CS 540 Introduction to Artificial Intelligence
Games I

University of Wisconsin-Madison

Spring 2023

Outline

Homeworks:

- Homework 9 due Thursday April 27
- Homework 10 due Thursday May 4

Class roadmap:

Tuesday, April 18	Games I
Thursday, April 20	Games II
Tuesday, April 25	Reinforcement Learning I
Thursday, April 27	Reinforcement Learning I
Tuesday, May 2	Review of RL + Games
Thursday, May 4	Ethics and Trust in AI

Outline

- Introduction to game theory
  - Properties of games, mathematical formulation
- Simultaneous-Move Games
  - Normal form, strategies, dominance, Nash equilibrium

So Far in The Course

We looked at techniques:

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.

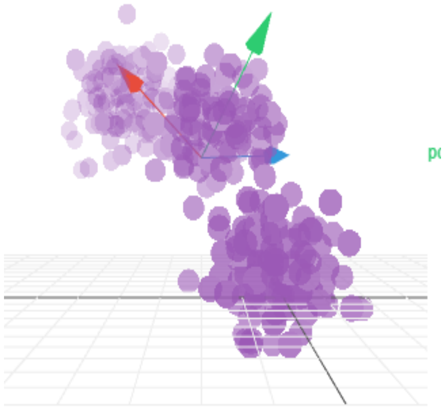

Victor Powell

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.
- **Supervised:** Train a model to make predictions. More structure (labels).

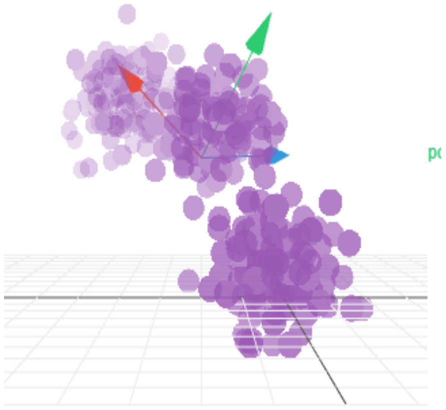

Victor Powell

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.
- **Supervised:** Train a model to make predictions. More structure (labels).

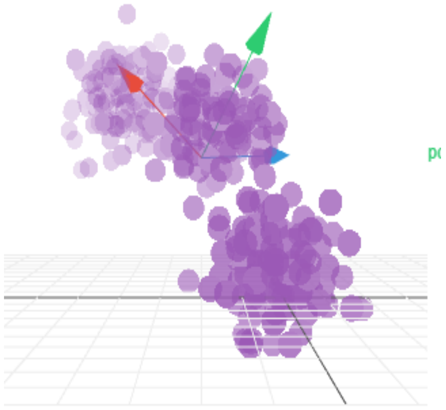

Victor Powell

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.
- **Supervised:** Train a model to make predictions. More structure (labels).
- **Planning and Games**: Much more structure.

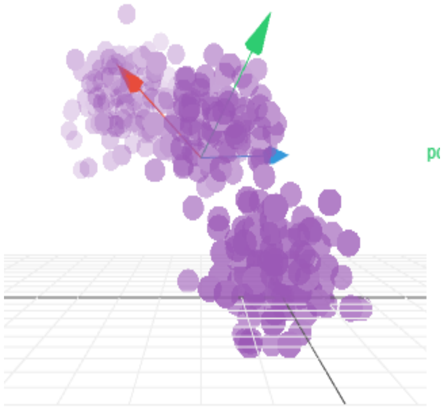

Victor Powell

So Far in The Course

We looked at techniques:

- **Unsupervised:** See data, do something with it. Unstructured.
- **Supervised:** Train a model to make predictions. More structure (labels).
- **Planning and Games**: Much more structure.

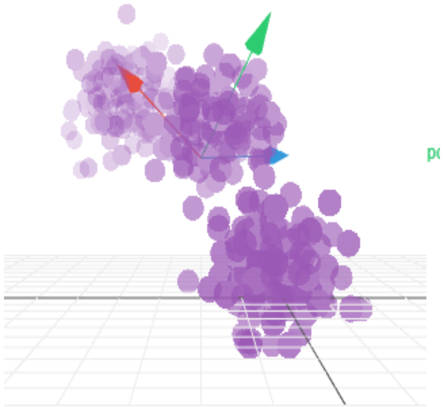

Victor Powell

More General Model

Suppose we have an **agent** **interacting** with the **world**

More General Model

Suppose we have an **agent** **interacting** with the **world**

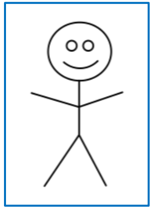

More General Model

Suppose we have an **agent interacting** with the **world**

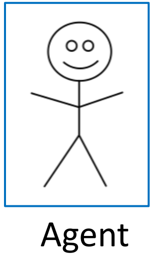

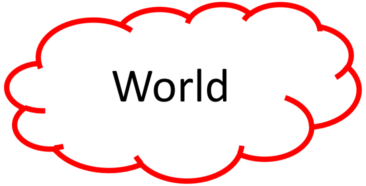

More General Model

Suppose we have an **agent** **interacting** with the **world**

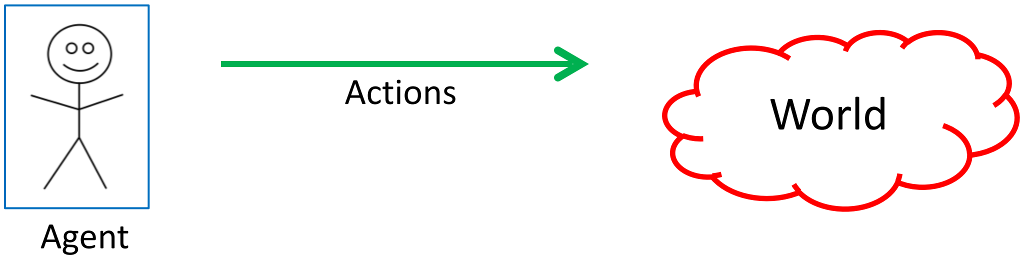

More General Model

Suppose we have an **agent** **interacting** with the **world**

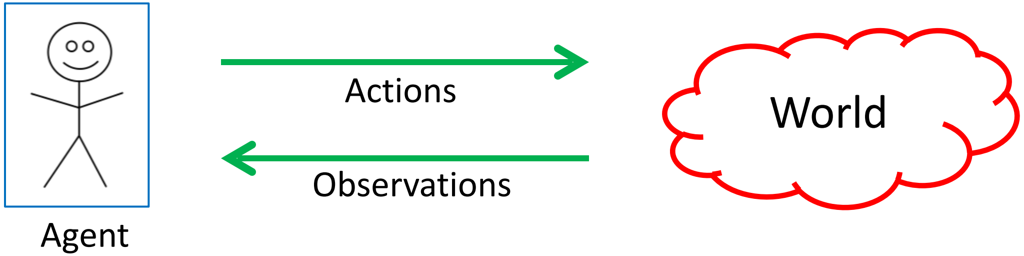

More General Model

Suppose we have an **agent** **interacting** with the **world**

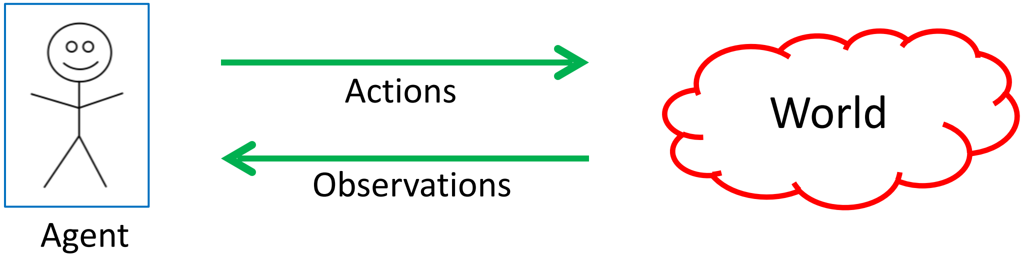

- Agent receives a reward based on state of the world

More General Model

Suppose we have an **agent** **interacting** with the **world**

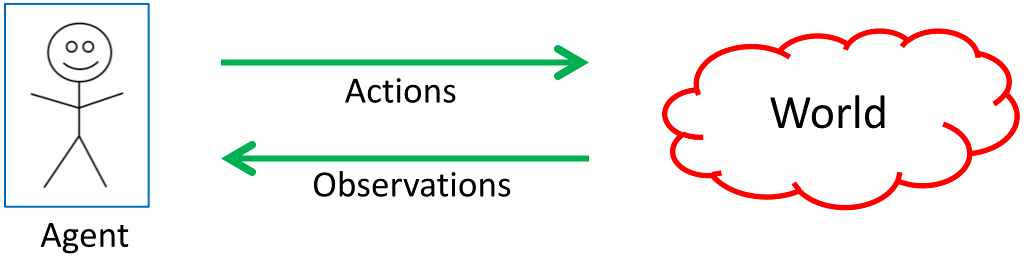

- Agent receives a reward based on state of the world
  - **Goal**: maximize reward / utility

More General Model

Suppose we have an **agent** **interacting** with the **world**

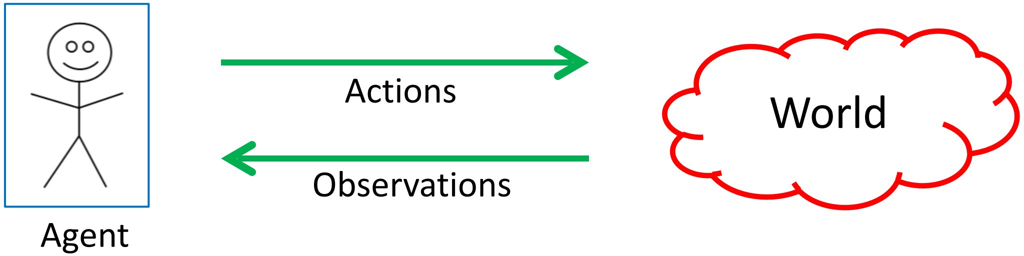

- Agent receives a reward based on state of the world
  - **Goal**: maximize reward / utility **($$$)**

More General Model

Suppose we have an **agent interacting** with the **world**

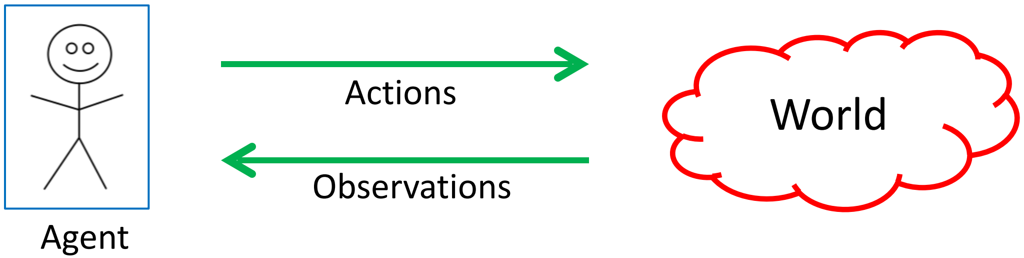

- Agent receives a reward based on state of the world
  - **Goal**: maximize reward / utility **($$$)**
  - Note: now **data** consists of actions, observations, and rewards

More General Model

Suppose we have an **agent** **interacting** with the **world**

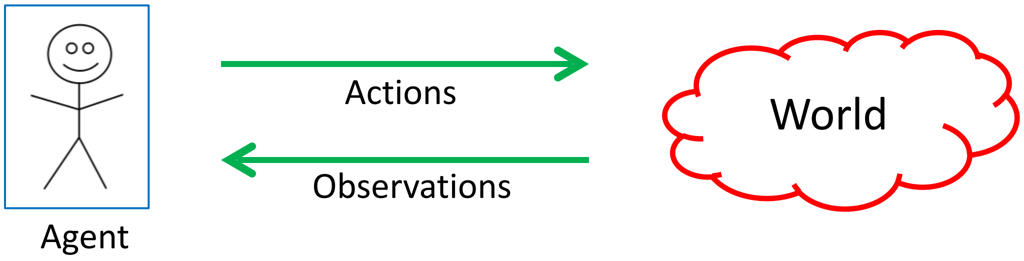

- Agent receives a reward based on state of the world
  - **Goal**: maximize reward / utility **($$$)**
  - Note: now **data** consists of actions, observations, and rewards
  - Setup for decision theory, reinforcement learning, planning

Games: Multiple Agents

Games setup: **multiple** agents

Games: Multiple Agents

Games setup: **multiple** agents

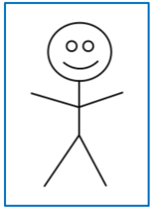

Player 1

Games: Multiple Agents

Games setup: **multiple** agents

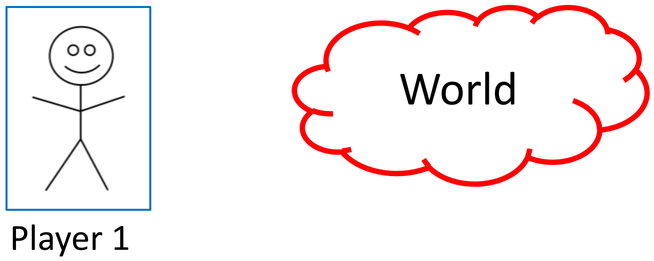

Games: Multiple Agents

Games setup: **multiple** agents

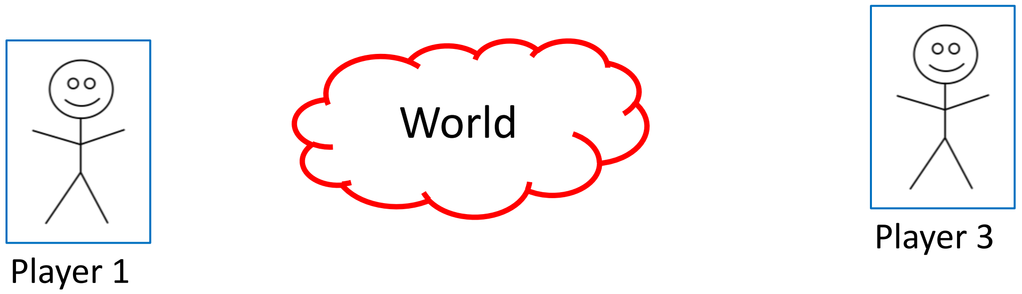

Games: Multiple Agents

Games setup: **multiple** agents

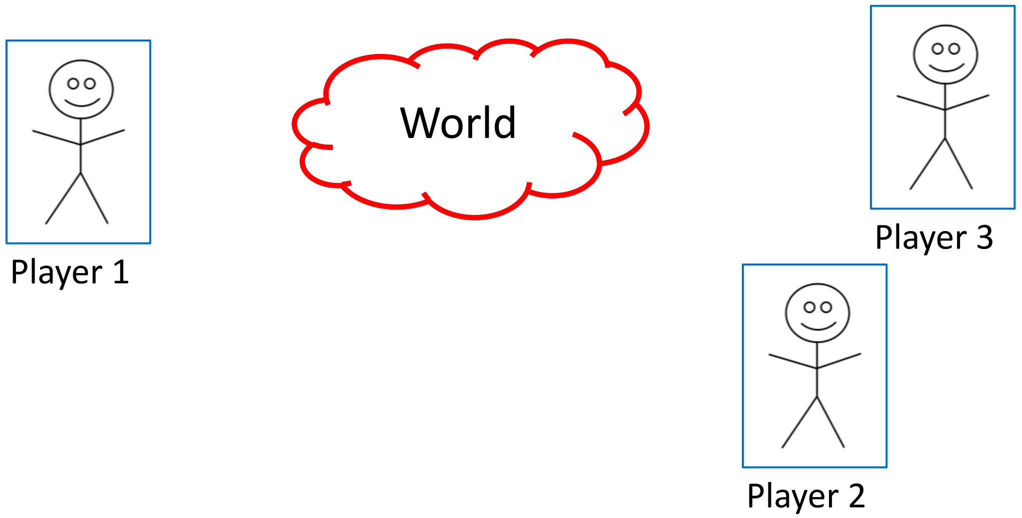

Games: Multiple Agents

Games setup: **multiple** agents

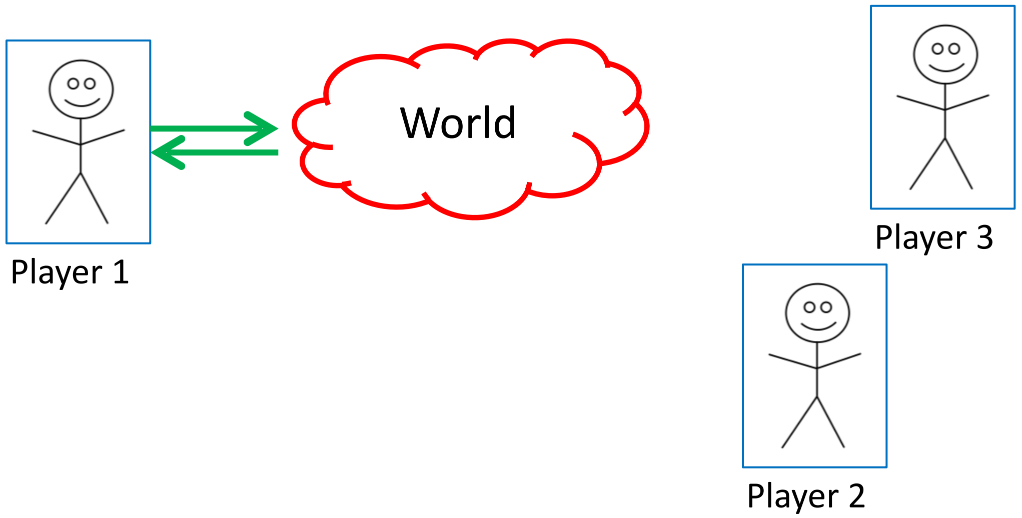

Games: Multiple Agents

Games setup: **multiple** agents

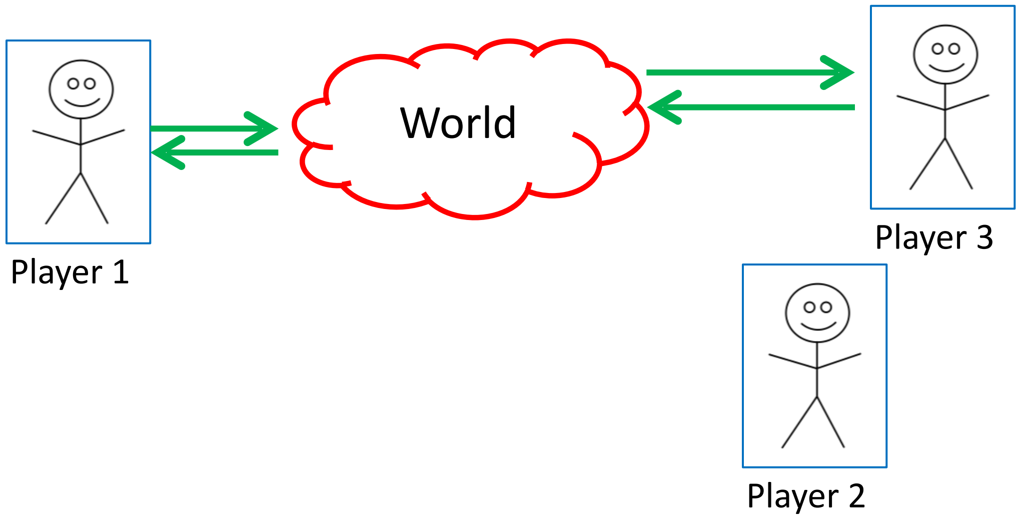

Games: Multiple Agents

Games setup: **multiple** agents

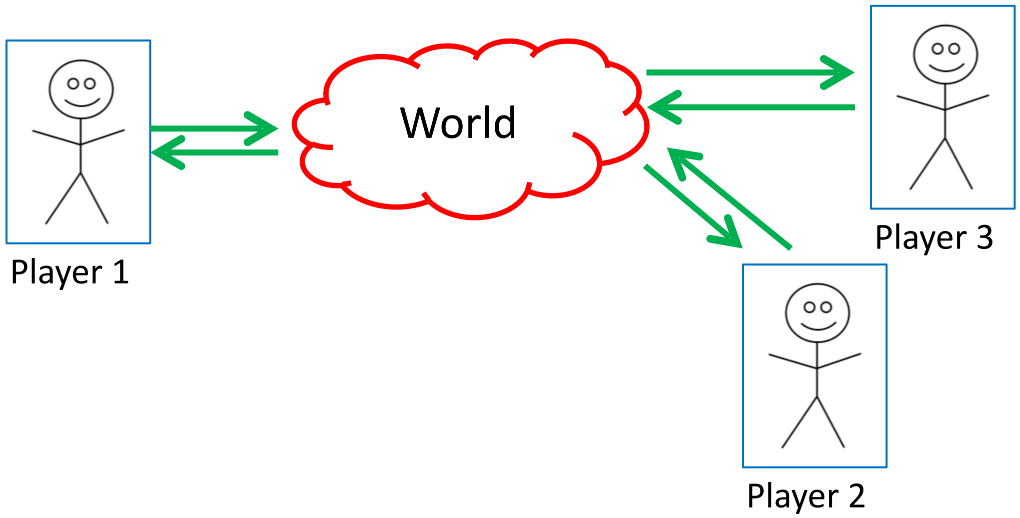

Games: Multiple Agents

Games setup: **multiple** agents

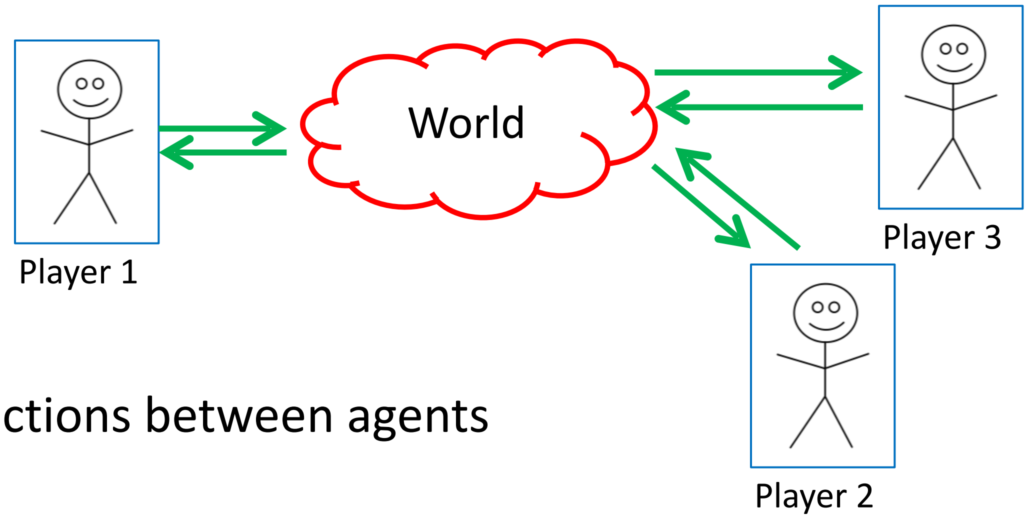

– Now: interactions between agents

Games: Multiple Agents

Games setup: **multiple** agents

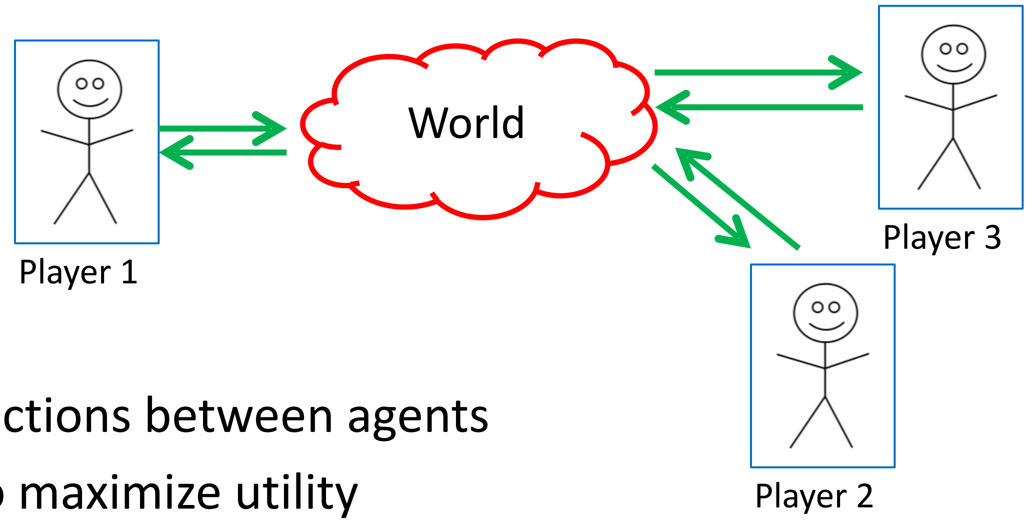

- Now: interactions between agents
- Still want to maximize utility

Games: Multiple Agents

Games setup: **multiple** agents

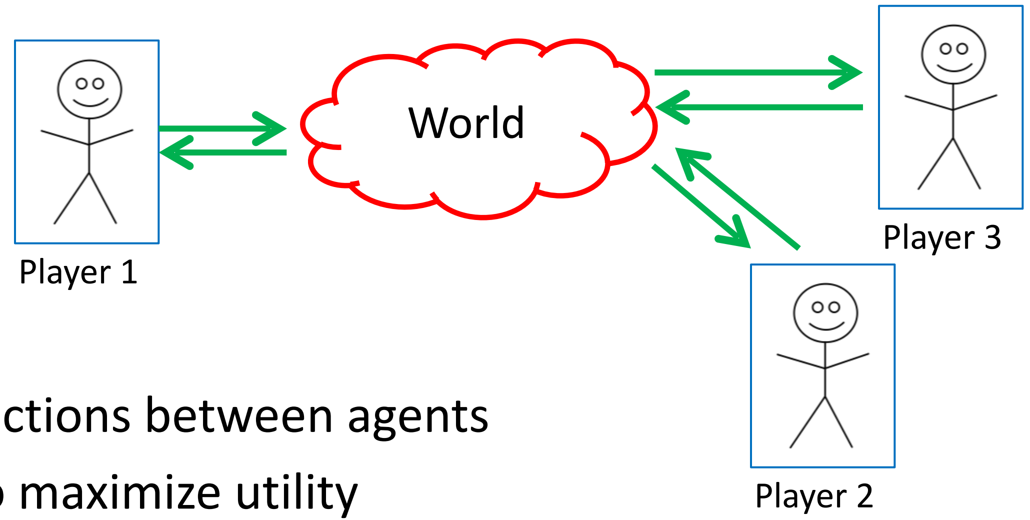

- Now: interactions between agents
- Still want to maximize utility
- Requires **strategic** decision making.

Modeling Games: Properties

Let’s work through **properties** of games

Modeling Games: Properties

Let's work through **properties** of games

- **Number** of agents/players

Modeling Games: Properties

Let’s work through **properties** of games

- **Number** of agents/players
- Action space: finite or infinite

Modeling Games: Properties

Let’s work through **properties** of games

- **Number** of agents/players
- Action space: finite or infinite
- **Deterministic** or **random**

Modeling Games: Properties

Let's work through **properties** of games

- **Number** of agents/players
- Action space: finite or infinite
- **Deterministic** or **random**
- Zero-sum or general-sum

Modeling Games: Properties

Let's work through **properties** of games

- **Number** of agents/players
- Action space: finite or infinite
- **Deterministic** or **random**
- Zero-sum or general-sum
- **Sequential** or **simultaneous moves**

Modeling Games: Properties

Let’s work through **properties** of games

- **Number** of agents/players
- Action space: finite or infinite
- **Deterministic** or **random**
- Zero-sum or general-sum
- **Sequential** or **simultaneous moves**

Wiki

Property 1: **Number** of players

Pretty clear idea: 1 or more players

Property 1: **Number** of players

Pretty clear idea: 1 or more players

- Usually interested in ≥ 2 players

Property 1: **Number** of players

Pretty clear idea: 1 or more players

- Usually interested in ≥ 2 players
- Typically a finite number of players

Property 1: **Number** of players

Pretty clear idea: 1 or more players

- Usually interested in ≥ 2 players
- Typically a finite number of players

Property 1: **Number** of players

Pretty clear idea: 1 or more players

- Usually interested in ≥ 2 players
- Typically a finite number of players

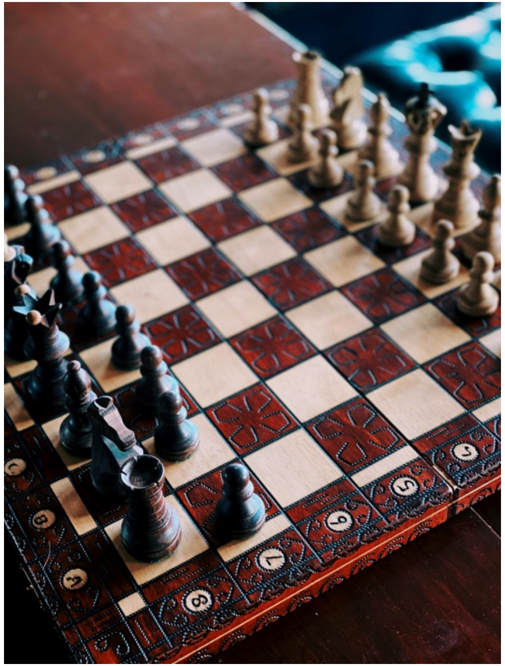

Property 2: Action Space

Property 2: Action Space

Action space: set of possible actions an agent can choose from.

Can be finite or infinite.

Examples:

- Rock-paper-scissors
- Tennis

Property 3: **Deterministic** or **Random**

Property 3: **Deterministic** or **Random**

- Is there **chance** in the game?
  - Poker
  - Scrabble
  - Chess

Property 3: Deterministic or Random

- Is there **chance** in the game?
  - Poker
  - Scrabble
  - Chess

Property 4: **Sum of payoffs**

Property 4: **Sum of payoffs**

- Two basic types: zero sum vs. general sum.

Property 4: **Sum of payoffs**

- Two basic types: zero sum vs. general sum.

Property 4: **Sum of payoffs**

- Two basic types: zero sum vs. general sum.

- Zero sum: one player’s win is the other’s loss
  - Pure competition.
  - Example: rock-paper-scissors

Property 4: **Sum of payoffs**

- Two basic types: zero sum vs. general sum.

- Zero sum: one player’s win is the other’s loss
  – Pure competition.
  – Example: rock-paper-scissors

Property 4: Sum of payoffs

- Two basic types: zero sum vs. general sum.

- Zero sum: one player’s win is the other’s loss
  - Pure competition.
  - Example: rock-paper-scissors

- General sum
  - Example: driving to work, prisoner’s dilemma

Property 5: **Sequential** or **Simultaneous Moves**

Property 5: **Sequential** or **Simultaneous Moves**

- Simultaneous: all players take action at the same time

Property 5: **Sequential** or **Simultaneous Moves**

- Simultaneous: all players take action at the same time
- Sequential: take turns (but payoff only revealed at end of game)

Quiz break:

Quiz break:

Give the properties of the game shown on the right:

Quiz break:

Give the properties of the game shown on the right:

- Number of players?

Quiz break:

Give the properties of the game shown on the right:

- Number of players?
- Deterministic or stochastic?

Quiz break:

Give the properties of the game shown on the right:

- Number of players?
- Deterministic or stochastic?
- Sum of pay-offs?

Quiz break:

Give the properties of the game shown on the right:

- Number of players?
- Deterministic or stochastic?
- Sum of pay-offs?
- Finite or infinite action-space?

Quiz break:

Give the properties of the game shown on the right:

- Number of players?
- Deterministic or stochastic?
- Sum of pay-offs?
- Finite or infinite action-space?
- Sequential or simultaneous?

Normal Form Game

Mathematical description of simultaneous games.

Normal Form Game

Mathematical description of simultaneous games.

- *n* players $\{1,2,\ldots,n\}$

Normal Form Game

Mathematical description of simultaneous games.

- n players $\{1,2,\ldots,n\}$
- Player i chooses strategy a_i from A_i.

Normal Form Game

Mathematical description of simultaneous games.

- n players $\{1,2,\ldots,n\}$
- Player i chooses strategy a_i from A_i.
- Strategy profile: $a = (a_1, a_2, \ldots, a_n)$

Normal Form Game

Mathematical description of simultaneous games.

- n players $\{1,2,\ldots,n\}$
- Player i chooses strategy a_i from A_i.
- Strategy profile: $a = (a_1, a_2, \ldots, a_n)$
- Player i gets rewards $u_i(a)$

Normal Form Game

Mathematical description of simultaneous games.

- n players $\{1,2,\ldots,n\}$
- Player i chooses strategy a_i from A_i.
- Strategy profile: $a = (a_1, a_2, \ldots, a_n)$
- Player i gets rewards $u_i(a)$
  - **Note**: reward depends on other players!

Normal Form Game

Mathematical description of simultaneous games.

- n players $\{1,2,\ldots,n\}$
- Player i chooses strategy a_i from A_i.
- Strategy profile: $a = (a_1, a_2, \ldots, a_n)$
- Player i gets rewards $u_i(a)$
  - **Note**: reward depends on other players!

- We consider the simple case where all reward functions are common knowledge.

Example of Normal Form Game

Ex: Prisoner’s Dilemma

Example of Normal Form Game

Ex: Prisoner's Dilemma

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Example of Normal Form Game

Ex: Prisoner's Dilemma

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

- 2 players, 2 actions: yields 2x2 payoff matrix

Example of Normal Form Game

Ex: Prisoner's Dilemma

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

- 2 players, 2 actions: yields 2x2 payoff matrix
- Strategy set: {Stay silent, betray}

Example of Normal Form Game

Ex: Prisoner's Dilemma

Player 2 / Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

- 2 players, 2 actions: yields 2x2 payoff matrix
- Strategy set: {Stay silent, betray}

Strictly Dominant Strategies

Strictly Dominant Strategies

Let’s analyze such games. Some strategies are better than others!

Strictly Dominant Strategies

Let's analyze such games. Some strategies are better than others!

- Strictly dominant strategy: if a_i strictly better than a_i' *regardless* of what other players do, a_i is **strictly dominant**

Strictly Dominant Strategies

Let's analyze such games. Some strategies are better than others!

- Strictly dominant strategy: if a_i strictly better than a_i' *regardless* of what other players do, a_i is **strictly dominant**
- I.e., $u_i\left(a_i,\ a_{-i}\right) > u_i(b,\ a_{-i}),\ \forall b \neq a_i,\ \ \forall a_{-i}$

Strictly Dominant Strategies

Let's analyze such games. Some strategies are better than others!

- Strictly dominant strategy: if a_i strictly better than a_i' *regardless* of what other players do, a_i is **strictly dominant**
- I.e., $u_i(a_i,\ a_{-i}) > u_i(b,\ a_{-i}),\ \forall b \neq a_i,\ \forall a_{-i}$

All of the other entries of *a* excluding *i*

Strictly Dominant Strategies

Let's analyze such games. Some strategies are better than others!

- Strictly dominant strategy: if a_i strictly better than a_i' *regardless* of what other players do, a_i is **strictly dominant**
- I.e., $u_i\left(a_i,\ a_{-i}\right) > u_i(b,\ a_{-i}),\ \forall b \neq a_i,\ \ \forall a_{-i}$

All of the other entries of *a* excluding *i*

- Sometimes a dominant strategy does not exist!

Strictly Dominant Strategies Example

Back to Prisoner’s Dilemma

Strictly Dominant Strategies Example

Back to Prisoner's Dilemma

- Examine all the entries: betray strictly dominates

Strictly Dominant Strategies Example

Back to Prisoner's Dilemma

- Examine all the entries: betray strictly dominates
- Check:

Strictly Dominant Strategies Example

Back to Prisoner's Dilemma

- Examine all the entries: betray strictly dominates
- Check:

Player 2 / Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy Equilibrium

Dominant Strategy Equilibrium

a^* is a (strictly) dominant strategy equilibrium (DSE), if all players have a strictly dominant strategy a_i^*

Dominant Strategy Equilibrium

$a*$ is a (strictly) dominant strategy equilibrium (DSE),
if all players have a strictly dominant strategy a_i^*

- Rational players will play at DSE, if one exists.

Dominant Strategy Equilibrium

$a*$ is a (strictly) dominant strategy equilibrium (DSE), if all players have a strictly dominant strategy a_i^*

- Rational players will play at DSE, if one exists.

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy Equilibrium

$a*$ is a (strictly) dominant strategy equilibrium (DSE), if all players have a strictly dominant strategy a_i^*

- Rational players will play at DSE, if one exists.

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy: Absolute Best Responses

Dominant Strategy: Absolute Best Responses

Player i's best response to strategy

$$a_{-i}:\quad BR\left(a_{-i}\right) = \underset{a}{\text{argmax}}\, u_i(a,\ a_{-i})$$

Dominant Strategy: Absolute Best Responses

Player i's best response to strategy

$$a_{-i}: \quad BR(a_{-i}) = \underset{a}{\operatorname{argmax}}\, u_i(a,\ a_{-i})$$

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy: Absolute Best Responses

Player i's best response to strategy

$$a_{-i}: \quad BR(a_{-i}) = \underset{a}{\operatorname{argmax}} u_i(a,\ a_{-i})$$

BR(player2=silent) = betray

Player 2 / Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy: Absolute Best Responses

Player i's best response to strategy

$$a_{-i}: \quad BR\left(a_{-i}\right) = \underset{a}{\mathrm{argmax}}\, u_i(a,\ a_{-i})$$

BR(player2=silent) = betray

BR(player2=betray) = betray

Player 2 / Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Dominant Strategy: Absolute Best Responses

Player i's best response to strategy

$$a_{-i}: \quad BR(a_{-i}) = \underset{a}{\operatorname{argmax}}\, u_i(a,\ a_{-i})$$

BR(player2=silent) = betray

BR(player2=betray) = betray

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

a_i^* is the dominant strategy for player i, if

$$a_i^* = BR(a_{-i}),\ \forall\ a_{-i}$$

Dominant Strategy Equilibrium

Dominant Strategy Equilibrium does not always exist.

Dominant Strategy Equilibrium

Dominant Strategy Equilibrium does not always exist.

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Nash Equilibrium

a^* is a Nash equilibrium if no player has an incentive to **unilaterally deviate**

Nash Equilibrium

$a*$ is a Nash equilibrium if no player has an incentive to **unilaterally deviate**

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i$$

Nash Equilibrium

$a*$ is a Nash equilibrium if no player has an incentive to **unilaterally deviate**

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i$$

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Nash Equilibrium

$a*$ is a Nash equilibrium if no player has an incentive to **unilaterally deviate**

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i$$

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Nash Equilibrium

$a*$ is a Nash equilibrium if no player has an incentive to **unilaterally deviate**

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i$$

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Nash Equilibrium: Best Response to Each Other

Nash Equilibrium: Best Response to Each Other

a^* is a Nash equilibrium:

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Equivalently, for each player i:

$$a_i^* \in BR\left(a_{-i}^*\right) = argmax_b\, u_i(b,\ a_{-i}^*)$$

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Equivalently, for each player i:

  $$a_i^* \in BR\left(a_{-i}^*\right) = argmax_b\, u_i(b,\ a_{-i}^*)$$

- Compared to DSE (a DSE is a NE, the other way is generally not true):

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i: u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Equivalently, for each player i:

  $$a_i^* \in BR\left(a_{-i}^*\right) = argmax_b\, u_i(b,\ a_{-i}^*)$$

- Compared to DSE (a DSE is a NE, the other way is generally not true):

$$a_i^* = BR\left(a_{-i}\right),\ \forall\ a_{-i}$$

Nash Equilibrium: Best Response to Each Other

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i: u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Pure Nash equilibrium:

Nash Equilibrium: Best Response to Each Other

a^* is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Pure Nash equilibrium:
  - A **pure strategy** is a deterministic choice (no randomness).

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Pure Nash equilibrium:
  - A **pure strategy** is a deterministic choice (no randomness).
  - Later: we will consider **mixed** strategies

Nash Equilibrium: Best Response to Each Other

$a*$ is a Nash equilibrium:

$$\forall i,\ \forall b \in A_i : u_i(a_i^*, a_{-i}^*) \geq u_i(b, a_{-i}^*)$$

(no player has an incentive to **unilaterally deviate**)

- Pure Nash equilibrium:
  - A **pure strategy** is a deterministic choice (no randomness).
  - Later: we will consider **mixed** strategies
  - In pure Nash equilibrium, players can only play pure strategies.

Finding (pure) Nash Equilibria by hand

- As player 1: For each column, find the best response, underscore it.

Finding (pure) Nash Equilibria by hand

- As player 1: For each column, find the best response, underscore it.

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- As player 1: For each column, find the best response, underscore it.

Player 2 Player 1	*L*	*R*
T	<u>2</u>, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- As player 1: For each column, find the best response, underscore it.

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- As player 2: for each row, find the best response, upper-score it.

Player 2 / Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- As player 2: for each row, find the best response, upper-score it.

Player 2 Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- As player 2: for each row, find the best response, upper-score it.

Player 2 / Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Finding (pure) Nash Equilibria by hand

- Entries with both lower and upper bars are pure NEs.

Player 2 / Player 1	*L*	*R*
T	2, 1	0, 0
B	0, 0	1, 2

Pure Nash Equilibrium may not exist

So far, pure strategy: each player picks a deterministic strategy. But:

Pure Nash Equilibrium may not exist

So far, pure strategy: each player picks a deterministic strategy. But:

Player 2 Player 1	*rock*	paper	*scissors*
rock	0, 0	-1, 1	1, -1
paper	1, -1	0, 0	-1, 1
scissors	-1, 1	1, -1	0, 0

Mixed Strategies

Mixed Strategies

Can also randomize actions: “**mixed**”

Mixed Strategies

Can also randomize actions: "**mixed**"

- Player i assigns probabilities x_i to each action

Mixed Strategies

Can also randomize actions: "**mixed**"

- Player i assigns probabilities x_i to each action

$$x_i(a_i), \text{ where } \sum_{a_i \in A_i} x_i(a_i) = 1, x_i(a_i) \geq 0$$

Mixed Strategies

Can also randomize actions: "**mixed**"

- Player i assigns probabilities x_i to each action

$$x_i(a_i), \text{ where } \sum_{a_i \in A_i} x_i(a_i) = 1, x_i(a_i) \geq 0$$

- Now consider **expected rewards**

Mixed Strategies

Can also randomize actions: "**mixed**"

- Player i assigns probabilities x_i to each action

$$x_i(a_i), \text{ where } \sum_{a_i \in A_i} x_i(a_i) = 1, x_i(a_i) \geq 0$$

- Now consider **expected rewards**

Mixed Strategies

Can also randomize actions: "**mixed**"

- Player i assigns probabilities x_i to each action

$$x_i(a_i), \text{ where } \sum_{a_i \in A_i} x_i(a_i) = 1, x_i(a_i) \geq 0$$

- Now consider **expected rewards**

$$u_i(x_i, x_{-i}) = E_{a_i \sim x_i, a_{-i} \sim x_{-i}} u_i(a_i, a_{-i}) = \sum_{a_i} \sum_{a_{-i}} x_i(a_i) x_{-i}(a_{-i}) u_i(a_i, a_{-i})$$

Mixed Strategy Nash Equilibrium

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

$$u_i(x_i^*, x_{-1}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall x_i \in \Delta_{A_i}, \forall i \in \{1, \ldots, n\}$$

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

$$u_i(x_i^*, x_{-1}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall x_i \in \Delta_{A_i}, \forall i \in \{1, \ldots, n\}$$

Better than doing
anything else,
"**best response**"

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

$$u_i(x_i^*, x_{-1}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall x_i \in \Delta_{A_i}, \forall i \in \{1, \ldots, n\}$$

Better than doing anything else, “**best response**”

Space of probability distributions over strategies.

Mixed Strategy Nash Equilibrium

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

$$u_i(x_i^*, x_{-1}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall x_i \in \Delta_{A_i}, \forall i \in \{1, \ldots, n\}$$

Better than doing anything else, "**best response**"

Space of probability distributions over strategies.

- Intuition: nobody can **increase expected reward** by changing only their own strategy.

Mixed Strategy Nash Equilibrium

Example: $x_1(\,.\,) = x_{2(\,.\,)} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

Mixed Strategy Nash Equilibrium

Example: $x_1(\,.\,) = x_{2(\,.\,)} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

Player 2 Player 1	*rock*	paper	*scissors*
rock	0, 0	-1, 1	1, -1
paper	1, -1	0, 0	-1, 1
scissors	-1, 1	1, -1	0, 0

Finding Mixed NE in 2-Player Zero-Sum Game

Example: Two Finger Morra. Show 1 or 2 fingers. The “even player” wins the sum if the sum is even, and vice versa.

Finding Mixed NE in 2-Player Zero-Sum Game

Example: Two Finger Morra. Show 1 or 2 fingers. The "even player" wins the sum if the sum is even, and vice versa.

odd even	*f1*	f2
f1	2, -2	-3, 3
f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Two Finger Morra. Two-player zero-sum game. No pure NE:

odd / even	*f1*	f2
f1	2, -2	-3, 3
f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

		q	1-q
	odd even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q) \; \forall p'$.

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q) \forall p'$.

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q)\, \forall p'$.

But $u_1(p, q) = p u_1(f_1, q) + (1-p) u_1(f_2, q)$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q) \, \forall p'$.

But $u_1(p, q) = pu_1(f_1, q) + (1 - p)u_1(f_2, q)$

Average is no greater than components

	even \ odd	*f1* (q)	f2 (1-q)
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q) \forall p'$.

But $u_1(p, q) = pu_1(f_1, q) + (1-p)u_1(f_2, q)$

Average is no greater than components

➔ $u_1(p, q) = u_1(f_1, q) = u_1(f_2, q)$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

Suppose odd's mixed strategy at NE is (q, 1-q), and even's (p, 1-p)

By definition, p is best response to q: $u_1(p, q) \geq u_1(p', q)\, \forall p'$.

But $u_1(p, q) = p u_1(f_1, q) + (1 - p) u_1(f_2, q)$

Average is no greater than components

➔ $u_1(p, q) = u_1(f_1, q) = u_1(f_2, q)$

We want to find q such that equality holds.

Then even has no incentive to change strategy.

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

		q	1-q
	odd even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1\left(f_1,\ q\right) = u_1(f_2, q)$$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1(f_1,\ q) = u_1(f_2, q)$$

$$2q + (-3)(1-q) = (-3)q + 4(1-q)$$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1\left(f_1,\ q\right) = u_1(f_2, q)$$

$$2q + (-3)\left(1 - q\right) = (-3)q + 4\left(1 - q\right)$$

$$q = \frac{7}{12}$$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1(f_1,\ q) = u_1(f_2, q)$$

$$2q + (-3)(1 - q) = (-3)q + 4(1 - q)$$

$$q = \frac{7}{12}$$

Similarly, $u_2(p, f_1) = u_2(p, f_2)$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1(f_1,\ q) = u_1(f_2, q)$$

$$2q + (-3)(1-q) = (-3)q + 4(1-q)$$

$$q = \frac{7}{12}$$

Similarly, $u_2(p, f_1) = u_2(p, f_2)$

$$p = \frac{7}{12}$$

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Finding Mixed NE in 2-Player 2-action Zero-Sum Game

$$u_1(f_1, q) = u_1(f_2, q)$$

$$2q + (-3)(1 - q) = (-3)q + 4(1 - q)$$

$$q = \frac{7}{12}$$

Similarly, $u_2(p, f_1) = u_2(p, f_2)$

$$p = \frac{7}{12}$$

At this NE, even gets -1/12, odd gets 1/12.

		q	1-q
	odd / even	*f1*	f2
p	*f1*	2, -2	-3, 3
1-p	f2	-3, 3	4, -4

Properties of Nash Equilibrium

Major result: (John Nash ’51)

Properties of Nash Equilibrium

Major result: (John Nash ’51)

- Every finite (players, actions) game has at least one Nash equilibrium

Properties of Nash Equilibrium

Major result: (John Nash ’51)

- Every finite (players, actions) game has at least one Nash equilibrium
  - But not necessarily **pure** (i.e., deterministic strategy)

Properties of Nash Equilibrium

Major result: (John Nash '51)

- Every finite (players, actions) game has at least one Nash equilibrium
  - But not necessarily **pure** (i.e., deterministic strategy)
- Could be more than one

Properties of Nash Equilibrium

Major result: (John Nash ’51)

- Every finite (players, actions) game has at least one Nash equilibrium
  - But not necessarily **pure** (i.e., deterministic strategy)
- Could be more than one
- Searching for Nash equilibria: computationally **hard**.

Properties of Nash Equilibrium

Major result: (John Nash ’51)

- Every finite (players, actions) game has at least one Nash equilibrium
  - But not necessarily **pure** (i.e., deterministic strategy)
- Could be more than one
- Searching for Nash equilibria: computationally **hard**.
  - Exception: two-player zero-sum games (can be found with linear programming).

Break & Quiz

Q 2.1: Which of the following is false?

(i) Rock/paper/scissors has a dominant pure strategy

(ii) There is a Nash equilibrium for rock/paper/scissors

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 2.1: Which of the following is **false**?

(i) Rock/paper/scissors has a dominant pure strategy

(ii) There is a Nash equilibrium for rock/paper/scissors

- A. Neither
- **B. (i) but not (ii)**
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 2.1: Which of the following is **false**?

(i) Rock/paper/scissors has a dominant pure strategy

(ii) There is a Nash equilibrium for rock/paper/scissors

- A. Neither (i is false: easy to check that there's no deterministic dominant strategy)
- **B. (i) but not (ii)**
- C. (ii) but not (i) (i is false: easy to check that there's no deterministic dominant strategy)
- D. Both (There is a mixed strategy Nash Eq.)

Break & Quiz

Q 2.2: Which of the following is true

(i) Nash equilibria require each player to know other players' strategies

(ii) Nash equilibria require rational play

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 2.2: Which of the following is **true**

(i) Nash equilibria require each player to know other players' strategies

(ii) Nash equilibria require rational play

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- **D. Both**

Break & Quiz

Q 2.2: Which of the following is true

(i) Nash equilibria require each player to know other players' strategies

(ii) Nash equilibria require rational play

- A. Neither (See below)
- B. (i) but not (ii) (Rational play required: i.e., what if prisoners desire longer jail sentences?)
- C. (ii) but not (i) (The basic assumption of Nash equilibria is knowing all of the strategies involved)
- D. **Both**

Summary

- Intro to game theory
  - Characterize games by various properties
- Mathematical formulation for simultaneous games
  - Normal form, dominance, Nash equilibria, mixed vs pure