C&O 355
Mathematical Programming
Fall 2010
Lecture 12

N. Harvey

Daniel Kleitman

First person with finite Erdos-Bacon number?

- What movie is this?
- What is the location for this scene?
- Who are the people in this scene?

Two-Player "Zero-sum" Games

- Alice and Bob play the following game
  - Alice has m possible actions she can make
  - Bob has n possible actions he can make
  - They both simultaneously choose an action
  - If Alice chooses action i and Bob chooses action j then Bob pays Alice M_{ij} dollars
- "Zero-sum" means amount Alice wins equals amount Bob loses

Bob

Alice

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

M

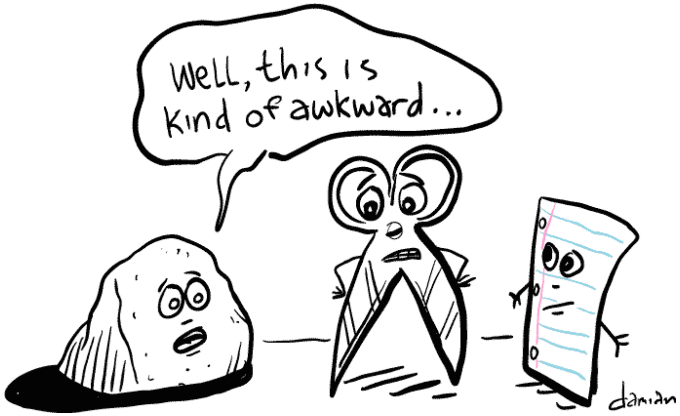

http://www.toddlittleton.net/

Randomized Strategies

- As we know from Rock-Scissors-Paper, we seek **randomized** choices of actions that optimize the **expected** payoff
- Suppose Alice chooses a probability distribution $x \in \mathbb{R}^m$ over her actions and Bob chooses a distribution $y \in \mathbb{R}^n$ over his actions. (So $x \geq 0$, $\sum_i x_i = 1$, $y \geq 0$, $\sum_i y_i = 1$.)
- The **expected** amount Bob pays Alice is
$$\sum_i \sum_j M_{i,j} x_i y_j = x^T M y$$
- Alice is paranoid and thinks Bob might know her distribution x. She wants x that pays well **no matter which** y Bob chooses, and even if Bob knows x.
- So Alice wants to solve: $\max_x \min_y x^T M y$
(over all distributions x and y)
- Similarly, Bob wants to solve: $\min_y \max_x x^T M y$

Von Neumann's Theorem

- We know: $\max_x \min_y \; x^T M y \le \min_y \max_x \; x^T M y$ (Asst 2, Q6)
  (over all distributions x and y)
- Does equality hold?
  In other words, is Alice's paranoid strategy actually an optimal strategy?

- **Theorem:** [Von Neumann, 1926] Yes! There exist x^* and y^* s.t.
  $$\max_x \min_y \; x^T M y \; = \; x^{*T} M y^* \; = \; \min_y \max_x \; x^T M y$$
- Can be proven using LP duality. (Similar to Asst 2, Q7.)
- Conversely, LP duality can be proven by Von Neumann's Thm
- The strategies x^* and y^* are called **Nash equilibria**
  Nash's work came later, showing equilibria in more general, non-zero-sum games.

John Nash,
Nobel Laureate

Hollywood's
depiction of
John Nash

Saddle Points

- It is **not** generally true that

$$\sup_{x \in P} \inf_{y \in Q} f(x,y) = \inf_{y \in Q} \sup_{x \in P} f(x,y).$$

- Consider P = [-1,1], Q = $\mathbb{R}$ and f(x,y) = x^2 + xy

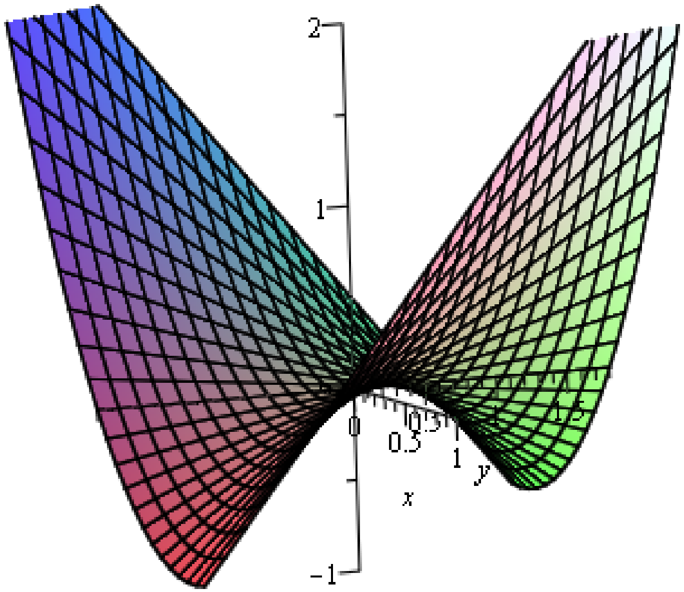

Saddle Points

- It is **not** generally true that

$$\sup_{x \in P} \inf_{y \in Q} f(x,y) = \inf_{y \in Q} \sup_{x \in P} f(x,y).$$

- Consider P = [-1,1], Q = $\mathbb{R}$ and f(x,y) = x^2 + xy
- Looking along the y-axis,

$$\sup_{x \in P} \inf_{y \in Q} f(x,y) = 0$$

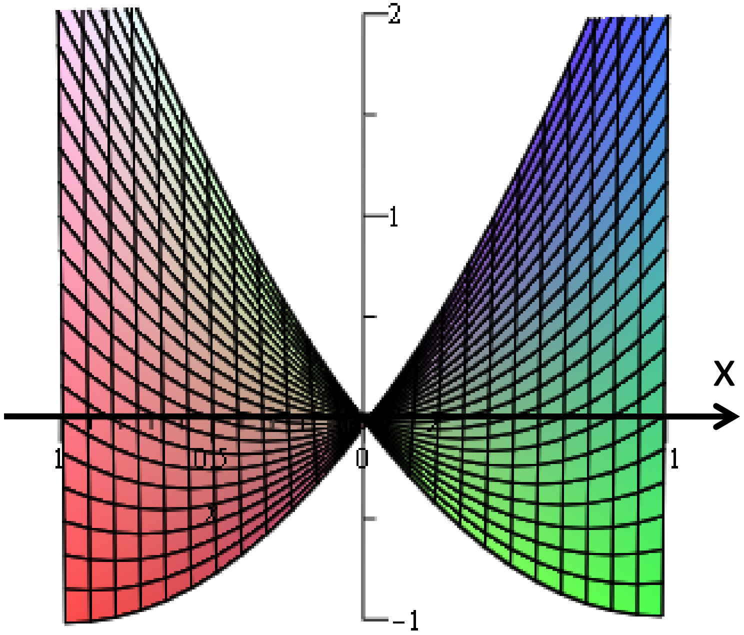

Saddle Points

- It is **not** generally true that

$$\sup_{x \in P} \inf_{y \in Q} f(x,y) = \inf_{y \in Q} \sup_{x \in P} f(x,y).$$

- Consider P = [-1,1], Q = $\mathbb{R}$ and f(x,y) = x^2 + xy
- Looking along the y-axis,

$$\sup_{x \in P} \inf_{y \in Q} f(x,y) = 0$$

- Looking along the x-axis,

$$\inf_{y \in Q} \sup_{x \in P} f(x,y) = 1$$

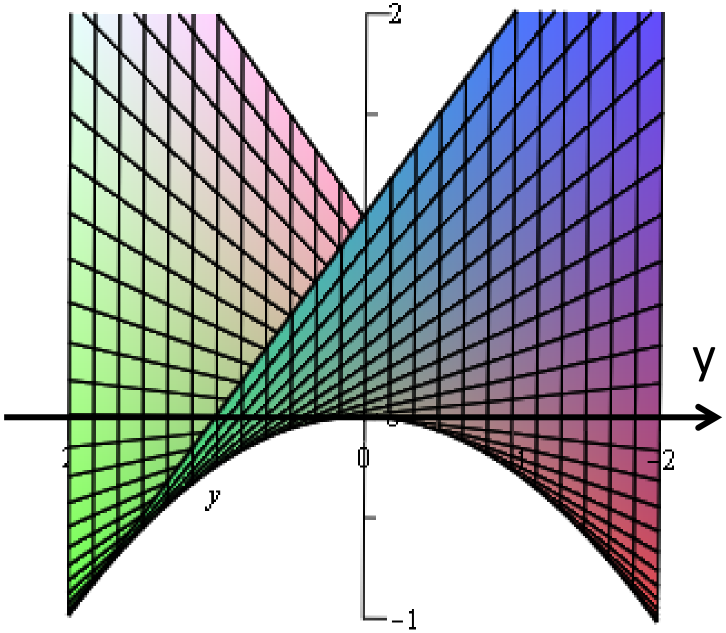

Saddle Points

- It is **not** generally true that

$$\sup_{x\in P}\inf_{y\in Q} f(x,y) = \inf_{y\in Q}\sup_{x\in P} f(x,y).$$

- $x^*\in P$ and $y^*\in Q$ are called **saddle points** if

$$\sup_{x\in P}\inf_{y\in Q} f(x,y) = f(x^*,y^*) = \inf_{y\in Q}\sup_{x\in P} f(x,y).$$

- The question **"when do saddle points exist?"** is closely related to the question **"when does a non-linear program satisfy strong duality?"**
  - Nash Equilibria of zero-sum games are saddle points
  - Note: we only defined duals of linear programs

Learning a Near-Optimal Strategy for Alice

- Alice maintains weights on her actions $w \in \mathbb{R}^m$, $w \geq 0$
- For T rounds
  - Alice normalizes w to get distribution $x = w / \sum_i w_i$
  - Alice tells Bob she's going to play with distribution x
  - Bob chooses an action j that is optimal against x
  - Alice plays a random action chosen according to x, and Bob plays j
  - Alice increases the weight w_i if $M_{i,j} > 0$ and decreases the weight w_i if $M_{i,j} < 0$

- **Amazing Fact:** If this is done carefully, then x is now a nearly optimal strategy for Alice

Multiplicative Weights Update Method

- Alice increases and decreases her weights by a small **multiplicative factor**
- This method is useful for many problems
  - fast algorithms (max flow, ...)
  - machine learning (AdaBoost, Winnow)
  - complexity theory (Yao's XOR Lemma)
  - quantum computing (QIP = PSPACE)
  - derandomization (pessimistic estimators, ...)
  - online optimization
  - ...

```
procedure FindEquilibrium($M$, $\delta$)
input: An $m \times n$ matrix $M$ and desired error $\delta$. Assume $M_{i,j} \in [-1, 1]\ \forall i, j$.
output: Distributions $\hat{x} \in \mathbb{R}^m$ and $\hat{y} \in \mathbb{R}^n$ satisfying

    $\min_y \hat{x}^\top M y \ \geq\ v - \delta \qquad \text{and} \qquad \max_x x^\top M \hat{y} \ \leq\ v + \delta,$

where $v$ is the value of the game.
Set $\epsilon = \delta/3$ and $T = (\ln m)/\epsilon^2$
Set $w_i^{(1)} = 1$ for every $i = 1, \dots, m$
For $t = 1, \dots, T$
    Set $x^{(t)} = w^{(t)} / \sum_{i=1}^m w_i^{(t)}$
    Let $j^{(t)}$ be a value of $j$ minimizing $(x^{(t)} M)_j$
    Let $y^{(t)}$ be the vector with 1 in coordinate $j^{(t)}$ and other coordinates 0
    Set
        $w_i^{(t+1)} \ = \ \begin{cases} w_i^{(t)} \cdot (1+\epsilon)^{M_{i,j^{(t)}}} & (\text{if } M_{i,j^{(t)}} \geq 0) \\ w_i^{(t)} \cdot (1-\epsilon)^{-M_{i,j^{(t)}}} & (\text{if } M_{i,j^{(t)}} < 0) \end{cases}$

Set $\hat{x} = \sum_{t=1}^T x^{(t)}/T$ and $\hat{y} = \sum_{t=1}^T y^{(t)}/T$
Return $\hat{x}$ and $\hat{y}$
```