Zero-Sum Games

Jay Daigle

jaydaigle@gwu.edu

https://jaydaigle.net/politics

The George Washington University

March 26, 2026

Battle of the Bismarck Sea

- In WWII, the Japanese Navy needed to resupply New Guinea
- Had to sail around New Britain, either North or South path
- US controlled New Britain, wanted to bomb Japanese Convoy

US Strategy

- US can focus its search to the North or the South
- If the Japanese go North:
  - Guess North: 2 days of bombing
  - Guess South: 1 day of bombing
- If the Japanese go South:
  - Guess North: 2 days of bombing
  - Guess South: 3 days of bombing
- What should the US do?

Conflict and Game Theory

- Game theory models strategic interactions
- How do you make choices when other people with different goals are also making choices?
- Developed in early 20th century
- Reached prominence after World War II:
  - Avoid another world war
  - Avoid nuclear war
- Summarize strategic situation in a mathematical model
- Find optimal strategy or strategies
- Can we find ways to cooperate?

Zero-Sum Games

Definition

- A two-person zero-sum game is a game featuring two players, in which
  - Each player adopts a strategy
  - The combination of strategies determines a number called the payoff.
- We can think of the payoff as the amount of money player 2 has to pay player 1
- If they payoff is negative, this means player 1 has to pay player 2.

Zero-Sum Games

- Can represent a two-player zero-sum game with a matrix, which is a grid of possible payoffs
- We call player 1 Row and player 2 Column
- If Row has m choices and Column has n choices, we get a $m \times n$ matrix
- The entry on row i, column j is $u_{i,j}$

	$j = 1$	$j = 2$	$j = 3$
$i = 1$	$u_{1,1}$	$u_{1,2}$	$u_{1,3}$
$i = 2$	$u_{2,1}$	$u_{2,2}$	$u_{2,3}$
$i = 3$	$u_{3,1}$	$u_{3,2}$	$u_{3,3}$

Battle of the Bismarck Sea

- Two players: US and Japan
- Zero-sum: US wants maximum bombing and Japan wants minimum

		Japan (Column)	
		North	South
United States (Row)	North	2	2
	South	1	3

Roshambo

Example (Roshambo, or Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

- Positive numbers are good for player 1 Row
- Negative numbers are good for player 2 Column

Naive Strategies

Definition

- The outcome that gives a player their best possible payoff is the primary outcome.
- For Row this is the largest entry in the payoff matrix
- For Column it is the smallest, or most negative, entry.

Definition

- In the naive method, a player chooses the strategy corresponding to their primary outcome.
- This is called the player's naive strategy, or sometimes the greedy strategy or optimistic strategy.
- If both players play their naive strategies, we get the doubly naive outcome.

Naive Strategies

Example (The Battle of the Bismarck Sea)

		Japan (Column) North	Japan (Column) South
United States (Row)	North	2	2
	South	1	3

Primary outcome for USA

- 3 days of bombing
- Searching to the South

Primary outcome for Japan

- 1 day of bombing
- Going to the North

- Doubly naive outcome: one day of bombing

Naive Strategies

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Primary outcome for Row

- Win, 1 point
- Any strategy

Primary outcome for Column

- Loss, -1 point
- Any strategy

- Doubly naive outcome: Any??

Prudent Strategies

Definition

- The worst payoff a player can get from a given strategy is that strategy's guarantee.
- (The player is guaranteed to get at least that good a payoff.)

Definition

- In the prudent method, a player chooses the strategy with the best guarantee, called the prudent strategy.
- We might also call this the pessimistic strategy, because it wants to minimize the damage from the worst-case scenario.
- If both players play their prudent strategies, we get the doubly prudent outcome.

Prudent Strategies

Example (The Battle of the Bismarck Sea Min-Max Diagram)

Japan (Column)

		North	South	
United States (Row)	North	2	2	2
	South	1	3	1
		2	3	

Prudent strategy for USA

- Search North

Prudent strategy for Japan

- Go to the North

- Doubly prudent outcome: two days of bombing

Naive and Prudent Strategies

Naive strategy

- Maximax: maximize your maximum payoff.
- "Hope chess": hope your opponent plays into your hands
- Dangerous if your opponent is smart!

Prudent strategy

- Minimax: minimize your maximum loss
- maximin: Maximize your minimum payoff
- Non-obvious mathematical conclusion: those two ideas are the same.

Prudent strategies

Proposition

- *Let r be the guarantee of Row's prudent strategy, and c be the guarantee of Column's prudent strategy.*
- *If Row plays a prudent strategy, their payoff will be at least r*
- *But if Row plays a non-prudent strategy, it is possible their payoff will be lower than r.*
- *Similarly, if Column plays a prudent strategy, their payoff will be at most c*
- *But if Column plays a non-prudent strategy, it is possible their payoff will be greater than c.*
- *(Remember Column prefers smaller payoffs)*

Prudent strategies

Proposition

- *If Row plays a prudent strategy their payoff will be at least r*
- *If Column plays strategy their payoff will be at most c*

Proof.

- Suppose row i and column j are prudent strategies
- The payoff to this pair of strategies is $u_{i,j}$
- Since i is prudent, we must have $u_{i,j} \geq r$
- Since j is prudent, we must have $u_{i,j} \leq c$

□

Prudent strategies

Proposition

- *If Row plays a prudent strategy their payoff will be at least r*
- *If Column plays prudent strategy their payoff will be at most c*

Corollary

If r is the guarantee of Row's prudent strategy and c is the guarantee of Column's prudent strategy, then $r \leq c$.

Proof.

- If row i and column j are prudent, then
  - $u_{i,j} \geq r$
  - $u_{i,j} \leq c$.
- Thus $r \leq u_{i,j} \leq c$.

□

Best Response

Discussion Question

If you know what your opponent will do, how should that affect your choices?

Definition

- A strategy choice by one player is called the best response to an opponent strategy if it gives the best payoff against that strategy.
- The best response to a naive strategy is called the counter-naive strategy.
- The best response to a prudent strategy is called the counter-prudent strategy.

Best Response

Example

		Japan (Column)	
		North	South
United States (Row)	North	2	2
	South	1	3

- If Japan goes North, US should search North
- If Japan goes South, US should search South
- If US searches South, Japan should go North
- If US searches North, Japan should . . . ?

Best Response and Flow Diagrams

- Summarize our best responses with a flow diagram
- A map that gives us the best response to any strategy
- Each column: vertical arrow points to *largest* entry (or entries)
- Each row: horizontal arrow points to *smallest* entry

Example (Battle of the Bismarck Sea)

	North	South
North	2	2
South	1	3

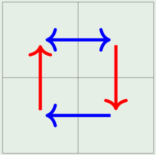

Best Response and Flow Diagrams

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

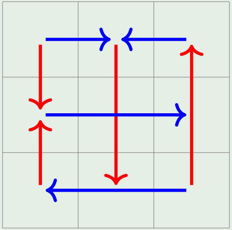

Backward Induction

- If I know what you'll do, I can choose the best response
- But if you know that, you can choose the best response to my choice
- But if I know that, I can choose the best response to that choice
- But if you know that, you can . . .
- We call this reasoning process backward induction
- Where does it stop?

Backward Induction: Rock Paper Scissors

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

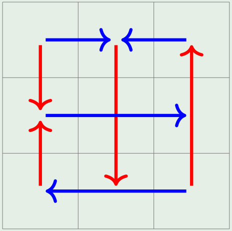

- If Row plays Rock, then Column should play Paper
- But if Column plays Paper, then Row should play Scissors
- But if Row plays Scissors, then Column should play Rock
- But if . . .

Backward Induction: Bismarck Sea

Example (Battle of the Bismarck Sea)

	North	South
North	2	2
South	1	3

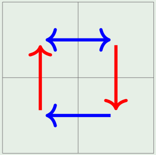

- If Japan goes South, then the US should go South
- But if the US goes South, then Japan should go North
- But if Japan goes North, the US should go North
- But if the US goes North . . .
- Japan is fine sticking with North!

Saddle Points

Definition

- A saddle point is an outcome such that the strategy for each player is the best response to the strategy of the opponent, simultaneously.
- A saddle point strategy is a strategy that corresponds to a saddle point outcome.

- An outcome is a saddle point if and only if all the arrows in its row and column point to it.
- The point in row k, column ℓ is a saddle point if

$$u_{k,\ell} \leq u_{k,j} \qquad \text{for any column } j$$
$$u_{k,\ell} \geq u_{i,\ell} \qquad \text{for any row } i.$$

Saddle Points

Theorem

- *A two-person zero-sum game has a saddle point if and only if* $r = c$*.*
- *In that case, the saddle point is a doubly prudent outcome, and the payoff is* $r = c$*.*

Proof.

- "If and only if": need to prove two things.
- Assume $r = c$, and prove that the game has a saddle point
- Assume the game has a saddle point, and prove $r = c$.

Saddle Points

Theorem

- *A two-person zero-sum game has a saddle point if and only if $r = c$.*
- *In that case, the saddle point is a doubly prudent outcome, and the payoff is $r = c$.*

Proof.

- Suppose $r = c$.
- Let row k and column ℓ have guarantee $r = c$
- $r \leq u_{k,j}$ for all j, so ℓ is a best response to row k.
- $r \geq u_{i,\ell}$ for all i, so j is a best response to column ℓ.
- Thus (k, ℓ) is a saddle point.

Saddle Points

Theorem

- *A two-person zero-sum game has a saddle point if and only if $r = c$.*
- *In that case, the saddle point is a doubly prudent outcome, and the payoff is $r = c$.*

Proof.

- Conversely, suppose (k, ℓ) is a saddle point
- $u_{k,\ell}$ must be the smallest entry in row k,
  - Otherwise there'd be a better response than ℓ.
- Thus $u_{k,\ell}$ is the guarantee of row k
- But r is the largest guarantee of any row, so $r \geq u_{k,\ell}$.

Saddle Points

Theorem

- *A two-person zero-sum game has a saddle point if and only if* $r = c$.
- *In that case, the saddle point is a doubly prudent outcome, and the payoff is* $r = c$.

Proof.

- Suppose (k, ℓ) is a saddle point
- $u_{k,\ell}$ must be the largest entry in column ℓ
- Thus $u_{k,\ell}$ is the guarantee of column ℓ
- But $r = c$ is the smallest guarantee of any column, so $r \leq u_{k,\ell}$.

Saddle Points

Theorem

- *A two-person zero-sum game has a saddle point if and only if $r = c$.*
- *In that case, the saddle point is a doubly prudent outcome, and the payoff is $r = c$.*

Proof.

- Suppose (k, ℓ) is a saddle point
- Because this is the smallest entry in row k, saw that $r \geq u_{k,\ell}$
- Because this is the largest entry in column ℓ, saw that $r \leq u_{k,\ell}$.
- Can only both be true if $r = u_{k,\ell}$.