roccurves.png>

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
 - Nash Equilibrium: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
 - Pareto Optimal outcome: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
 - Dominant Strategy: an action that is rational regardless of what the other player does
- Example: Paparazzi Game (Antisymmetric Games)
 - No pure-strategy Nash equilibrium

Prisoner's dilemma

- Two criminals have been arrested and the police visit them separately
- If one player testifies against the other and the other refuses, the one who testified goes free and the one who refused gets a 10-year sentence
- If both players testify against each other, they each get a 5-year sentence
- If both refuse to testify, they each get a 1-year sentence



Bob:
Testify

Bob:
Refuse

	Alice: Testify	Alice: Refuse
Bob: Testify		
Bob: Refuse		

By Monogram Pictures,
Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=50338507>

Prisoner's dilemma

- Two criminals have been arrested and the police visit them separately
- If one player testifies against the other and the other refuses, the one who testified goes free and the one who refused gets a 10-year sentence
- If both players testify against each other, they each get a 5-year sentence
- If both refuse to testify, they each get a 1-year sentence



	Alice: Testify	Alice: Refuse
Bob: Testify	-5	-10
Bob: Refuse	-10	-1

By Monogram Pictures,
Public Domain,
<https://commons.wikimedia.org/w/index.php?curid=50338507>

Pareto optimality

If you were permitted to discuss options with the other player, what are the different possible outcomes that might result from that discussion?

- If Bob's needs are considered most important, the (-10,0) outcome might result.
- If Alice's needs are considered more important, the (0,-10) outcome might result.
- If their needs are equally important, the (-1,-1) outcome might result.

A ***Pareto optimal*** outcome is an outcome whose cost to player A can only be reduced by increasing the cost to player B.

	Alice: Testify	Alice: Refuse
Bob: Testify	-5	-10
Bob: Refuse	-10	-1



Nash equilibrium

If you knew in advance what your opponent was going to do, what would you do?

- If Bob knew that Alice was going to refuse, then it be rational for Bob to testify (he'd get 0 years, instead of 1).
- If Alice knew that Bob was going to testify, then it would be rational for her to testify (she'd get 5 years, instead of 10).
- If Bob knew that Alice was going to testify, then it would be rational for him to testify (he'd get 5 years, instead of 10).

A **Nash equilibrium** is an outcome such that foreknowledge of the other player's action does not cause either player to change their action.

	Alice: Testify	Alice: Refuse
Bob: Testify	-5	-10
Bob: Refuse	0	-1

Diagram illustrating a 2x2 payoff matrix for a game between Alice and Bob. The matrix shows payoffs for each combination of actions. A blue arrow points from the top-right cell (-10) to the top-left cell (-5), indicating that if Alice knows Bob will testify, she should testify. An orange arrow points from the bottom-right cell (-1) to the bottom-left cell (0), indicating that if Alice knows Bob will refuse, she should refuse. The top-left cell (-5) is highlighted in green, representing the Nash equilibrium outcome where both players testify.



Dominant strategy

If you didn't know in advance what your opponent was going to do, what would you do?

- If Bob knew that Alice was going to refuse, then it be rational for Bob to testify (he'd get 0 years, instead of 1).
- If Bob knew that Alice was going to testify, then it would still be rational for him to testify (he'd get 5 years, instead of 10).

A **dominant strategy** is an action that maximizes reward, for one player, regardless of what the other player does.

	Alice: Testify	Alice: Refuse
Bob: Testify	-5	0
Bob: Refuse	-10	-1

Blue arrows point from the top-right cell (-10) to the top-left cell (-5) and from the bottom-right cell (-1) to the bottom-left cell (0), indicating that for both Alice and Bob, the left column (Testify) is the dominant strategy.



What makes it a Prisoner's Dilemma?

We use that term to mean a game in which

- Defecting is the **dominant strategy** for each player, therefore
- (Defect,Defect) is the only **Nash equilibrium**, even though
- (Defect,Defect) is not a **Pareto-optimal solution**.

	Defect	Cooperate
Defect	Lose Lose	Lose Big Win Big
Cooperate	Win Big Lose Big	Win Win

http://en.wikipedia.org/wiki/Prisoner's_dilemma

Try the quiz!

Try the quiz!

Existence of Nash equilibria

- If a player has a dominant strategy, there exists a Nash equilibrium in which the player plays that strategy and the other player plays the *best response* to that strategy.
- If both players have dominant strategies, there exists a Nash equilibrium in which they play those strategies.

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
 - Nash Equilibrium: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
 - Pareto Optimal outcome: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
 - Dominant Strategy: an action that is rational regardless of what the other player does
- Example: Paparazzi Game (Antisymmetric Games)
 - No pure-strategy Nash equilibrium

Does every game have a Nash equilibrium?

- Answer: It's a trick question
- Strategies we've discussed today are all **pure strategies**, meaning that each player chooses exactly one action. There are games that have no pure-strategy Nash equilibrium.
- Next time we will discuss **mixed strategies**, meaning that each player chooses an action at random. If you allow mixed strategies, then every game has a Nash equilibrium.



No pure-strategy equilibrium: The Paparazzi Game



- Karrie Kingfisher wants coffee
- Espresso Royale? Café Paradiso?
- If Pablo Paparazzo guesses correctly, he gets to take her photo, and she wastes time trying to buy it from him

		Pablo Action	
		CP	ER
Karrie Action	CP	-4 / +5	+2 / 0
	ER	+2 / 0	-4 / +5

Images:

Jyotsnu Vasudeva

https://commons.wikimedia.org/wiki/File:Megaceryle_alcyon_f_Sam_Smith_Park_Toronto.jpg

https://commons.wikimedia.org/wiki/File:Paparazzi_La_Dolce_Vita.jpg

Alejandro Aragon



No pure-strategy equilibrium: The Paparazzi Game



- If Karrie knows Pablo will be at Paradiso, it's rational for her to go to Espresso

		Pablo Action	
		CP	ER
Karrie Action	CP	-4 / +5	+2 / 0
	ER	+2 / 0	-4 / +5

A blue arrow points to the cell where Karrie chooses ER and Pablo chooses CP, with a value of +2.

Images:
 Jyotsnu Vasudeva
https://commons.wikimedia.org/wiki/File:Megaceryle_alcyon_f_Sam_Smith_Park_Toronto.jpg
https://commons.wikimedia.org/wiki/File:Paparazzi_La_Dolce_Vita.jpg
 Alejandro Aragon



No pure-strategy equilibrium: The Paparazzi Game



- If Karrie knows Pablo will be at Paradiso, it's rational for her to go to Espresso
- If Pablo knows Karrie will be at ER, he should go to ER too

		Pablo Action	
		CP	ER
Karrie Action	CP	-4 / +5	+2 / 0
	ER	+2 / -4	0 / +5

A blue arrow points down from the top-left cell (-4 / +5) to the bottom-left cell (+2 / -4). A grey arrow points right from the bottom-left cell (+2 / -4) to the bottom-right cell (0 / +5).

Images:
 Jyotsnu Vasudeva
https://commons.wikimedia.org/wiki/File:Megaceryle_alcyon_f_Sam_Smith_Park_Toronto.jpg
https://commons.wikimedia.org/wiki/File:Paparazzi_La_Dolce_Vita.jpg
 Alejandro Aragon



No pure-strategy equilibrium: The Paparazzi Game



- If Karrie knows Pablo will be at Paradiso, it's rational for her to go to Espresso
- If Pablo knows Karrie will be at ER, he should go to ER too
- ...but then Karrie goes to CP...

		Pablo Action	
		CP	ER
Karrie Action	CP	-4 / +5	+2 / 0
	ER	+2 / -4	0 / +5

Blue arrows point from the (CP, CP) cell to the (ER, CP) cell and from the (ER, ER) cell to the (CP, ER) cell. A grey arrow points from the (ER, CP) cell to the (ER, ER) cell.

Images:

Jyotsnu Vasudeva

https://commons.wikimedia.org/wiki/File:Megaceryle_alcyon_f_Sam_Smith_Park_Toronto.jpg

https://commons.wikimedia.org/wiki/File:Paparazzi_La_Dolce_Vita.jpg

Alejandro Aragon



No pure-strategy equilibrium: The Paparazzi Game



- If Karrie knows Pablo will be at Paradiso, it's rational for her to go to Espresso
- If Pablo knows Karrie will be at ER, he should go to ER too
- ...but then Karrie goes to CP...
- ...but then Pablo goes to CP...

		Pablo Action	
		CP	ER
Karrie Action	CP	-4 / +5	+2 / 0
	ER	+2 / -4	0 / +5

The table shows a game between Karrie and Pablo. Karrie's actions are CP and ER. Pablo's actions are CP and ER. The payoffs are given as (Karrie, Pablo).
 - If Karrie chooses CP and Pablo chooses CP, Karrie gets -4 and Pablo gets +5.
 - If Karrie chooses CP and Pablo chooses ER, Karrie gets +2 and Pablo gets 0.
 - If Karrie chooses ER and Pablo chooses CP, Karrie gets +2 and Pablo gets -4.
 - If Karrie chooses ER and Pablo chooses ER, Karrie gets 0 and Pablo gets +5.
 Blue arrows point from the (CP, CP) cell to the (ER, CP) cell, and from the (ER, ER) cell to the (CP, ER) cell. Grey arrows point from the (CP, ER) cell to the (ER, ER) cell, and from the (CP, CP) cell to the (CP, ER) cell.

Images:

Jyotsnu Vasudeva

https://commons.wikimedia.org/wiki/File:Megaceryle_alcyon_f_Sam_Smith_Park_Toronto.jpg

https://commons.wikimedia.org/wiki/File:Paparazzi_La_Dolce_Vita.jpg

Alejandro Aragon

Existence of Nash equilibria

- If there are no dominant strategies, then equilibrium is determined by the symmetry of rewards:
 - Some games have pure-strategy Nash equilibrium
 - Some games have only mixed-strategy Nash equilibrium, which we'll discuss next time

Summary

- Dominant strategy
 - A strategy that's optimal for one player, regardless of what the other player does
 - Not all games have dominant strategies
- Nash equilibrium
 - An outcome (one action by each player) such that, knowing the other player's action, each player has no reason to change their own action
 - Every game with a finite set of actions has at least one Nash equilibrium, though it might be a mixed-strategy equilibrium.
- Pareto optimal
 - An outcome such that neither player would be able to win more without simultaneously forcing the other player to lose more
 - Every game has at least one Pareto optimal outcome. Usually there are many, representing different tradeoffs between the two players.