CSCI 357: Algorithmic Game Theory

Lecture 3: Game Theory II

Shikha Singh


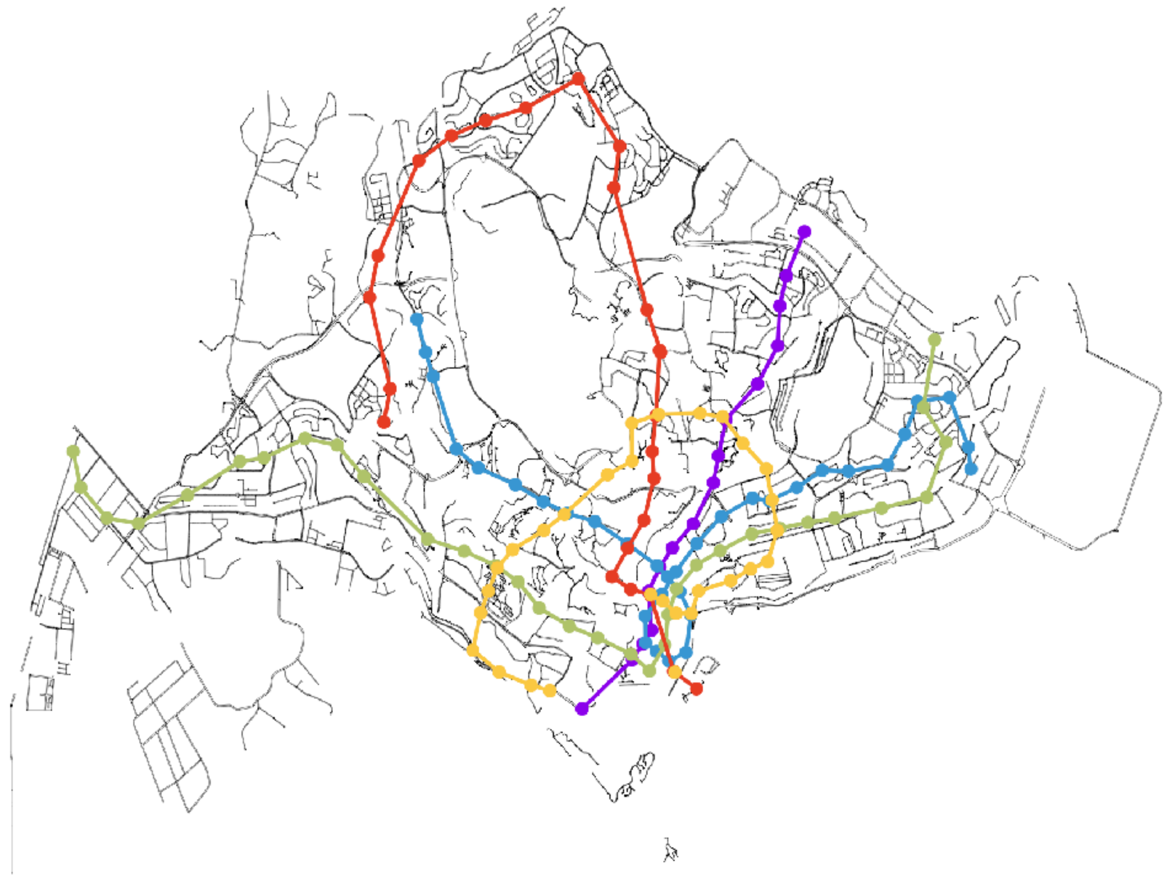

Announcements and Logistics

- **Assignment 1** due Thursday 10 pm
  - Office hours 4-5 pm (TCL 304), Tai's hours: 7.30-9 pm (TCL 206)
  - Submission on Gradescope: assign questions to pages
- Course calendar mystery (solved)
  - Williams email (logged in/ default) to view on course website
  - Alternate option: add to your google calendar directly
- **Assignment 2** will be released tomorrow
  - Questions based on today's + Monday's lecture

Questions?

Last Time

- Solution concepts in game theory:
  - Dominant strategy equilibrium, Pure Nash equilibrium
- Iterated elimination of dominated actions
  - If it leads to a single outcome: must be a pure Nash (HW 1)
  - Not always possible
  - Computationally expensive (HW 1)
  - Application: 2/3rds of average game
- Any questions/comments on these?

Today

- Today we will cover some more topics in game theory
  - Mixed-strategy Nash equilibrium
  - Downsides of pure Nash/ mixed Nash
  - Complexity of finding equilibrium
- Towards mechanism design
  - High-level of how its different and the challenges
- Next week: Auctions

Mixed-Strategy Nash

No Pure Nash Equilibrium

- Consider the following matching pennies game:
  - P_1 and P_2 choose heads or tails simultaneously and independently of each other
  - If they both pick the same action: P_1 wins
  - Otherwise, P_2 wins
- Games like these are called **zero-sum games**
- What happens if players play "pure" strategies?
- How would *you* play such a game?

	Head	*Tail*
Head	$1, -1$	$-1, \quad 1$
Tail	$-1, \quad 1$	$1, -1$

Mixed Strategies

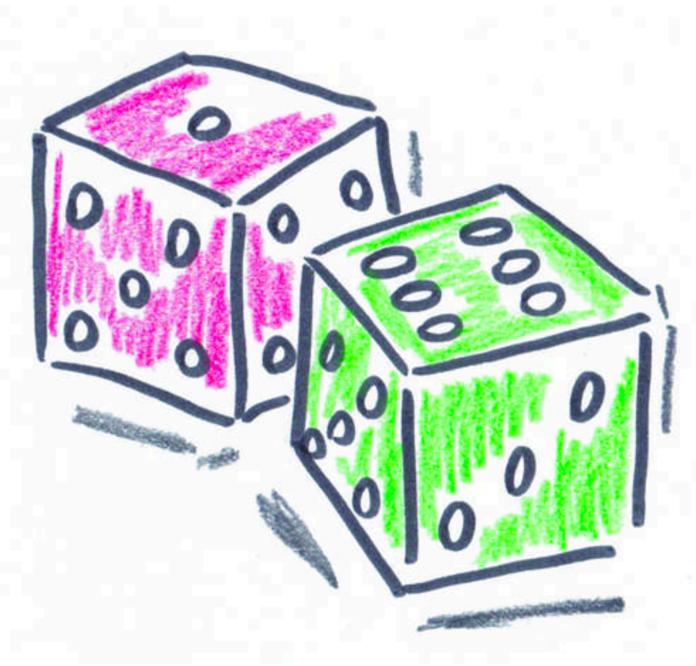

- Randomize to create uncertainty
- That is the idea behind mixed strategies
- Player strategies are now a probability distribution over actions
    - E.g. P_1 picks H with prob $q \in [0,1]$, T with prob $1 - q$
    - Similarly, P_2 picks H with prob $p \in [0,1]$, T with prob $1 - p$
- Overall strategy profile is a probability distribution over the sample space $S = \{(H,H),(H,T),(T,T),(T,H)\}$
- Player's want to maximize their "expected utility"

		p	$1 - p$
		$Head$	$Tail$
q	$Head$	$1, -1$	$-1, \quad 1$
$1 - q$	$Tail$	$-1, \quad 1$	$1, -1$

Expected Utility

- **Expected utility.** Let $p(a)$ denote the probability assigned to action profile a by strategy profile s, then the expected utility to player i for strategy profile s is:
  - $$u_i(s) = \sum_{a \in A} u_i(a) \cdot \Pr(a \mid s)$$

  where $\Pr(a \mid s) = s_1(a) \cdot s_2(a) \ \cdots \ s_n(a)$

Joint probability that this action profile is played; uses **independence**

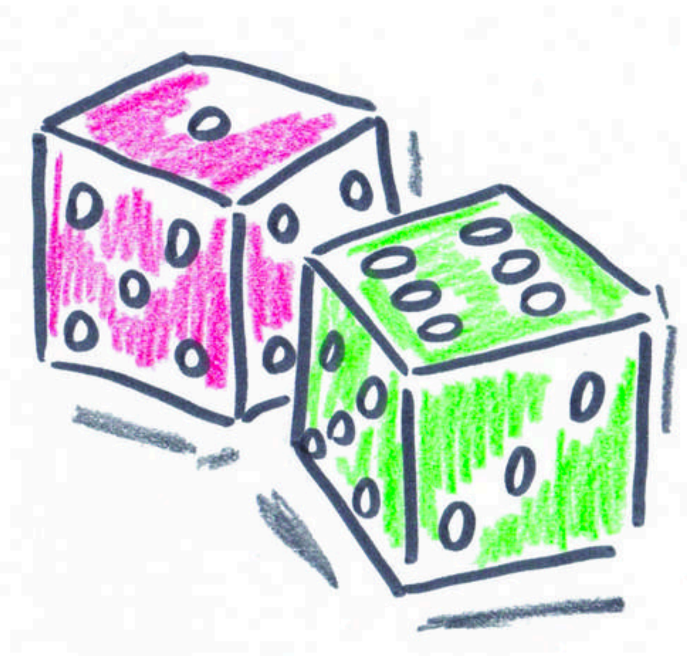

- For example, let $s = ((0.5{,}0.5), (0.5{,}0.5))$
- Then, $u_1 = 1 \cdot 1/2 \cdot 1/2 + (-1) \cdot 1/2 \cdot 1/2$
  $+ (-1) \cdot 1/2 \cdot 1/2 + 1 \cdot 1/2 \cdot 1/2 = 0$
- Similarly, $u_2(s) = 0$ by symmetry

		0.5 *Head*	0.5 *Tail*
0.5	*Head*	$1, -1$	$-1, 1$
0.5	*Tail*	$-1, 1$	$1, -1$

Mixed-Strategy Nash Definition

- Definition of best response & Nash equilibrium same as before
- Replace action profiles by strategy profiles & utility by expected utility
- **Best response** definition
  - $s_i^* \in \mathrm{BR}(s_{-i})$ iff $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \forall s_i \in S_i$
- **Nash equilibrium** definition
  - A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a mixed-strategy Nash equilibrium iff $\forall i \;\; s_i \in \mathrm{BR}(s_{-i})$

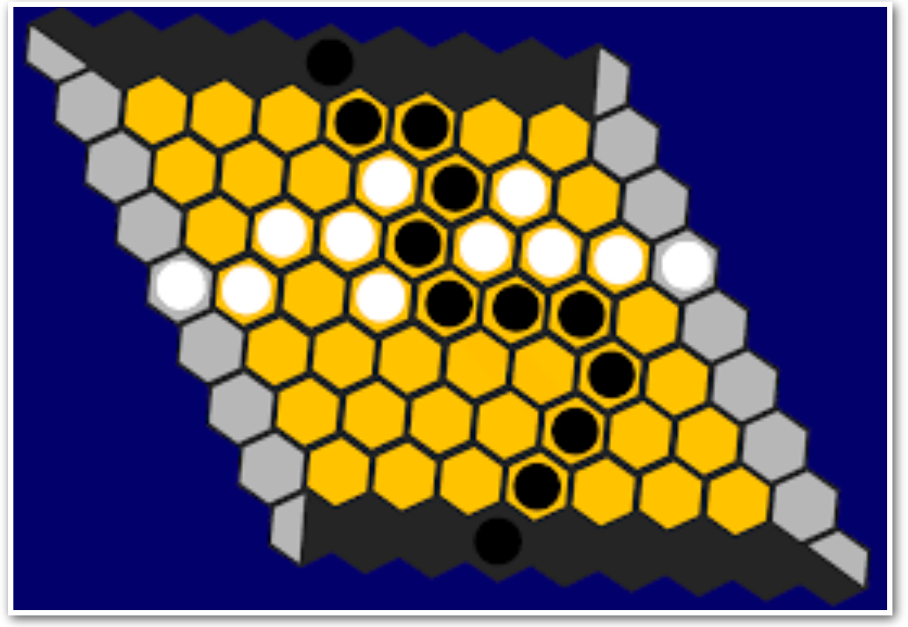

Theorem (Nash, 1950). Every finite game has a mixed-strategy Nash equilibrium.

Finding a Mixed Nash Equilibrium

- Nash's theorem tells us that one always **exists**
- Does not tell **how** we can find/compute it
- In general, it is **not obvious how to compute Nash equilibria**
- But if we can guess the support of the strategies, we can reason about the equilibrium
  - **Support** of a randomized/mixed strategy are the actions that are played with non-zero probability

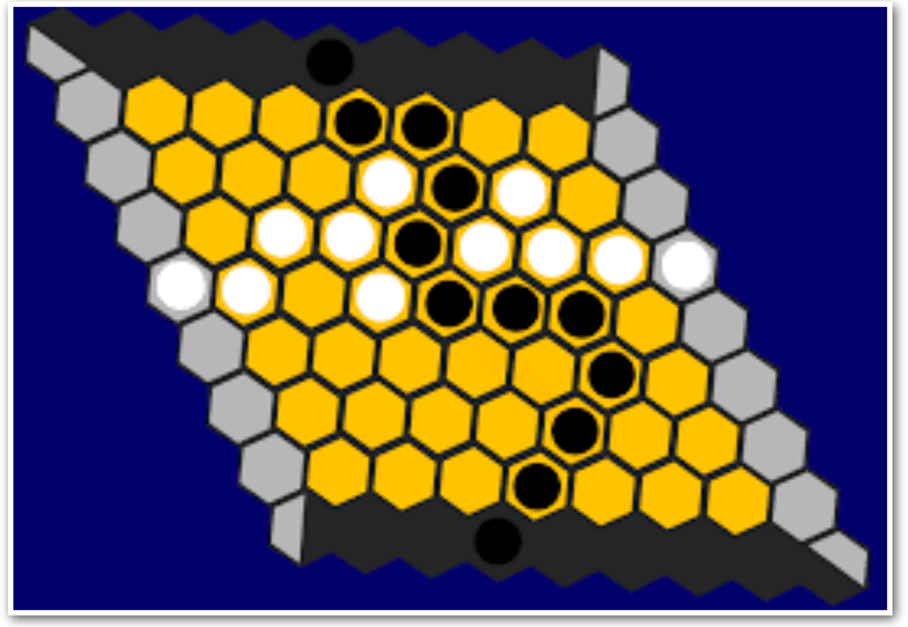

Theorem (Nash, 1950). Every finite game has a mixed-strategy Nash equilibrium.

Finding a Mixed Nash Equilibrium

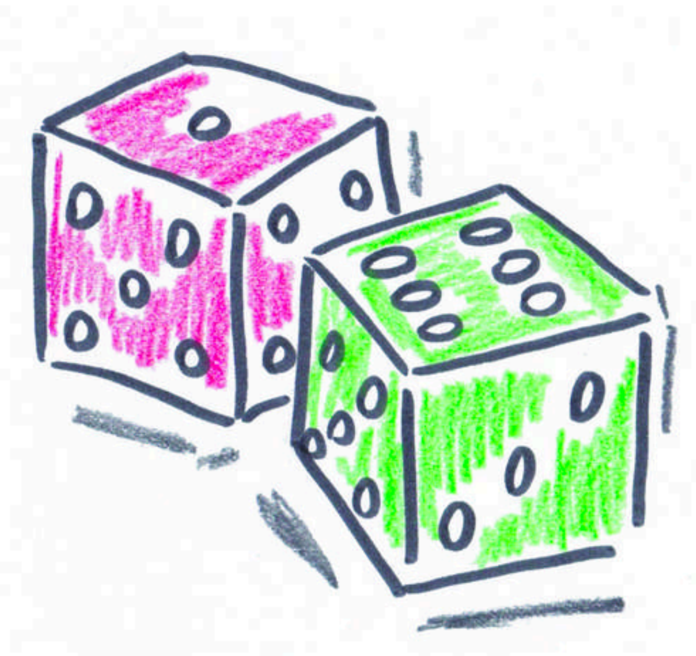

- Suppose P_2 plays B with prob q, and S with prob $1-q$
- Expected utility of player 1 if it only plays pure strategy B?
  - $2 \cdot q + 0 \cdot (1-q) = 2q$
- Expected utility of player 1 if it only plays pure strategy S?
  - $0 \cdot q + 1 \cdot (1-q) = 1-q$
- If $2q > 1-q$, then P_1 should always play B
- If $2q < 1-q$, then P_1 should always play S
- When $2q = 1-q$, P_1 is indifferent
- To be in an eqm, P_2 should choose $q = 1/3$

		q *Bach*	$1-q$ *Stravinsky*
p	*Bach*	$2, 1$	$0, 0$
$1-p$	*Stravinsky*	$0, 0$	$1, 2$

Finding a Mixed Nash Equilibrium

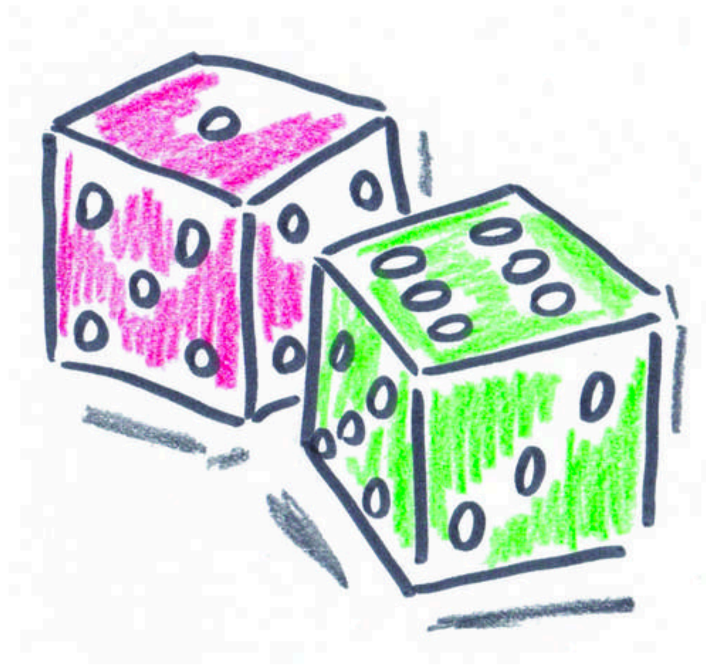

- Similarly, say P_1 plays B with prob p and S with prob $1-p$
- Expected utility of P_2 if it only plays pure strategy B?
  - $1 \cdot p + 0 \cdot (1-p) = p$
- Expected utility of P_2 if it only plays pure strategy S?
  - $0 \cdot p + 2 \cdot (1-p) \ = 2(1-p)$
- If $p > 2(1-p)$, then P_2 should always play B
- If $p < 2(1-p)$, then P_2 should always play S
- When $p = 2(1-p)$, P_2 is indifferent
- Thus, P_1 should choose $p = 2/3$

	q *Bach*	$1-q$ *Stravinsky*
p *Bach*	$2, 1$	$0, 0$
$1-p$ *Stravinsky*	$0, 0$	$1, 2$

Best Response Correspondence

- We can represent these choice of pure and mixed-strategies of the player in a best-response correspondence graph

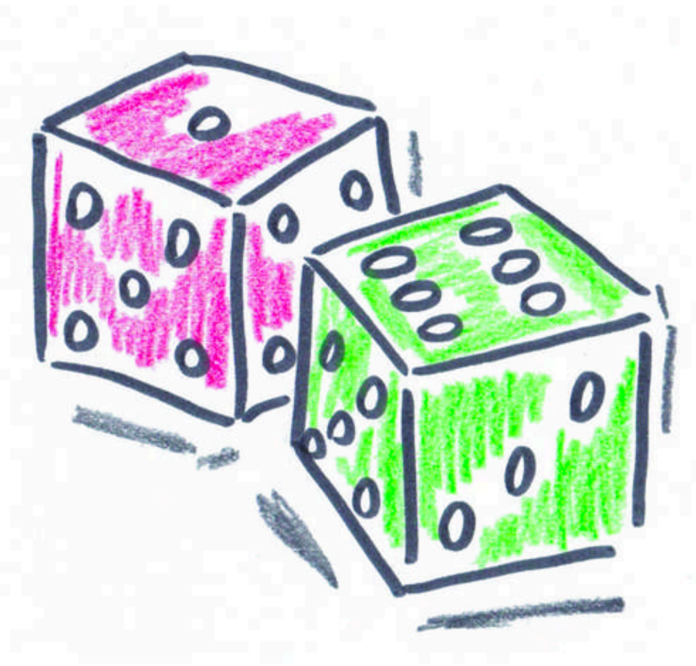

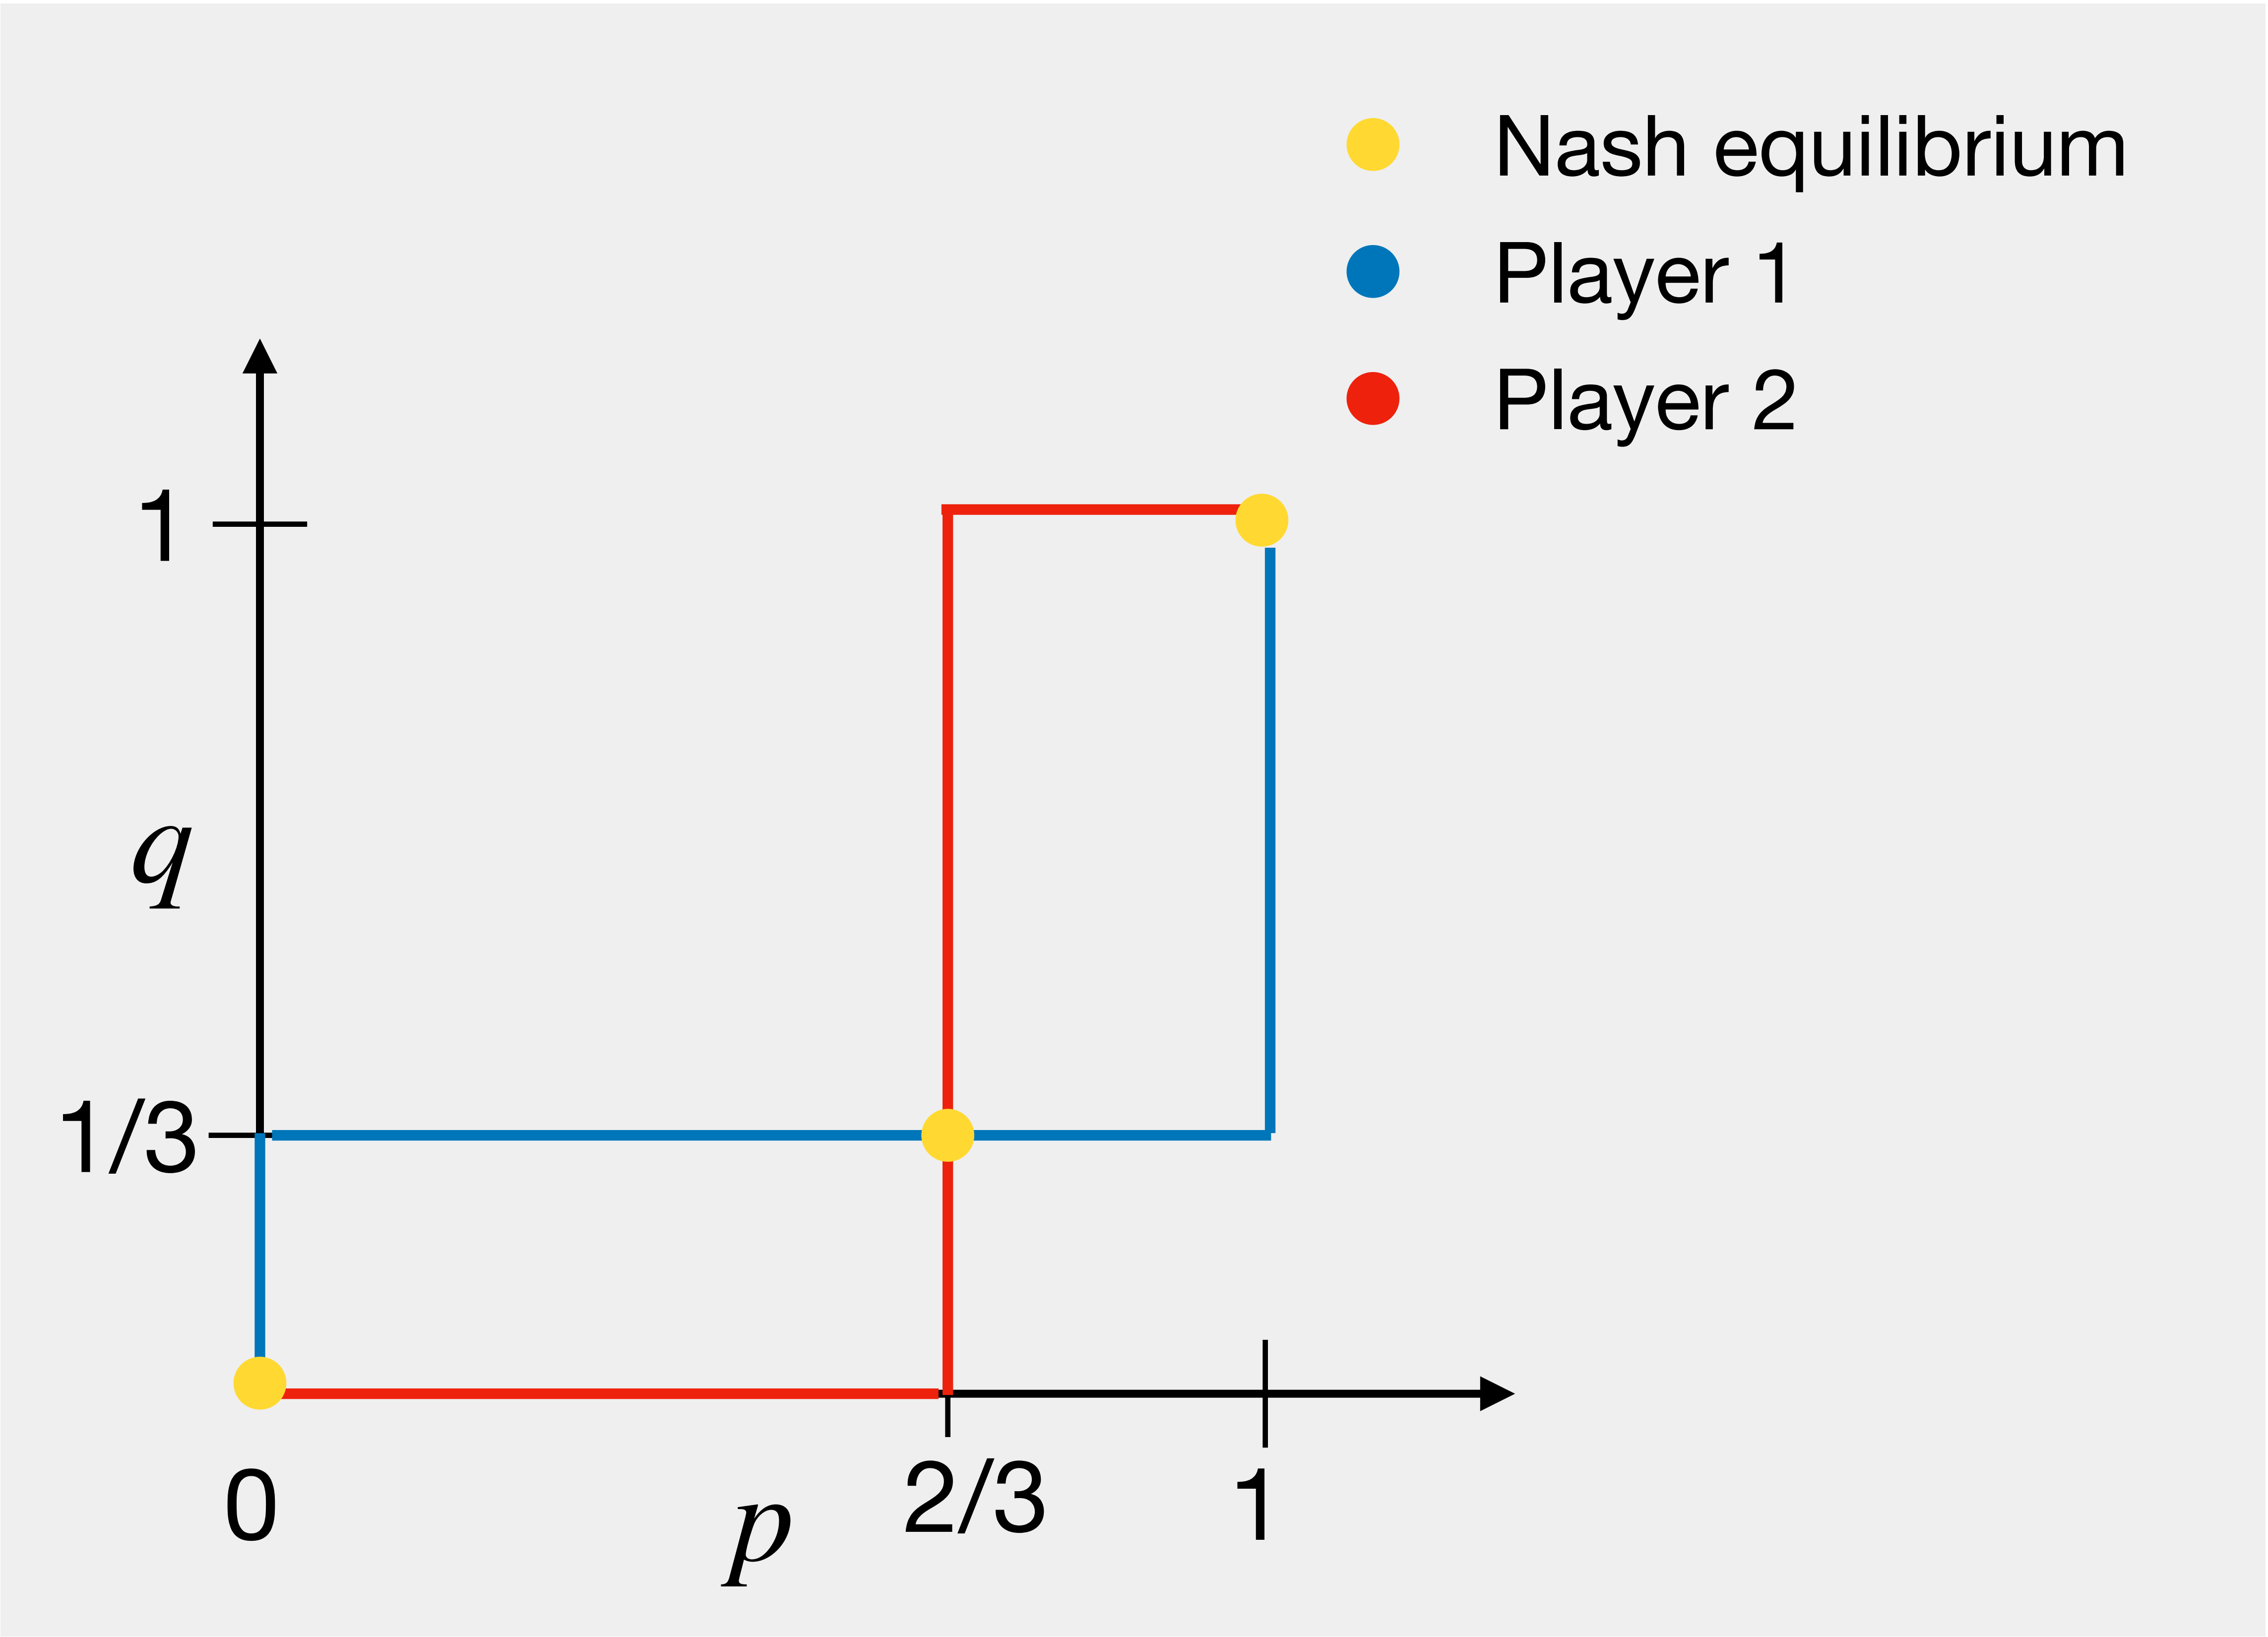

		1/3 *Bach*	2/3 *Stravinsky*
2/3	*Bach*	$2, 1$	$0, 0$
1/3	*Stravinsky*	$0, 0$	$1, 2$

Mixed-Nash Takeway

All actions in the support of a mixed strategy Nash equilibrium $s = (s_i, s_{-i})$ have the same expected utility

- Suppose one action $j \in A_i$ gives more expected utility to i than other actions $j' \in A_i$, then assigning more probability to j only improves $i's$ expected utility
- s_i is not a best response to s_{-i}
- Contradiction to s being a Nash equilibrium

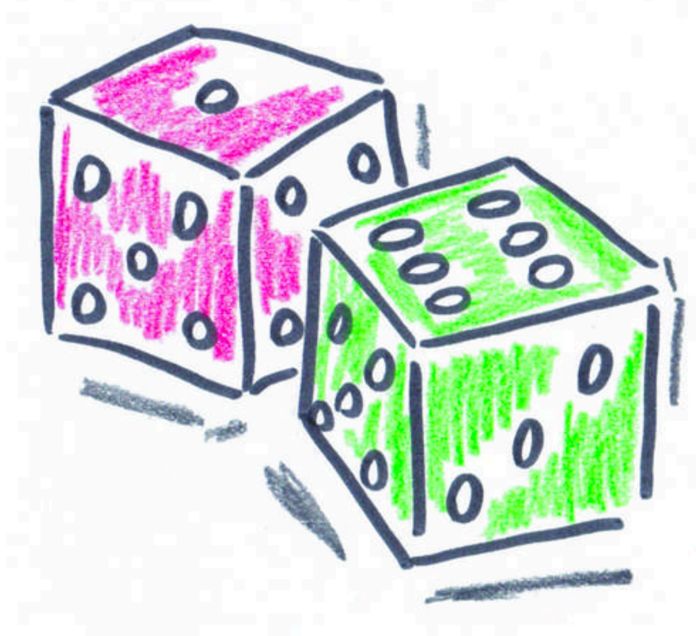

- In particular, in a mixed-strategy Nash, players are "indifferent" between the actions in their support
    - To reach a mixed-strategy Nash, players play in such a way as to make the opponent indifferent between their actions

Interpreting Mixed Strategies

- Why would players randomize over their actions?
- To create uncertainty/confuse the opponent
    - Imagine you are the government trying to secure n different checkpoints but have resources only for a few
    - If you deterministically choose which checkpoints to secure, attackers can pick the rest and always win
    - Thus, you need to randomize
    - Similarly tax audits are also randomized
- Randomize when uncertain about other's action
    - Games with multiple equilibria
        - E.g., BoS Game
- **Question.** Do players actually reach a mixed-Nash equilibrium in practice?

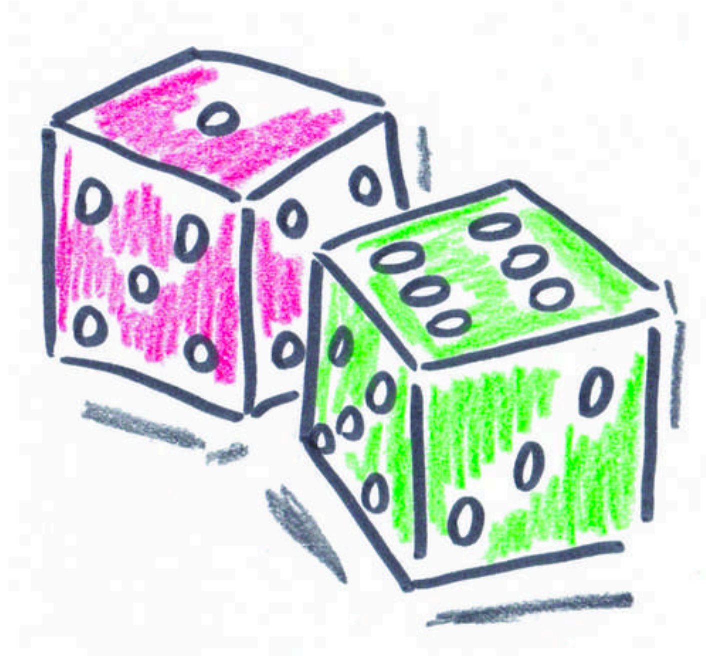

Soccer Penalty Kicks

- Mixed strategies in soccer penalty kicks
- Consider a goalie and kicker in soccer penalty kicks
  - Goal of kicker is to be unpredictable
- How do equilibrium strategies adjust to skills?
  - Should a kicker who kicks penalty kicks worse to the right than the left, kick more often to the left than right?

Soccer Penalty Kicks

- Start with the simple version where abilities are equal on left/right
- Randomizing 50-50 for both is the Nash

Goalie

Kicker		0.5 L	0.5 R
0.5	L	0, 1	1, 0
0.5	R	1, 0	0, 1

Soccer Penalty Kicks

- Now, we have a kicker who sometimes misses when they kick right
- **Question.** What is the mixed strategy Nash in this game?
    - Try this at home!

		Goalie: L	Goalie: R
Kicker	L	0, 1	1, 0
	R	$\frac{3}{4}$, $\frac{1}{4}$	0, 1

Soccer Penalty Kicks

- **Notice.** Even though the kicker is weak on the right side, in equilibrium, he prefers playing on the right compared to left

If goalie continues to play 50-50, the kicker could always go left and win 1/2 the time instead of 3/7.

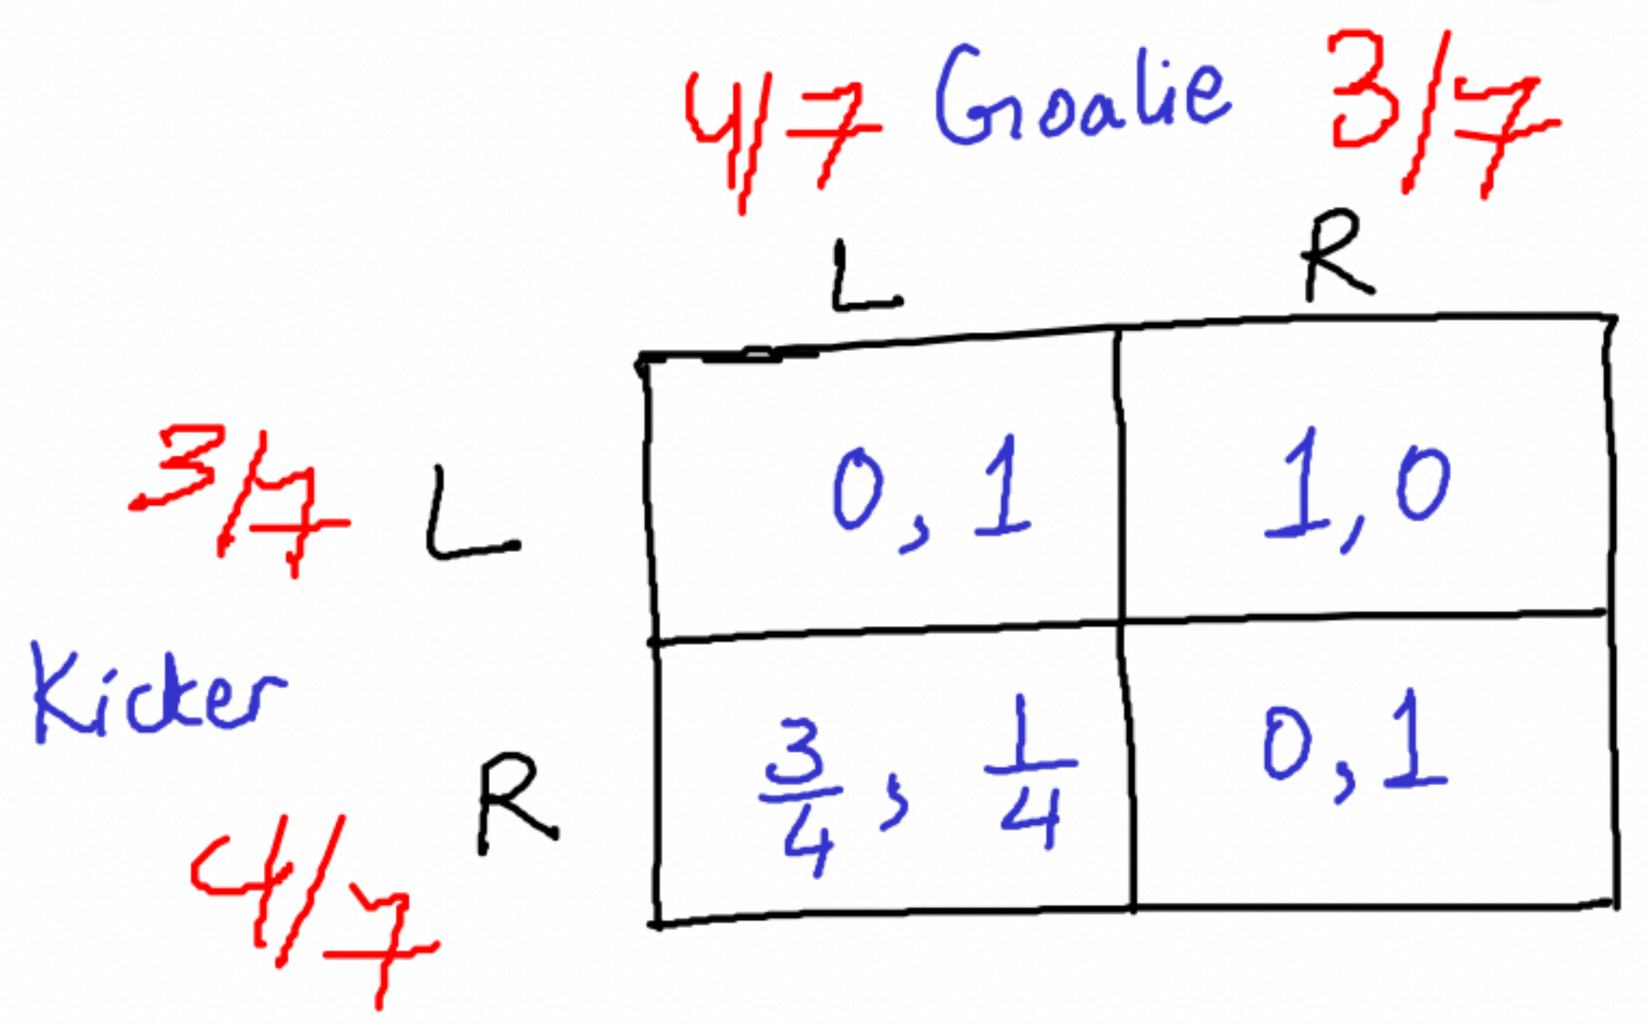

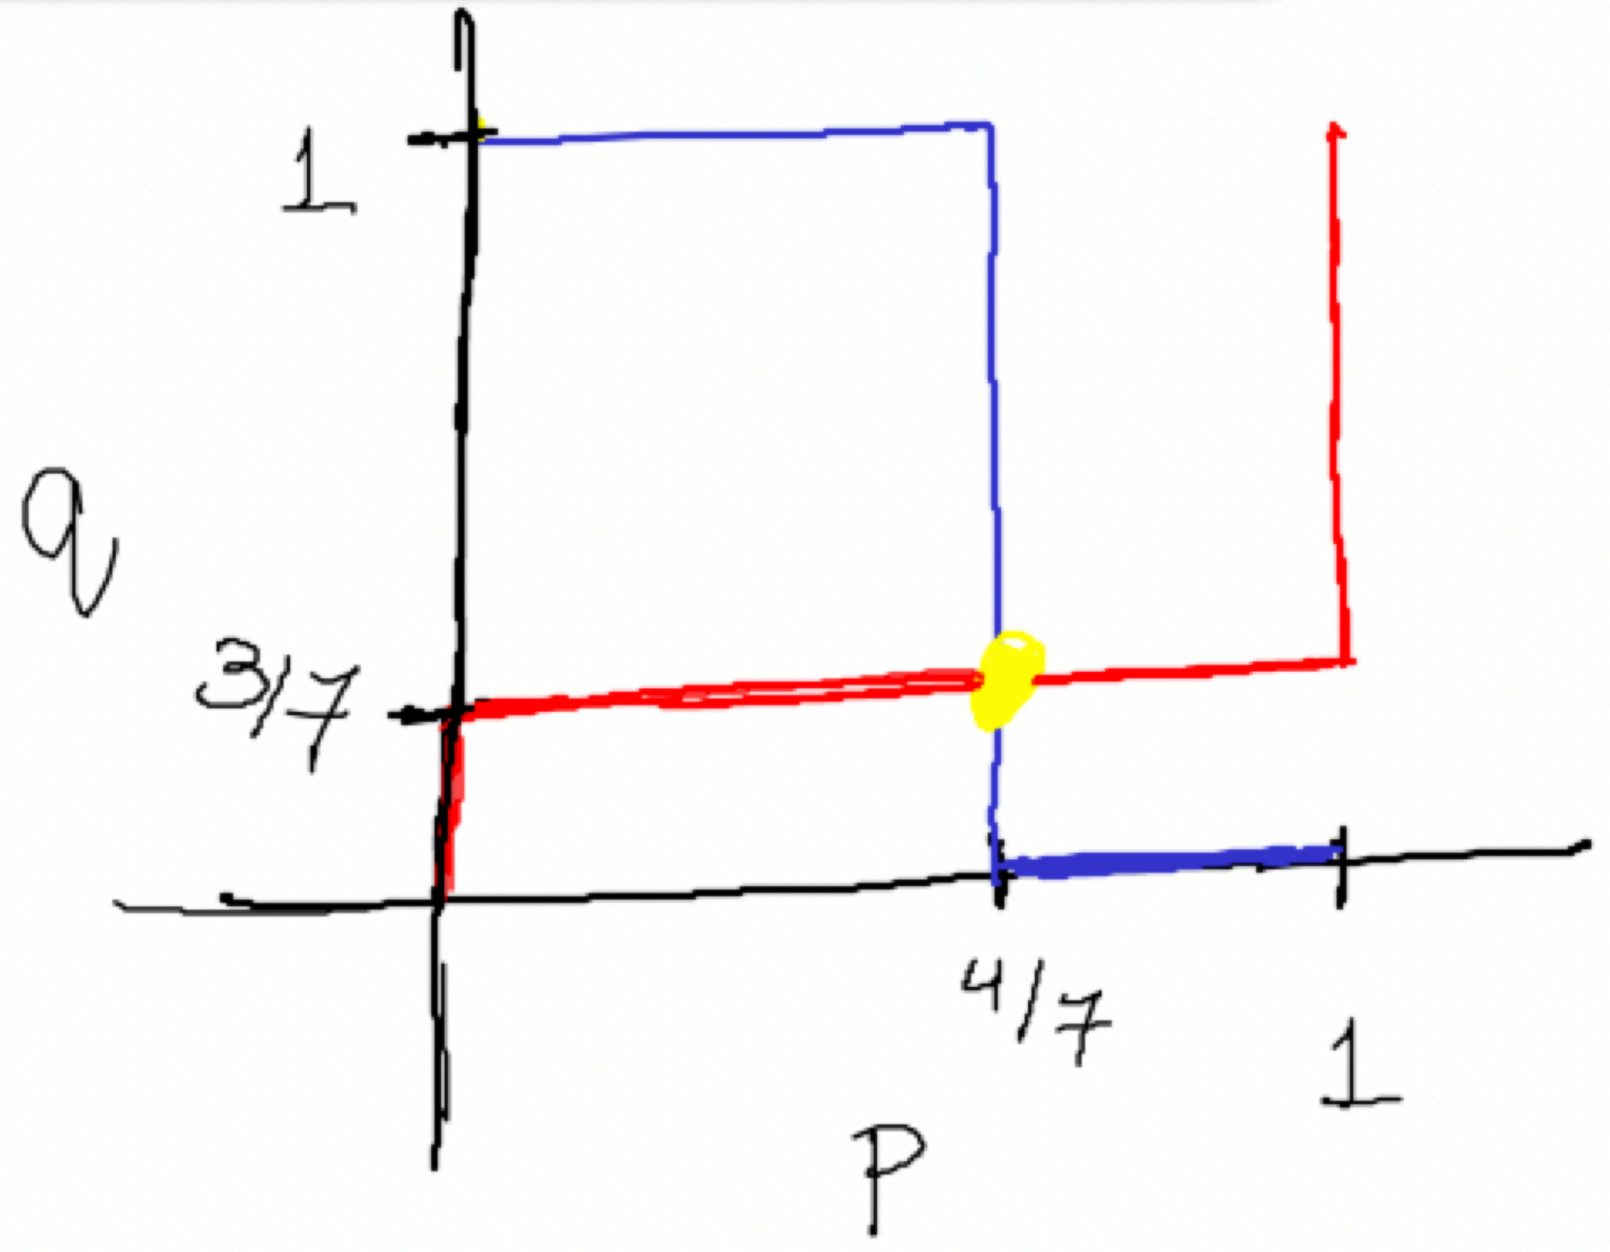

- **Question.** Do players really do this?

Data: Professional Soccer

- Ignacio Palacios-Heurta (2003) "Professionals Play Minimax", Review of Economic Studies
- Studied 1417 penalty kicks from FIFA games: Spain, England, Italy
- Ignacio considers left, center or right choices and which leg the kicker used
    - We will restrict to left and right choices
- Averaged win data (payoffs) from actual plays:

0.58, .42	0.95, .05
0.93, 0.07	0.70, 0.30

Biases exist!

Data: Professional Soccer

- What is the mixed Nash here?
- What was actually observed on average in 1417 games?

		0.42	0.58
		L	*R*
0.38	*L*	0.58, .42	0.95, .05
0.62	*R*	0.93, 0.07	0.70, 0.30

Data: Professional Soccer

- What is the mixed Nash here?
- What was actually observed on average in 1417 games?
- Players are playing the Nash equilibria in practice!

		0.42	0.58
		L	*R*
0.38	*L*	0.58, .42	0.95, .05
0.62	*R*	0.93, 0.07	0.70, 0.30

		0.42	0.58
		L	*R*
0.40	*L*	0.58, .42	0.95, .05
0.60	*R*	0.93, 0.07	0.70, 0.30

Mixed Nash in Nature

- Mixed-Nash has proved to be a good predictor of behavior in sports
- Other sports studied as well:
  - **Minimax Play and Wimbeldon**, M. Walker and J Wooders
- Randomized play seen in nature and evolutionary studies as well

Minimax Play at Wimbledon

By Mark Walker and John Wooders*

In many strategic situations it is important that one's actions not be predictable by one's opponent, or by one's opponents. Indeed, the origins of modern game theory lie in the attempt to understand such situations. The theory of mixed-strategy play, including von Neumann's Minimax Theorem and the more general notion of a Nash equilibrium in mixed strategies, remains the cornerstone of our theoretical understanding of strategic situations that require

are soundly rejected. Our results therefore provide some evidence that play by highly motivated and highly experienced players may conform more closely to the theory of mixed-strategy equilibrium than the play that has been observed in experiments.

We begin from the observation that games are not easy to play, or at least to play well. This is especially true of games requiring unpredictable play. Consider poker—say, five-card draw

Mixed Nash in Mechanism Design

- Mechanism design is the algorithmic design of games to achieve certain desirable properties: e.g. ensure players reach a "good" equilibrium
- Game theory vs mechanism design
  - Maskin's commentary to give you intuition about the differences
- Suppose you have designing a game
  - Making the rules and assigning utilities
  - Suppose your goal was to achieve some outcome
  - Would mixed Nash pose any challenges?

Complexity Of FindNash

Why Should We Care?

- If we believe in equilibrium theory, efficient algorithms can help make predictions
- If equilibria are supposed to model behavior, computational tractability is an important modeling prerequisite

"*If your laptop cannot find the equilibrium, then how can the market?*" - Kamal Jain, eBay

Pure Strategy Nash

- Suppose it exists, what is the complexity of finding it?
- Start with 2 player game, let each player have m actions
  - How many outcomes/action profiles?
    - m choices for each player, so m^2 combinations
- Verifying a Nash means check all possible deviations
  - How many such deviations to check per player?
  - $(m-1)$ per player
- Overall complexity: $O(m^3)$
- What is the input size? $O(m^2)$
- Running time of algorithm is $O(x^{3/2})$ in size of input x

Pure Strategy Nash

- What if we have n players with m actions per player?
- Number of outcomes/actions profiles?
  - m^n
- How many deviations we need to check to verify a Nash?
  - $n(m-1)$
- Overall complexity is now $O(nm^{n+1})$
- Size of input?
  - $O(nm^n)$
- So technically, we can find pure Nash using a brute force technique in polynomial time, but this is not very reassuring given the complexity

Succinct Games

- Not all games require enumerating all possible outcomes
- Example of succinct games:
  - Symmetric utility games: congestion/ network games
- Turns out pure Nash always exists in certain routing games
  - We will prove this in class!
  - And give an approximation result on how good these equilibria are
- Complexity of PNE in such games is topic of extensive research in theory CS
  - Under some conditions efficient iterative-best-response algorithms known for finding approximate solutions

Mixed-Strategy Nash

- We start with a special case

"Two-player zero-sum games are one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made."- Robert Aumann, 1987

Two-Player Zero-Sum Games

- A zero sum game is a strictly competitive game where the sum of the utilities of players in each outcome is zero
- In a two-player zero-sum game, player 1 chooses an action to maximize their payoff, under the assumption that player 2 will do their best to hurt them as much as possible (one person's gain is another's loss)
    - **Maximin strategy**: P_1 maximizes their payoff under the assumption that P_2 will try to minimize it
    - **Minimax strategy**: P_2 (receives $-u_1$) and tries to maximize it: that is, plays to minimize the maximum payment to P_1
- **Surprising fact.** Pessimistic optimization by both players lead them to a Nash equilibrium (uses LP Duality theory)

Example: Odd/Even Game

- Rick and Morty are playing the following game
  - Each of them put a $1\$$ or $2\$$ bill on the table
  - If the sum is odd, Rick wins, otherwise Morty wins
  - Winning players gets to keep the total sum
- Suppose Rick commits to strategy $(x,\ 1-x)$
- Expected utilities for Morty:
  - $u_2(1\$) = 2x - 3(1-x)$
  - $u_2(2\$) = -3x + 4(1-x)$

Morty

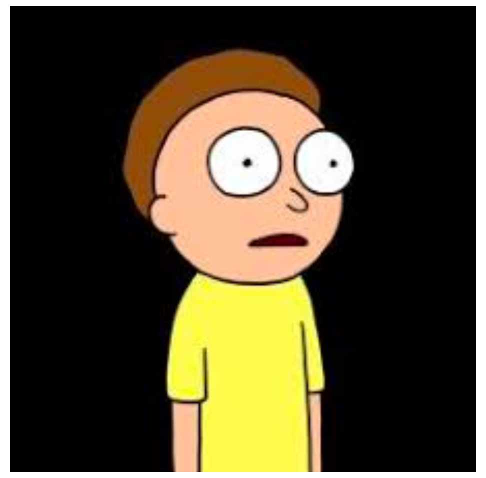

Rick

Rick \ Morty		$1\$$	$2\$$
x	$1\$$	$-2,\ 2$	$3,\ -3$
$1-x$	$2\$$	$3,\ -3$	$-4,\ 4$

Two-Player Zero-Sum Games

- Suppose Rick commits to strategy $(x,\ 1 - x)$
- Expected utilities for Morty:
  - $u_2(1\$) = 2x - 3(1 - x)$
  - $u_2(2\$) = -3x + 4(1 - x)$
- Morty's best response: $\max(2x - 3(1 - x),\ -3 + 4(1 - x))$
- Rick receives: $-\max(2x - 3(1 - x),\ -3 + 4(1 - x))$
  $= \min(-2x + 3(1 - x),\ 3x - 4(1 - x))$
- So Rick's best strategy is to pick the max x that minimizes the following:
  $(-2x + 3(1 - x),\ 3x - 4(1 - x))$

Morty

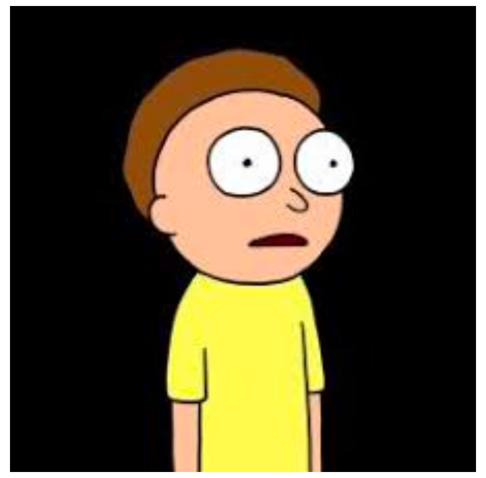

Rick

		$1\$$	$2\$$
x	$1\$$	$-2, 2$	$3, -3$
$1 - x$	$2\$$	$3, -3$	$-4, 4$

Two-Player Zero-Sum Games

- So Rick's best strategy is to pick the max x that minimizes the following: $(-2x + 3(1-x),\ 3x - 4(1-x))$
- If we plot these lines we get something like this

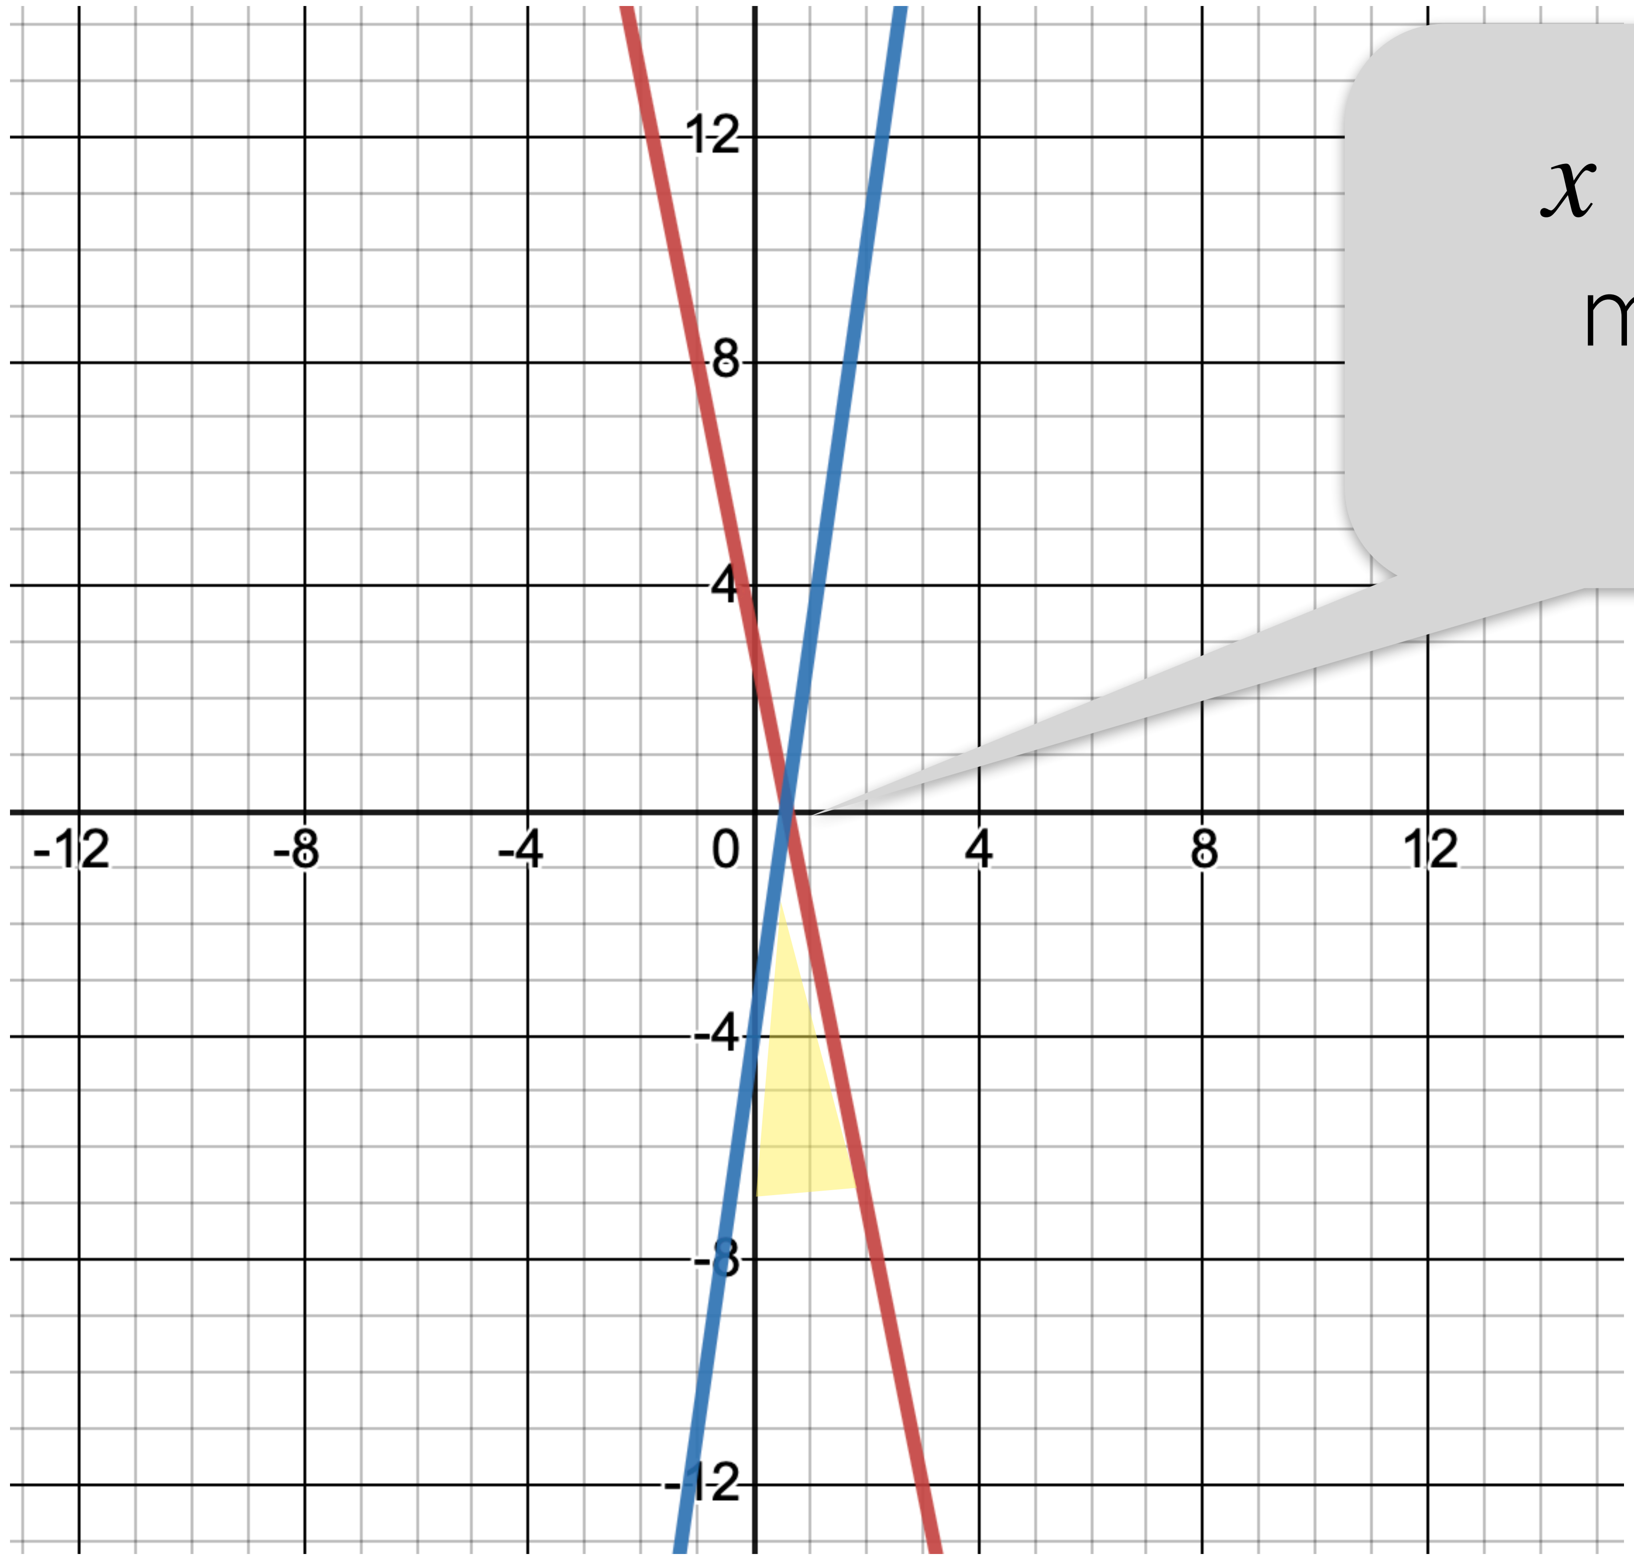

$x = 7/12$ is the maximum x that minimizes the two expressions; Called **maximin** strategy

Morty

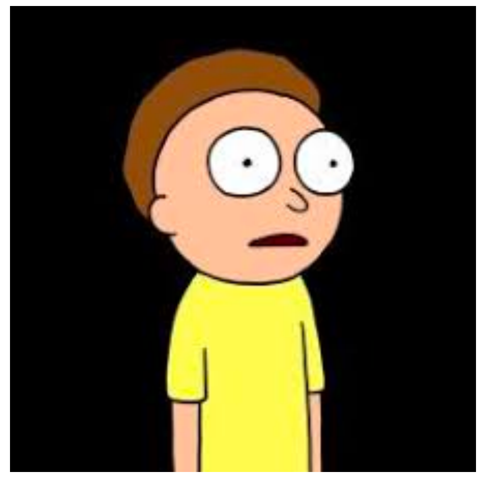

Rick

		1\$	2\$
x	1\$	$-2, 2$	$3, -3$
$1-x$	2\$	$3, -3$	$-4, 4$

Two-Player Zero-Sum Games

- Similarly if Morty commits to $(y, 1 - y)$, we can do the same reasoning, and should choose max y that minimize $(2y - 3y(1 - y), -3y + 4(1 - y))$
- Which gives us $y = 7/12$
- When there are m actions we can write a linear program to solve this optimization problem
- LPs are out of scope for this class but they can be solved in polynomial time!

Morty

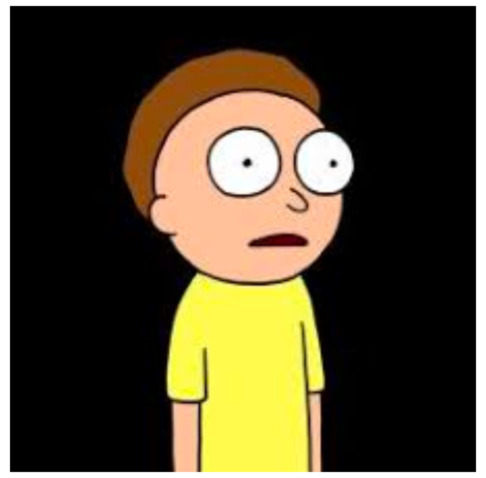

Rick

	y: 1\$	$1 - y$: 2\$
1\$	$-2, 2$	$3, -3$
2\$	$3, -3$	$-4, 4$

Two-Player Zero-Sum Games

- Turns out: $((7,12),(5,12),(7,12),(5,12))$ is also the Nash equilibrium of this game (you can verify that this makes the two players indifferent
- Essentially, both players can solve their own optimization problem pessimistically and reach a Nash!
    - Maxmin strategy same as minimax strategy
    - Why this works follows from **LP duality theory**
- **(Minimax Theorem)**
  Any maxmin (or minmax) strategy for player 1 and minmax (or maxmin) strategy for player 2 form a Nash
- Such a Nash can be computed efficiently using LPs

Morty

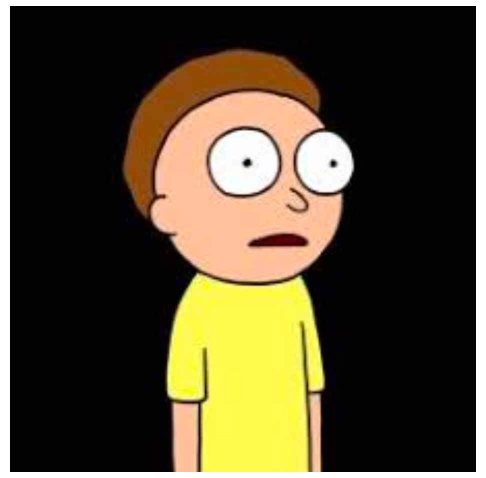

Rick

	1\$	2\$
1\$	$-2, 2$	$3, -3$
2\$	$3, -3$	$-4, 4$

Takeaways

- Two-player, zero-sum games are special
- Existence of a mixed-Nash equilibrium in such games follows from strong LP duality theory
- These techniques generalize to any two-player zero-sum game
- Since LPs are solvable in polynomial time, this allows us to efficiently compute the mixed-Nash equilibria of such games
- **Remember.** If you hear "maxmin" and "minmax" mentioned to you, know that these strategies characterize Nash equilibria in two-player zero-sum games

General Games

- What about finding mixed-strategy Nash in general games?
- Many heuristics have been studied that do well in practice
  - Lemke & Howson (1964)
  - Porter, Nudelman & Shoham (2004)
- Worst-case time is still exponential in the size of the payoff matrix
- Attempts to find polynomial time algorithms have failed

The Pavlovian Reaction

- "Is it NP-complete to find a Nash equilibrium?"
  - Probably not, because of unique property: **guaranteed to exist**
  - (It is NP-complete to find a slightly more general problem than just finding Nash, e.g., finding two Nash equilibria if they exist)

Complexity of Nash: PPAD

- Computer science's biggest contribution to the filed
- Computing a Nash equilibrium in general non-zero-sum games is **computationally intractable**, exactly as intractable as the class **PPAD**
- **PPAD**: "Polynomial Parity Arguments on Directed graphs"
- At a high level:
  - **F**NP problems are constructive versions of NP problems (**F**: "Functional")
  - **T**FNP is a subclass of FNP for which a solution is guaranteed (**T**: "Total")
  - **PPAD** is a subclass of **TFNP** where the proofs are based on parity arguments in directed graphs

Theorem [Daskalakis, Goldberg, Papadimitriou ’06]. FindNash problem is PPAD-complete.

Complexity of Nash: PPAD

- Where is **PPAD**?
  - Intermediate to **P** and **NP** complete!
- A natural and fundamental problem of "intermediate" difficulty, unlikely to be either in P or NP complete
  - Most problems we studied in algorithms were either in **P** or were **NP** complete
  - Very few other such computational problems (graph isomorphism)

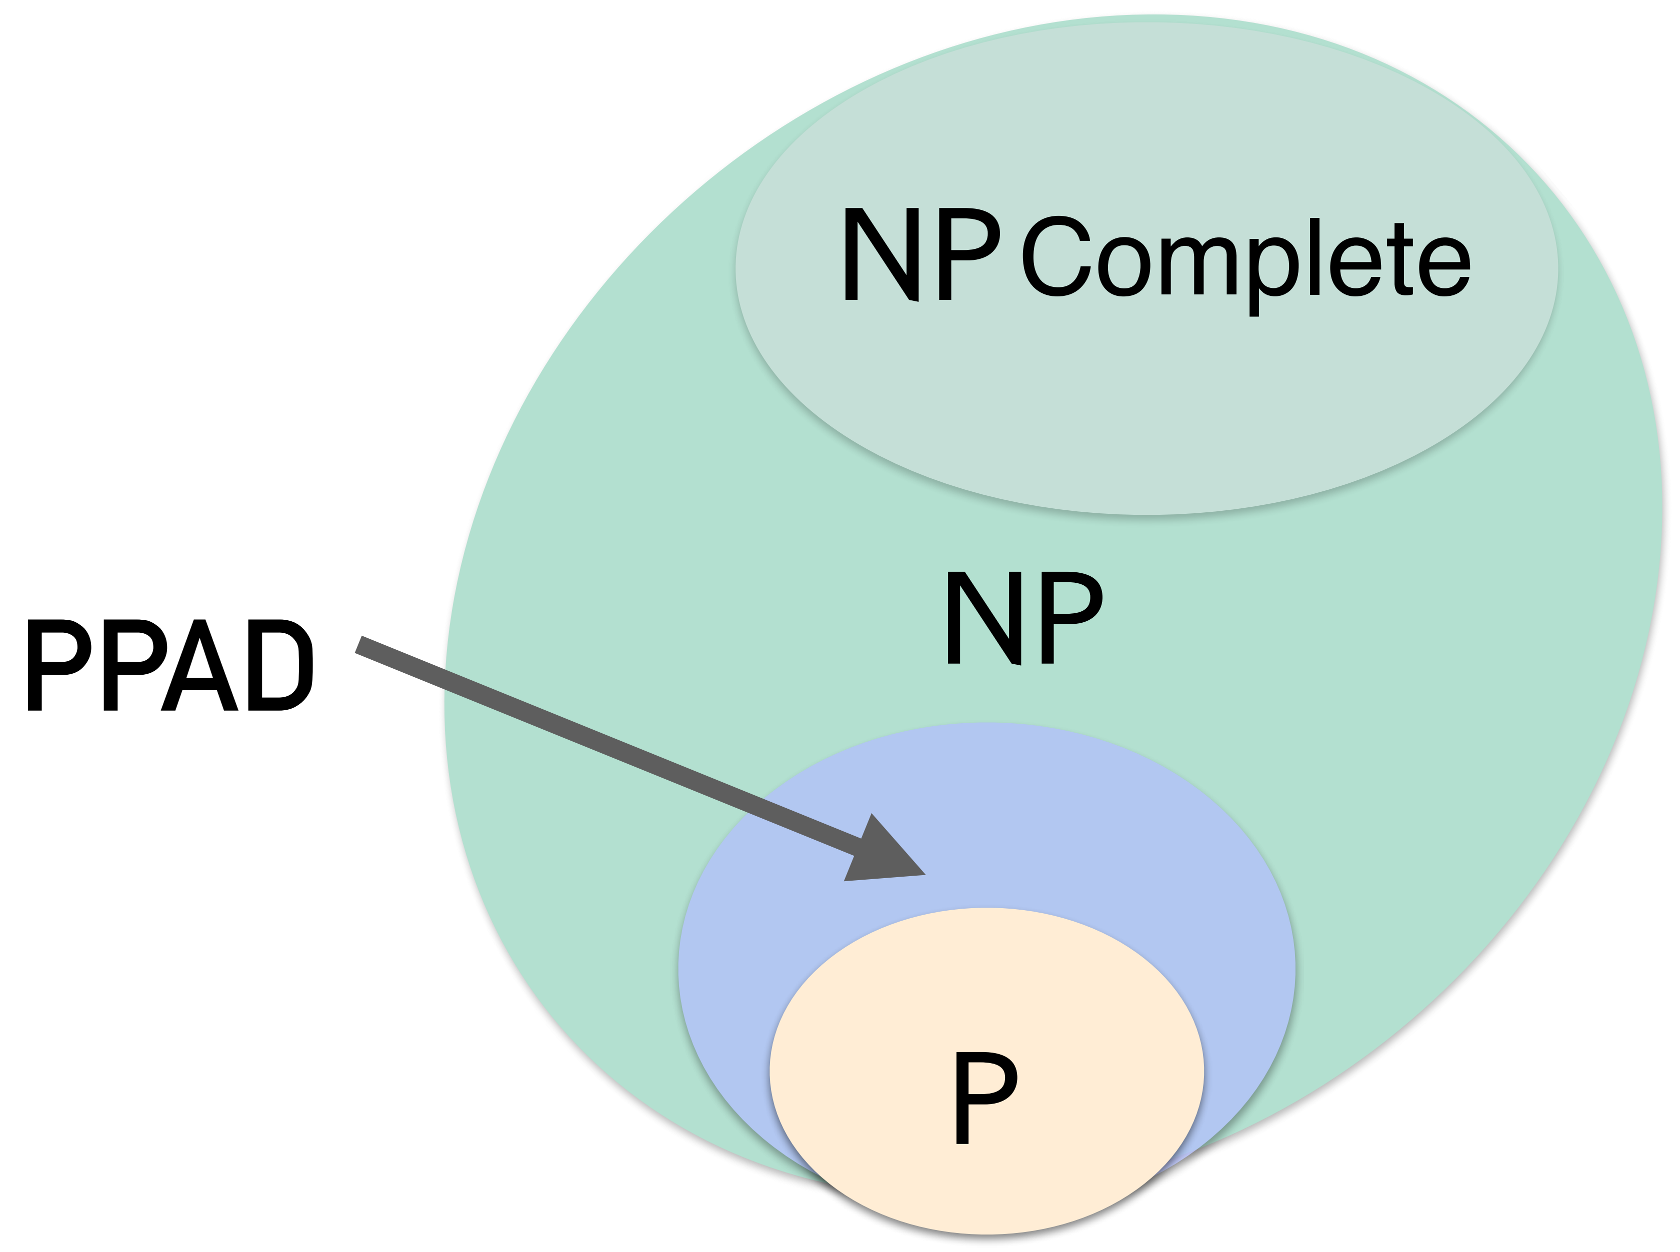

Theorem [Daskalakis, Goldberg, Papadimitriou '06].
FindNash problem is PPAD-complete

Other Problems in PPAD

- SPERNER: Takes an input a triangulation of a square, where the bottom edge does not have any vertices of color red, left edge does not have any vertices of color blue and the rest of the edges do not have vertices of color yellow.
- Problem is to find a trichromatic triangle, that is, a triangle whose vertices have all three colors
    - No matter how the internal nodes are colored, there exists a tri-chromatic triangle

SPERNER ∈ PPAD

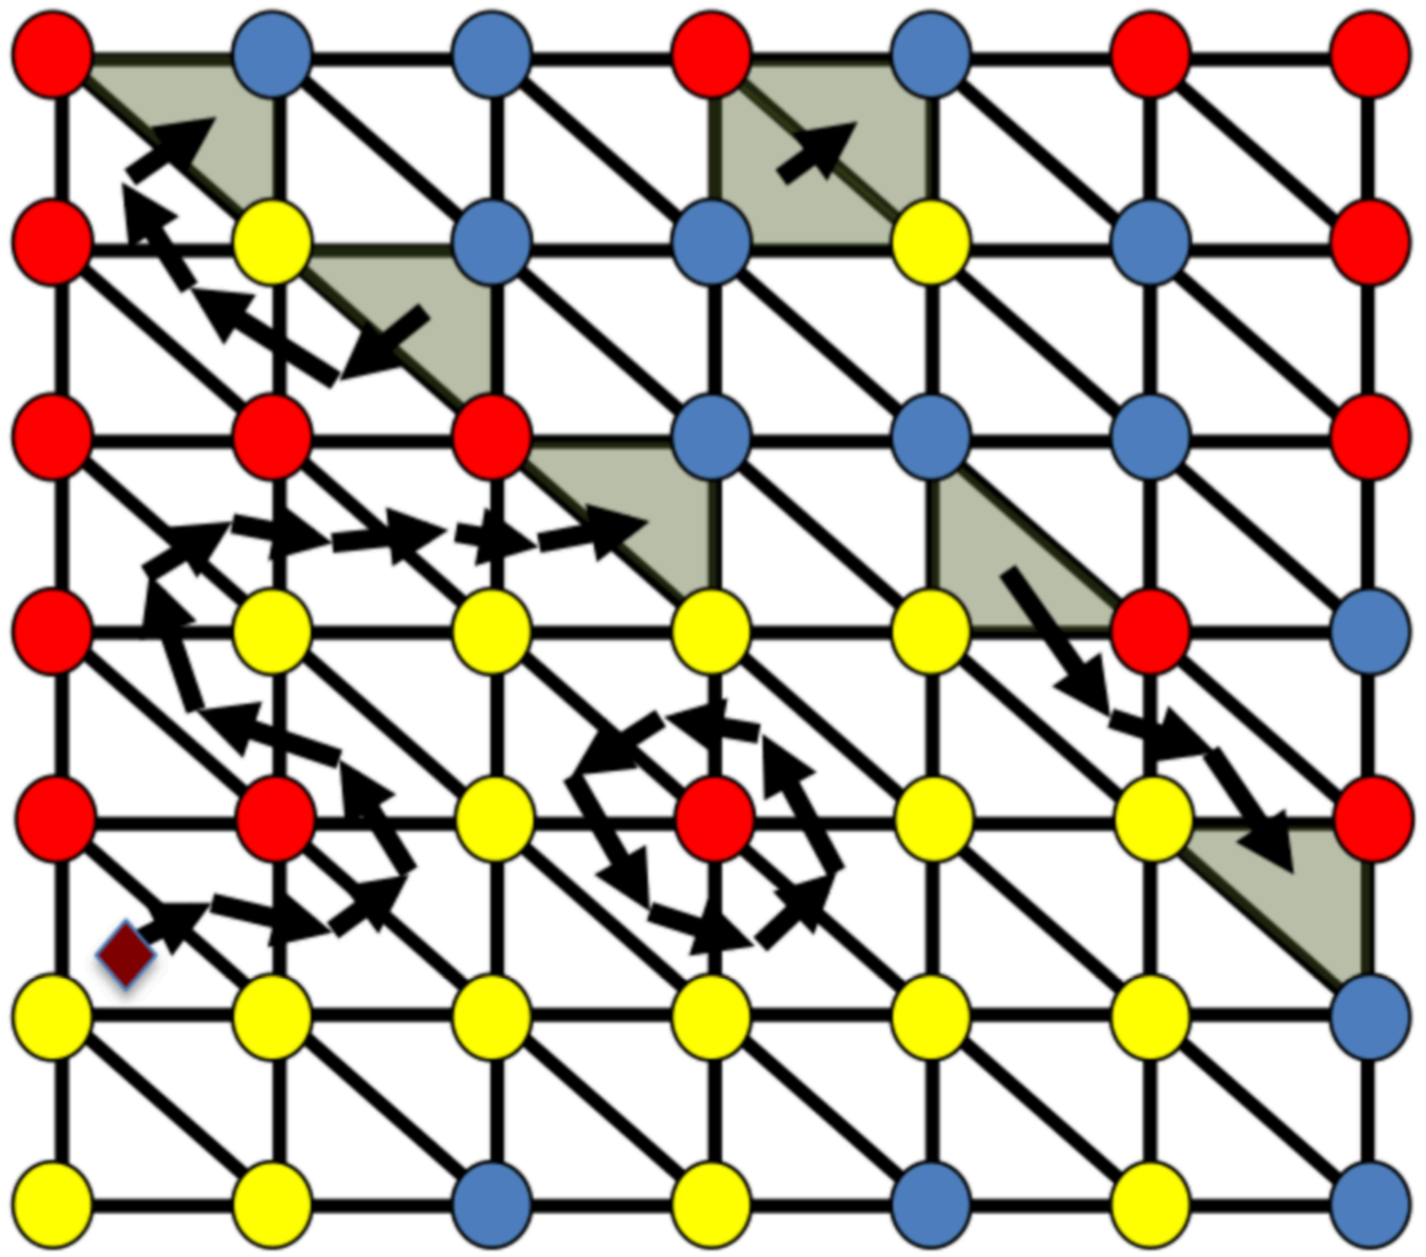

Other Problems in PPAD

- BROUWER:
  Let $f : D \rightarrow D$ be a continuous function from a convex and compact subset D of the Euclidean space to itself
- Then, there exists a point $x \in D$ such that $x = f(x)$
- Problem is to find such a point

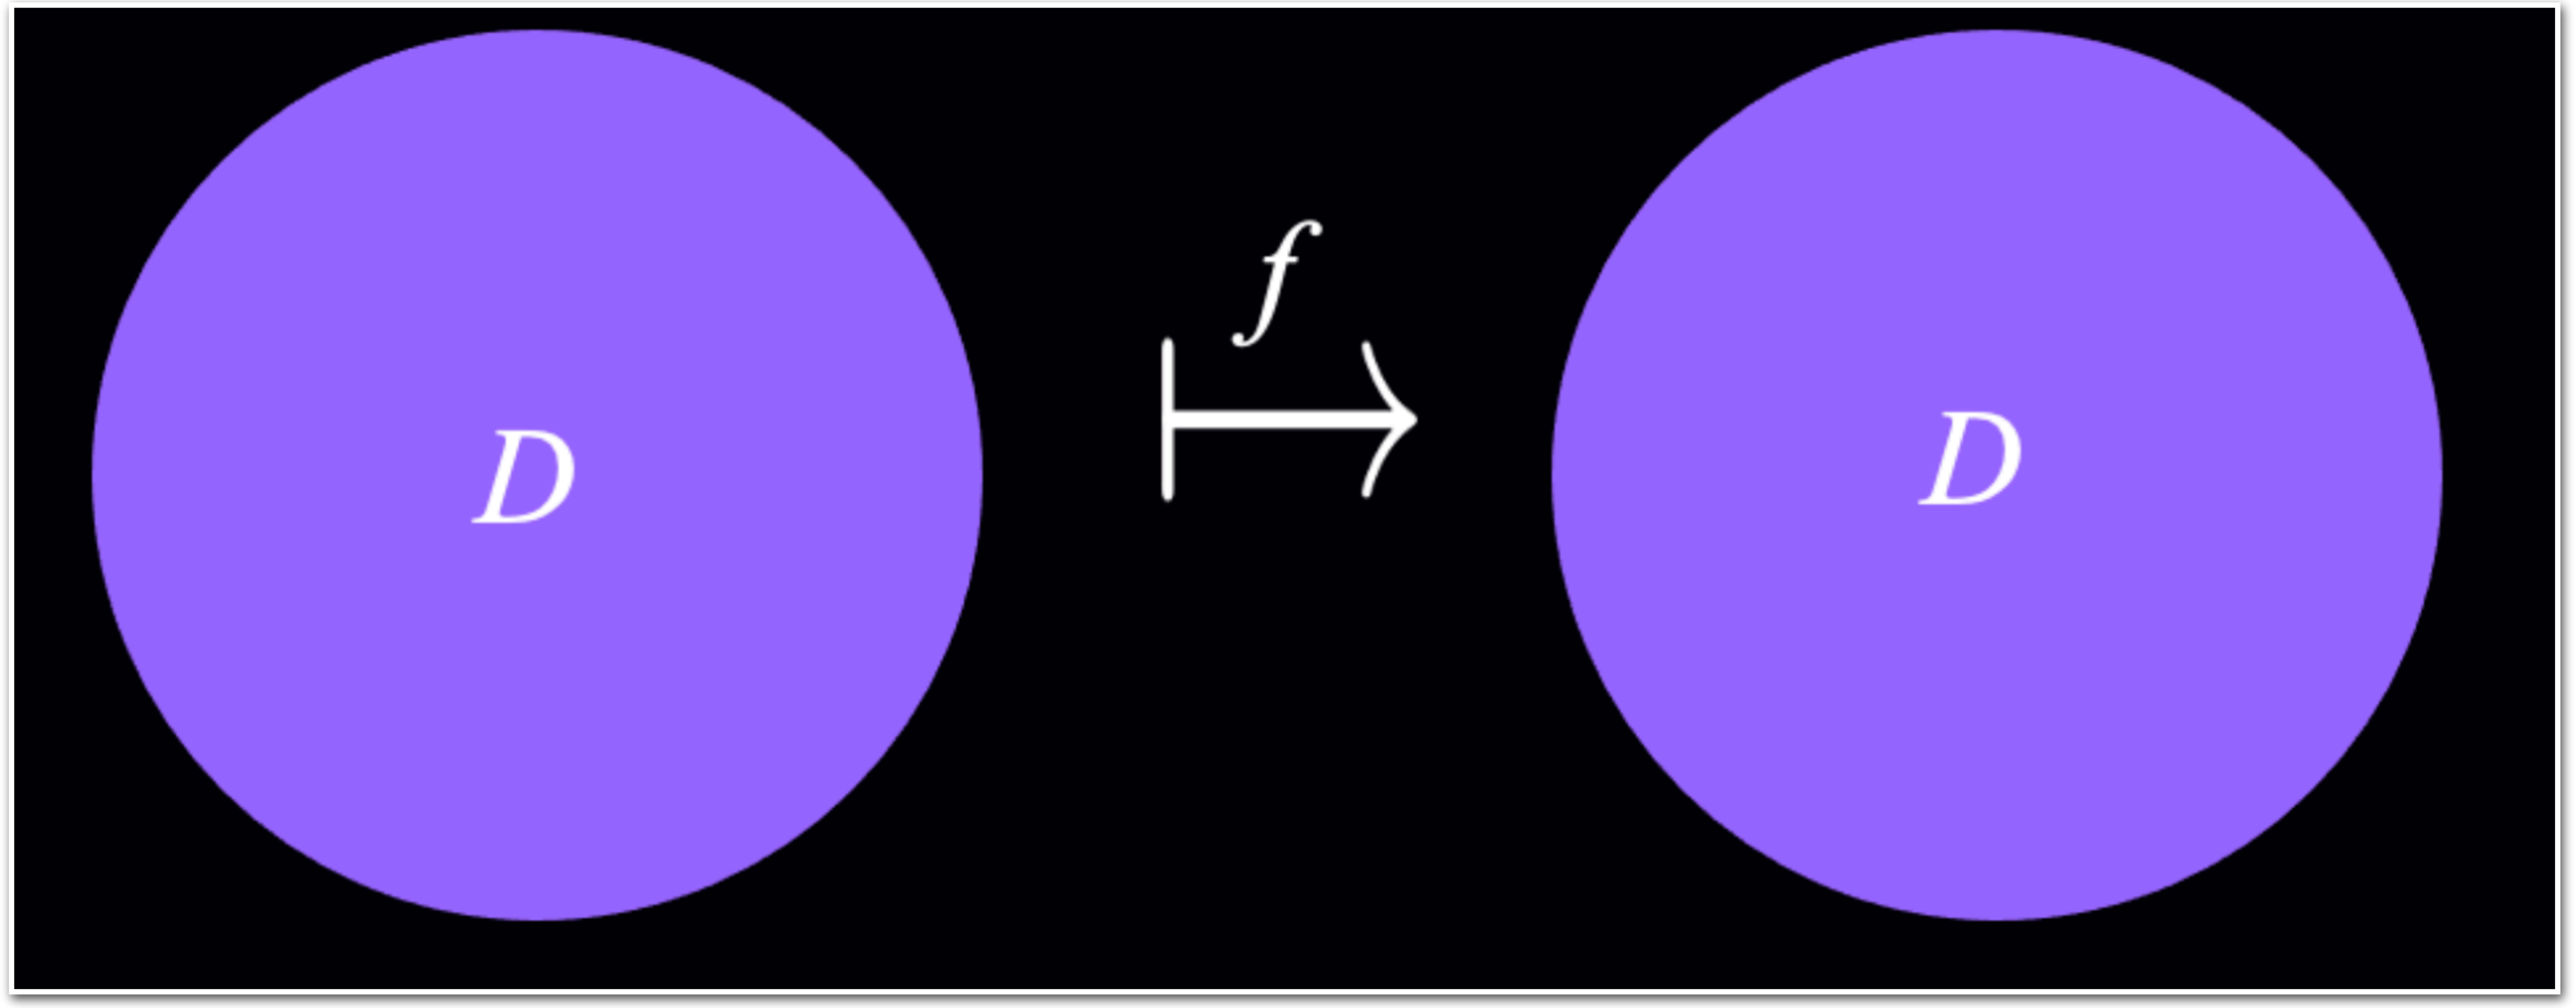

BROUWER $\in$ PPAD

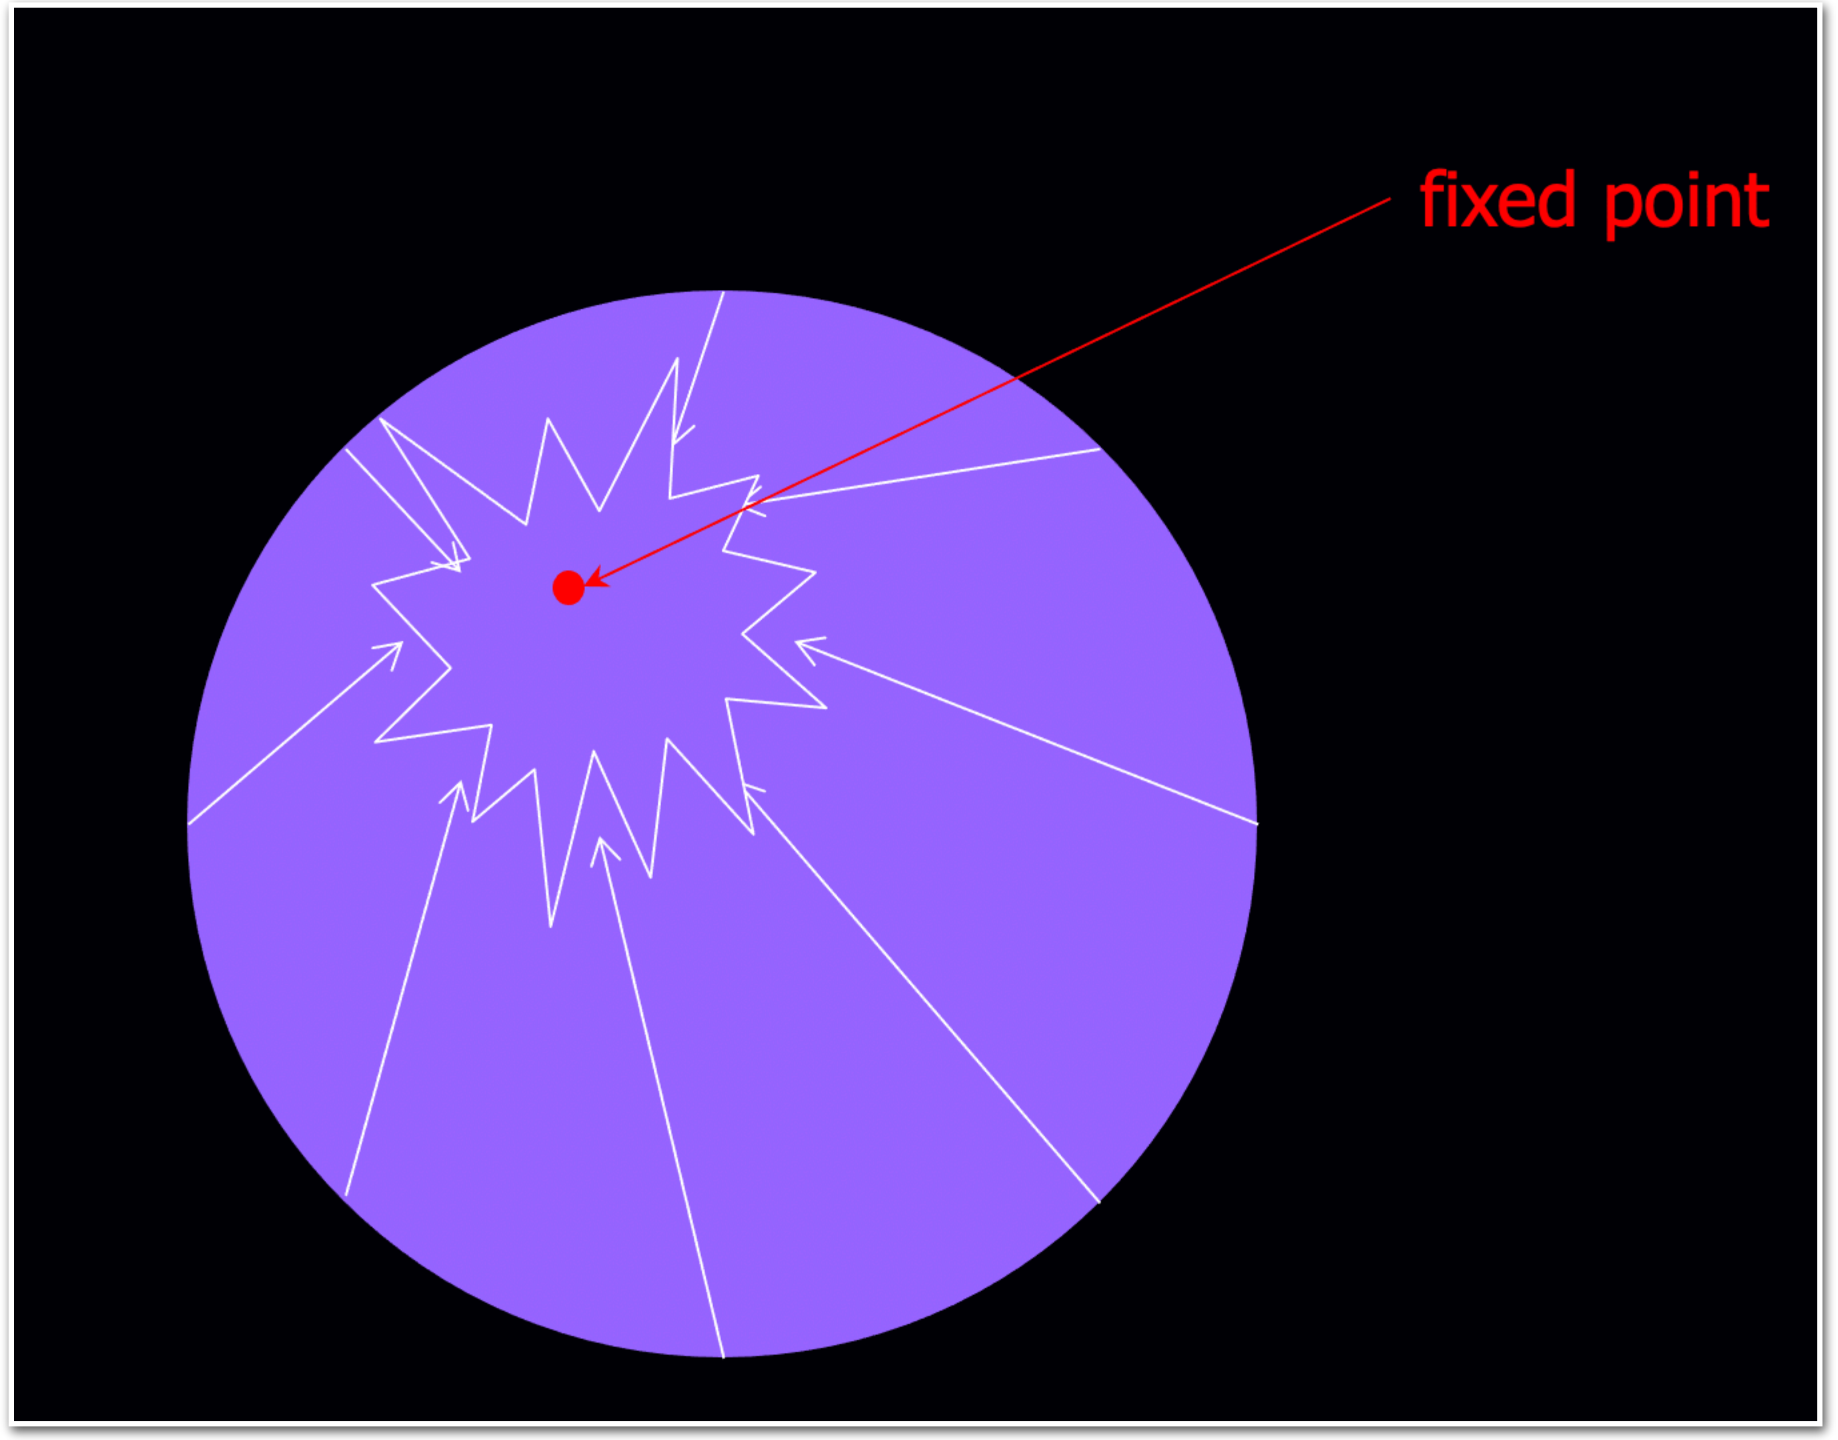

Game Over?

- If computing equilibrium is hard, can we expect agents to actually find and play it?
- Worst-case result: does not preclude fast solutions in practical instances
- Motivates other computationally tractable solution concepts such as correlated and coarse correlated equilibria
- Many iterative iterative learning algorithm are known that converge quickly to approximate equilibria as well
    - Shows how agents reach an equilibria
    - Concepts from online learning theory
- Motivates the field of **mechanism design**
    - Design mechanisms that admit tractable equilibrium concepts!

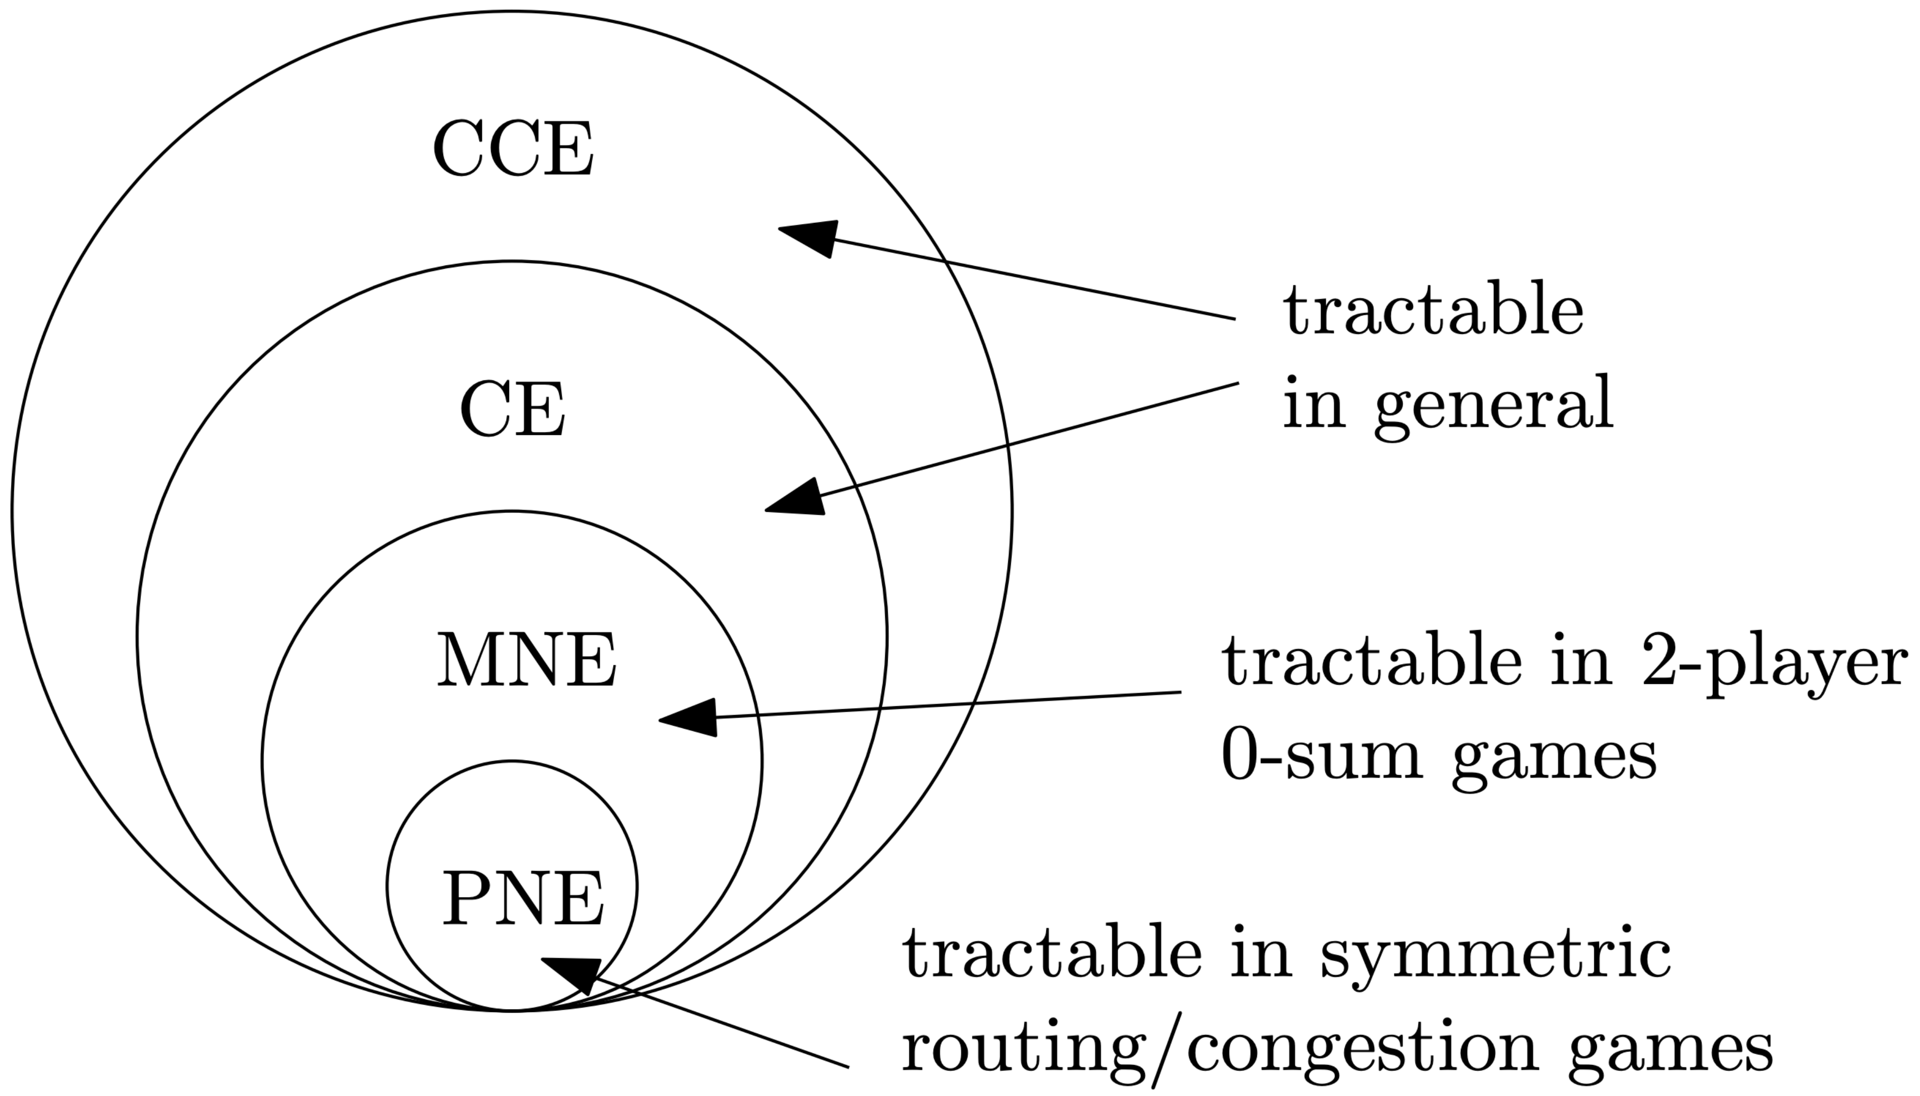

Mechanism Design

Markets with Money

Markets without Money

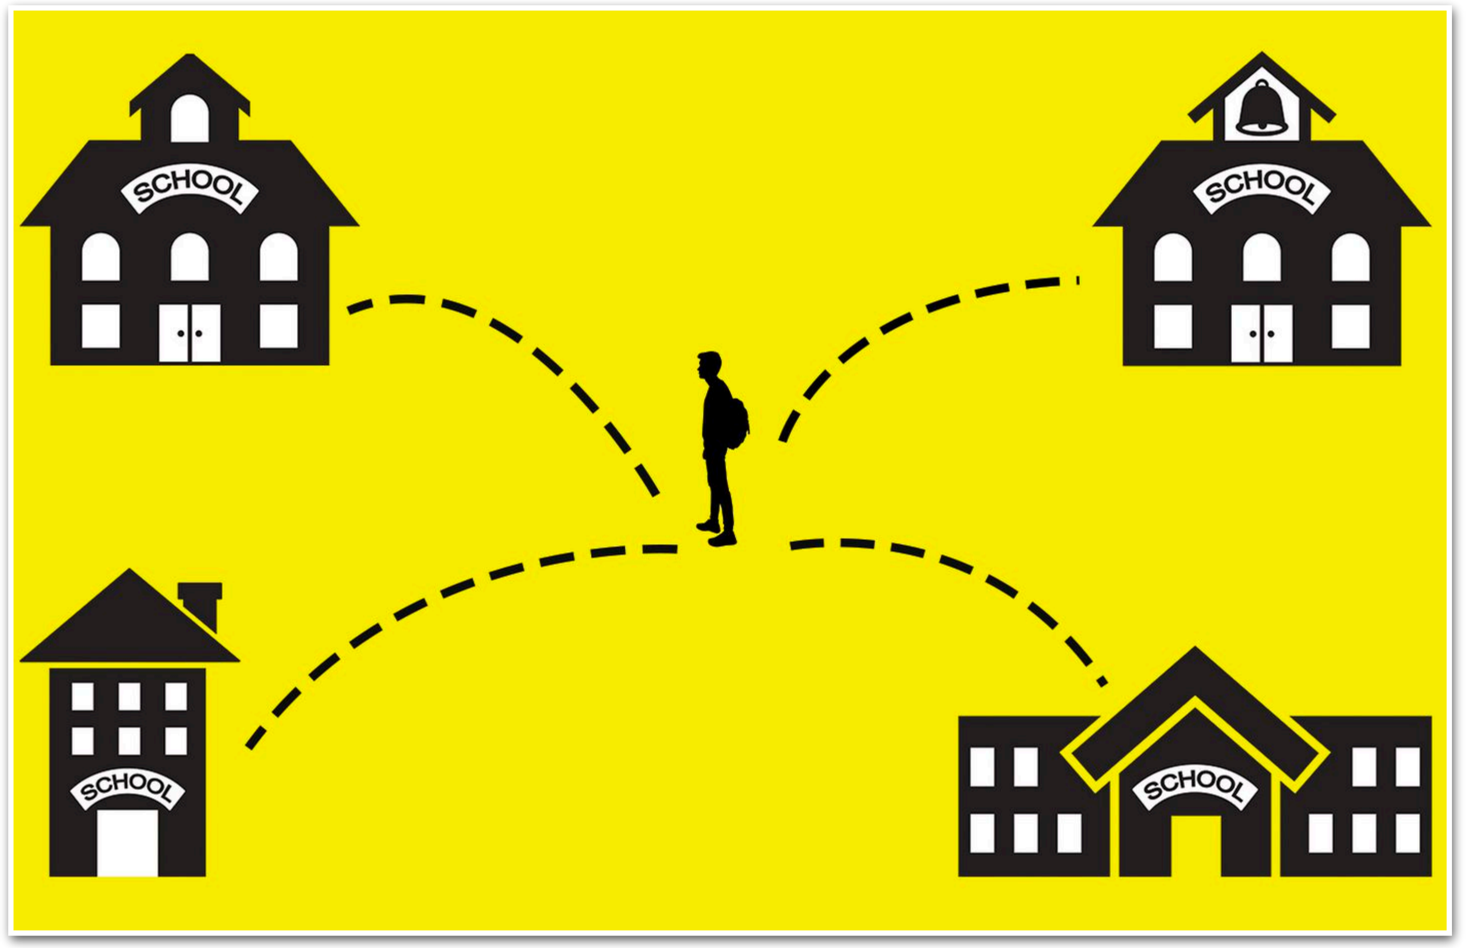

Utility: ordinal/ monetary equivalent

Auctions

Utility: Cardinal, based on preferences

Stable matchings/ Voting

Next Time: Auctions