Zero Sum Games and the Minimax Thoerem

Aaron Roth

University of Pennsylvania

February 11 2021

Overview

- Today we'll dive into zero sum games.

Overview

- Today we'll dive into zero sum games.
- They have a very special property: the minimax theorem.

Overview

- Today we'll dive into zero sum games.
- They have a very special property: the minimax theorem.
- And a close connection to the polynomial weights algorithm (and related algorithms)

Overview

- Today we'll dive into zero sum games.
- They have a very special property: the minimax theorem.
- And a close connection to the polynomial weights algorithm (and related algorithms)
- Playing the polynomial weights algorithm in a zero sum game leads to equilibrium (a plausible dynamic!)

Overview

- Today we'll dive into zero sum games.
- They have a very special property: the minimax theorem.
- And a close connection to the polynomial weights algorithm (and related algorithms)
- Playing the polynomial weights algorithm in a zero sum game leads to equilibrium (a plausible dynamic!)
- In fact, we'll use it to prove the minimax theorem.

Zero Sum Games

Definition

A two player zero sum game is any two player game such that for every $a \in A_1 \times A_2$, $u_1(a) = -u_2(a)$.(i.e. at every action profile, the utilities sum to zero)

Zero Sum Games

Definition

A two player zero sum game is any two player game such that for every $a \in A_1 \times A_2$, $u_1(a) = -u_2(a)$.(i.e. at every action profile, the utilities sum to zero)

1. Strictly adversarial games: The only way for player 1 to improve his payoff is to harm player 2, and vice versa.

Zero Sum Games

Definition

A two player zero sum game is any two player game such that for every $a \in A_1 \times A_2$, $u_1(a) = -u_2(a)$.(i.e. at every action profile, the utilities sum to zero)

1. Strictly adversarial games: The only way for player 1 to improve his payoff is to harm player 2, and vice versa.
2. Closely related to linear programming, adversarial machine learning, and lots of other things.

Example

Consider the "Presidential Election Game":

	Morality	Tax-Cuts
Economy	(3,-3)	(-1,1)
Society	(-2,2)	(1,-1)

Example

Consider the "Presidential Election Game":

	Morality	Tax-Cuts
Economy	(3,-3)	(-1,1)
Society	(-2,2)	(1,-1)

Since utilities always sum to zero, we can more economically write the game by specifying only the row player's utility:

Example

Consider the "Presidential Election Game":

	Morality	Tax-Cuts
Economy	(3,-3)	(-1,1)
Society	(-2,2)	(1,-1)

Since utilities always sum to zero, we can more economically write the game by specifying only the row player's utility:

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

Example

Consider the "Presidential Election Game":

	Morality	Tax-Cuts
Economy	(3,-3)	(-1,1)
Society	(-2,2)	(1,-1)

Since utilities always sum to zero, we can more economically write the game by specifying only the row player's utility:

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

The row player (Max) wishes to *maximize* the utility. The column player (Min) wishes to *minimize* the utility.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. What if Max has to go first and announce his strategy, letting Min best respond?

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. What if Max has to go first and announce his strategy, letting Min best respond?
2. If Max announces mixed strategy $(1/2, 1/2)$, what should Min do?

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. What if Max has to go first and announce his strategy, letting Min best respond?
2. If Max announces mixed strategy $(1/2, 1/2)$, what should Min do?
3. Min should pick the action that minimizes her cost! She can compute:

$$\mathrm{E}[\mathrm{Morality}] = \frac{1}{2} \cdot 3 + \frac{1}{2} \cdot (-2) = \frac{1}{2}$$

$$\mathrm{E}[\mathrm{Tax-Cuts}] = \frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 1 = 0$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. What if Max has to go first and announce his strategy, letting Min best respond?
2. If Max announces mixed strategy $(1/2, 1/2)$, what should Min do?
3. Min should pick the action that minimizes her cost! She can compute:

$$\mathrm{E}[\mathrm{Morality}] = \frac{1}{2} \cdot 3 + \frac{1}{2} \cdot (-2) = \frac{1}{2}$$

$$\mathrm{E}[\mathrm{Tax-Cuts}] = \frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 1 = 0$$

4. So she plays Tax-cuts.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. More generally, if Max announces he is going to play according to $(p, 1-p)$...

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. More generally, if Max announces he is going to play according to $(p, 1-p)$...
2. Min will play the strategy that minimizes her cost:

$$\arg\min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. More generally, if Max announces he is going to play according to $(p, 1-p)$...
2. Min will play the strategy that minimizes her cost:

$$\arg\min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

3. So what should Max do, knowing this?

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. More generally, if Max announces he is going to play according to $(p, 1-p)$...
2. Min will play the strategy that minimizes her cost:

$$\arg\min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

3. So what should Max do, knowing this?
4. He should play:

$$\arg\max_p \min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. More generally, if Max announces he is going to play according to $(p, 1-p)$...
2. Min will play the strategy that minimizes her cost:

$$\arg\min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

3. So what should Max do, knowing this?
4. He should play:

$$\arg\max_p \min(p \cdot 3 - 2(1-p), p \cdot (-1) + (1-p))$$

5. And if Min goes first, she should play:

$$\arg\min_q \max(q \cdot 3 - (1-q), q \cdot (-2) + (1-q))$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. Going first in a zero sum game is only a *disadvantage*. It reveals your strategy and lets your opponent respond optimally.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. Going first in a zero sum game is only a *disadvantage*. It reveals your strategy and lets your opponent respond optimally.
2. But how much of a disadvantage? Lets see. What should Max do?

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. Going first in a zero sum game is only a *disadvantage*. It reveals your strategy and lets your opponent respond optimally.
2. But how much of a disadvantage? Lets see. What should Max do?
3. He should play $(p, 1-p)$ to equalize the cost of Min's two actions: i.e. he should pick p such that:

$$3p - 2(1-p) = -p + (1-p) \Leftrightarrow 5p - 2 = 1 - 2p \Leftrightarrow p = \frac{3}{7}$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. Going first in a zero sum game is only a *disadvantage*. It reveals your strategy and lets your opponent respond optimally.
2. But how much of a disadvantage? Lets see. What should Max do?
3. He should play $(p, 1-p)$ to equalize the cost of Min's two actions: i.e. he should pick p such that:

$$3p - 2(1-p) = -p + (1-p) \Leftrightarrow 5p - 2 = 1 - 2p \Leftrightarrow p = \frac{3}{7}$$

4. So Max should play $(3/7, 4/7)$...

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. Going first in a zero sum game is only a *disadvantage*. It reveals your strategy and lets your opponent respond optimally.
2. But how much of a disadvantage? Lets see. What should Max do?
3. He should play $(p, 1-p)$ to equalize the cost of Min's two actions: i.e. he should pick p such that:

$$3p - 2(1-p) = -p + (1-p) \Leftrightarrow 5p - 2 = 1 - 2p \Leftrightarrow p = \frac{3}{7}$$

4. So Max should play $(3/7, 4/7)$...
5. And when Min best responds, he gets payoff $1 - 2p = 1/7$.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?
2. Min should play so as to equalize the payoff of Max's two options.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?
2. Min should play so as to equalize the payoff of Max's two options.
3. i.e. she should play $(q, 1-q)$ such that:

$$3q - (1-q) = -2q + (1-q) \Leftrightarrow 4q - 1 = -3q + 1 \Leftrightarrow q = \frac{2}{7}$$

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?
2. Min should play so as to equalize the payoff of Max's two options.
3. i.e. she should play $(q, 1-q)$ such that:

$$3q-(1-q)=-2q+(1-q) \Leftrightarrow 4q-1=-3q+1 \Leftrightarrow q=\frac{2}{7}$$

4. So Min should play $(2/7, 5/7)$.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?
2. Min should play so as to equalize the payoff of Max's two options.
3. i.e. she should play $(q, 1-q)$ such that:

$$3q - (1-q) = -2q + (1-q) \Leftrightarrow 4q - 1 = -3q + 1 \Leftrightarrow q = \frac{2}{7}$$

4. So Min should play $(2/7, 5/7)$.
5. When Max best responds, Min suffers cost $4q - 1 = 1/7$.

How to Play This Game

	Morality	Tax-Cuts
Economy	3	-1
Society	-2	1

1. And what if Min goes first?
2. Min should play so as to equalize the payoff of Max's two options.
3. i.e. she should play $(q, 1-q)$ such that:

$$3q - (1-q) = -2q + (1-q) \Leftrightarrow 4q - 1 = -3q + 1 \Leftrightarrow q = \frac{2}{7}$$

4. So Min should play $(2/7, 5/7)$.
5. When Max best responds, Min suffers cost $4q - 1 = 1/7$.
6. Hmmm...

Hmmm

1. So in this case it was no disadvantage to go first...

Hmmm

1. So in this case it was no disadvantage to go first...
2. The $\max\min$ value of the game exactly equal to the $\min\max$ value of the game.

Hmmm

1. So in this case it was no disadvantage to go first...
2. The $\max\min$ value of the game exactly equal to the $\min\max$ value of the game.
3. This immediately implies that $(3/7, 4/7), (2/7, 5/7)$ is a Nash equilibrium of the game, since both players are getting payoff/cost $1/7$, and we have just derived that both are best responding to one another.

Hmmm

1. So in this case it was no disadvantage to go first...
2. The max min value of the game exactly equal to the min max value of the game.
3. This immediately implies that $(3/7, 4/7), (2/7, 5/7)$ is a Nash equilibrium of the game, since both players are getting payoff/cost $1/7$, and we have just derived that both are best responding to one another.
4. Lets investigate further...

Order of Play

We use the notation $[n] = \{1, 2, \ldots, n\}$, and $\Delta[n]$ to denote the set of probability distributions over $[n]$:

Order of Play

We use the notation $[n] = \{1, 2, \ldots, n\}$, and $\Delta[n]$ to denote the set of probability distributions over $[n]$:

Definition

For an $n \times m$ matrix U (think about this as the payoff matrix in a two player zero sum game if you like):

$$\max\min(U) = \max_{p \in \Delta[n]} \min_{y \in [m]} \sum_{i=1}^{n} p_i \cdot U(i, y)$$

$$\min\max(U) = \min_{q \in \Delta[m]} \max_{x \in [n]} \sum_{j=1}^{m} q_j \cdot U(x, j)$$

Order of Play

We use the notation $[n] = \{1, 2, \ldots, n\}$, and $\Delta[n]$ to denote the set of probability distributions over $[n]$:

Definition

For an $n \times m$ matrix U (think about this as the payoff matrix in a two player zero sum game if you like):

$$\max\min(U) = \max_{p \in \Delta[n]} \min_{y \in [m]} \sum_{i=1}^{n} p_i \cdot U(i, y)$$

$$\min\max(U) = \min_{q \in \Delta[m]} \max_{x \in [n]} \sum_{j=1}^{m} q_j \cdot U(x, j)$$

If U is a zero sum game, then $\max\min(U)$ represents the payoff that Max can guarantee if he goes first, and $\min\max(U)$ represents the payoff that he can guarantee if Min goes first.

The Minimax Theorem

Recall going first is not an advantage. In math, for any game U:

$$\min\max(U) \geq \max\min(U)$$

The Minimax Theorem

Recall going first is not an advantage. In math, for any game U:

$$\min\max(U) \geq \max\min(U)$$

It turns out that in a zero sum game, it is also not a disadvantage:

Theorem (Von Neumann)

In any zero sum game U:

$$\min\max(U) = \max\min(U)$$

The Minimax Theorem

Recall going first is not an advantage. In math, for any game U:

$$\min\max(U) \geq \max\min(U)$$

It turns out that in a zero sum game, it is also not a disadvantage:

Theorem (Von Neumann)

In any zero sum game U:

$$\min\max(U) = \max\min(U)$$

Corollary

In any Nash equilibrium of a zero sum game, Max plays a maxmin strategy and Min plays a minmax strategy. Note that these can be computed without needing to reason about what the other player is doing.

The Minimax Theorem

Recall going first is not an advantage. In math, for any game U:

$$\min\max(U) \geq \max\min(U)$$

It turns out that in a zero sum game, it is also not a disadvantage:

Theorem (Von Neumann)

In any zero sum game U:

$$\min\max(U) = \max\min(U)$$

Corollary

In any Nash equilibrium of a zero sum game, Max plays a maxmin strategy and Min plays a minmax strategy. Note that these can be computed without needing to reason about what the other player is doing.

Corollary

All Nash equilibria in Zero sum games have the same payoff – the $\max\min$ *value of the game.*

The Minimax Theorem

1. None of these things are true of general games. (Consider Battle of the Sexes)

The Minimax Theorem

1. None of these things are true of general games. (Consider Battle of the Sexes)
2. It means Zero-sum games are easy to play: no need for counter-speculation.

The Minimax Theorem

1. None of these things are true of general games. (Consider Battle of the Sexes)
2. It means Zero-sum games are easy to play: no need for counter-speculation.
3. Its non-obvious. Von Neumann gave a complicated proof in 1928, writing "As far as I can see, there could be no theory of games ... without that theorem ... I thought there was nothing worth publishing until the Minimax Theorem was proved"

The Minimax Theorem

1. None of these things are true of general games. (Consider Battle of the Sexes)
2. It means Zero-sum games are easy to play: no need for counter-speculation.
3. Its non-obvious. Von Neumann gave a complicated proof in 1928, writing "As far as I can see, there could be no theory of games ... without that theorem ... I thought there was nothing worth publishing until the Minimax Theorem was proved"
4. Previously, Borell had proven it for the special case of 5×5 matrices, and thought it was false for larger matrices.

The Minimax Theorem

1. None of these things are true of general games. (Consider Battle of the Sexes)
2. It means Zero-sum games are easy to play: no need for counter-speculation.
3. Its non-obvious. Von Neumann gave a complicated proof in 1928, writing "As far as I can see, there could be no theory of games ... without that theorem ... I thought there was nothing worth publishing until the Minimax Theorem was proved"
4. Previously, Borell had proven it for the special case of 5×5 matrices, and thought it was false for larger matrices.
5. But well give an easy, constructive proof.

Proof

1. Suppose the theorem were false: there is some game U for which $\min\max(U) > \max\min(U)$.

Proof

1. Suppose the theorem were false: there is some game U for which $\min\max(U) > \max\min(U)$.
2. Write $v_1 = \min\max(U)$ and $v_2 = \max\min(U)$ (And so $v_1 = v_2 + \epsilon$ for some constant $\epsilon > 0$).

Proof

1. Suppose the theorem were false: there is some game U for which $\min\max(U) > \max\min(U)$.
2. Write $v_1 = \min\max(U)$ and $v_2 = \max\min(U)$ (And so $v_1 = v_2 + \epsilon$ for some constant $\epsilon > 0$).
3. In other words: if Min has to go first, then Max can guarantee payoff at least v_1, but if Max is forced to go first, then Min can force Max to have payoff only v_2.

Proof: A Thought Experiment

Lets consider what happens when Min and Max repeatedly play against each other as follows, for T rounds:

1. Min will play using the polynomial weights algorithm. i.e. at each round t, the weights w^t of the polynomial weights algorithm will form her mixed strategy, and she will sample an action at random from this distribution, updating based on the losses she experiences at that round.
2. Max will play the best response to Min's strategy. i.e. Max will play $x^t = \arg\max_x \mathrm{E}_{y \sim w^t}[U(x, y)]$.

Proof: A Thought Experiment

Lets consider what happens when Min and Max repeatedly play against each other as follows, for T rounds:

1. Min will play using the polynomial weights algorithm. i.e. at each round t, the weights w^t of the polynomial weights algorithm will form her mixed strategy, and she will sample an action at random from this distribution, updating based on the losses she experiences at that round.
2. Max will play the best response to Min's strategy. i.e. Max will play $x^t = \arg\max_x \mathrm{E}_{y \sim w^t}[U(x, y)]$.

What we know about each player's average payoffs when they play in this manner?

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

$$\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] \quad \leq \quad \frac{1}{T}\min_{y^*}\sum_{t=1}^{T}U(x^t, y^*) + \Delta(T)$$

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] &\leq \frac{1}{T}\min_{y^*}\sum_{t=1}^{T}U(x^t, y^*) + \Delta(T) \\
&= \min_{y^*}\sum_{t=1}^{T}\frac{1}{T}U(x^t, y^*) + \Delta(T)
\end{aligned}$$

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] &\leq \frac{1}{T}\min_{y^*}\sum_{t=1}^{T}U(x^t, y^*) + \Delta(T)\\ &= \min_{y^*}\sum_{t=1}^{T}\frac{1}{T}U(x^t, y^*) + \Delta(T)\\ &= \min_{y^*}\mathrm{E}_{x\sim\bar{x}}[U(x, y^*)] + \Delta(T)\end{aligned}$$

$\bar{x}$ is the mixed strategy that puts weight $1/T$ on each action x^t.
$\Delta(T)$ is the regret bound of the polynomial weights algorithm:

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] &\leq \frac{1}{T}\min_{y^*}\sum_{t=1}^{T} U(x^t, y^*) + \Delta(T) \\
&= \min_{y^*}\sum_{t=1}^{T}\frac{1}{T}U(x^t, y^*) + \Delta(T) \\
&= \min_{y^*}\mathrm{E}_{x\sim\bar{x}}[U(x, y^*)] + \Delta(T)
\end{aligned}$$

$\bar{x}$ is the mixed strategy that puts weight $1/T$ on each action x^t.
$\Delta(T)$ is the regret bound of the polynomial weights algorithm:

$$\Delta(T) = 2\sqrt{\frac{\log n}{T}}.$$

Proof: A Thought Experiment

We know from the guarantee of the polynomial weights algorithm:

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] &\leq \frac{1}{T}\min_{y^*}\sum_{t=1}^{T} U(x^t, y^*) + \Delta(T) \\ &= \min_{y^*}\sum_{t=1}^{T}\frac{1}{T}U(x^t, y^*) + \Delta(T) \\ &= \min_{y^*}\mathrm{E}_{x\sim\bar{x}}[U(x, y^*)] + \Delta(T)\end{aligned}$$

$\bar{x}$ is the mixed strategy that puts weight $1/T$ on each action x^t.
$\Delta(T)$ is the regret bound of the polynomial weights algorithm:

$$\Delta(T) = 2\sqrt{\frac{\log n}{T}}.$$

By definition, $\min_{y^*}\mathrm{E}_{x\sim\bar{x}}U(x, y^*) \leq \max\min(U) = v_2$ and so:

$$\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] \leq v_2 + \Delta(T)$$

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

$$\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] \quad = \quad \frac{1}{T}\sum_{t=1}^{T}\max_{x^*}\mathrm{E}_{y\sim w^t}[U(x^*, y)]$$

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t, y^t)] &= \frac{1}{T}\sum_{t=1}^{T}\max_{x^*}\mathrm{E}_{y\sim w^t}[U(x^*, y)] \\
&\geq \frac{1}{T}\sum_{t=1}^{T} v_1
\end{aligned}$$

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T} \mathrm{E}[U(x^t, y^t)] &= \frac{1}{T}\sum_{t=1}^{T} \max_{x^*} \mathrm{E}_{y\sim w^t}[U(x^*, y)] \\
&\geq \frac{1}{T}\sum_{t=1}^{T} v_1 \\
&= v_1
\end{aligned}$$

Combining these inequalities:

$$v_1 \leq v_2 + \Delta(T)$$

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

$$
\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T}\mathrm{E}[U(x^t,y^t)] &= \frac{1}{T}\sum_{t=1}^{T}\max_{x^*}\mathrm{E}_{y\sim w^t}[U(x^*,y)] \\
&\geq \frac{1}{T}\sum_{t=1}^{T} v_1 \\
&= v_1
\end{aligned}
$$

Combining these inequalities:

$$v_1 \leq v_2 + \Delta(T)$$

Since $v_1 = v_2 + \epsilon$:

$$\epsilon \leq \Delta(T)$$

Proof: A Thought Experiment

But: on each day t Max is best responding to Min's mixed strategy w^t. So...

$$\begin{aligned}
\frac{1}{T}\sum_{t=1}^{T} \mathrm{E}[U(x^t, y^t)] &= \frac{1}{T}\sum_{t=1}^{T} \max_{x^*} \mathrm{E}_{y\sim w^t}[U(x^*, y)] \\
&\geq \frac{1}{T}\sum_{t=1}^{T} v_1 \\
&= v_1
\end{aligned}$$

Combining these inequalities:

$$v_1 \leq v_2 + \Delta(T)$$

Since $v_1 = v_2 + \epsilon$:

$$\epsilon \leq \Delta(T)$$

Taking T large enough leads to contradiction.

Reflection

1. An amazing feature of Polynomial Weights: It guarantees that no matter what, you do as well *as if you had gotten to observe your opponent's strategy, and then best respond after the fact*.

Reflection

1. An amazing feature of Polynomial Weights: It guarantees that no matter what, you do as well *as if you had gotten to observe your opponent's strategy, and then best respond after the fact*.
2. Using the polynomial weights algorithm guarantees that a player using gets payoff quickly approaching the (unique) value of the game.

Reflection

1. An amazing feature of Polynomial Weights: It guarantees that no matter what, you do as well *as if you had gotten to observe your opponent's strategy, and then best respond after the fact*.
2. Using the polynomial weights algorithm guarantees that a player using gets payoff quickly approaching the (unique) value of the game.
3. It does so *without needing to know what the game is*. The game matrix is not an input to the PW algorithm!

Reflection

1. An amazing feature of Polynomial Weights: It guarantees that no matter what, you do as well *as if you had gotten to observe your opponent's strategy, and then best respond after the fact*.
2. Using the polynomial weights algorithm guarantees that a player using gets payoff quickly approaching the (unique) value of the game.
3. It does so *without needing to know what the game is*. The game matrix is not an input to the PW algorithm!
4. The only information needed is the realized payoffs are for the actions as it plays the game.

Thanks!

See you next class — stay healthy, and wear a mask!