TTIC 31260 Algorithmic Game Theory

Spring 2024

Avrim Blum

TTIC

Overall

Theory and algorithms for systems of interacting agents, each with their own interests in mind.

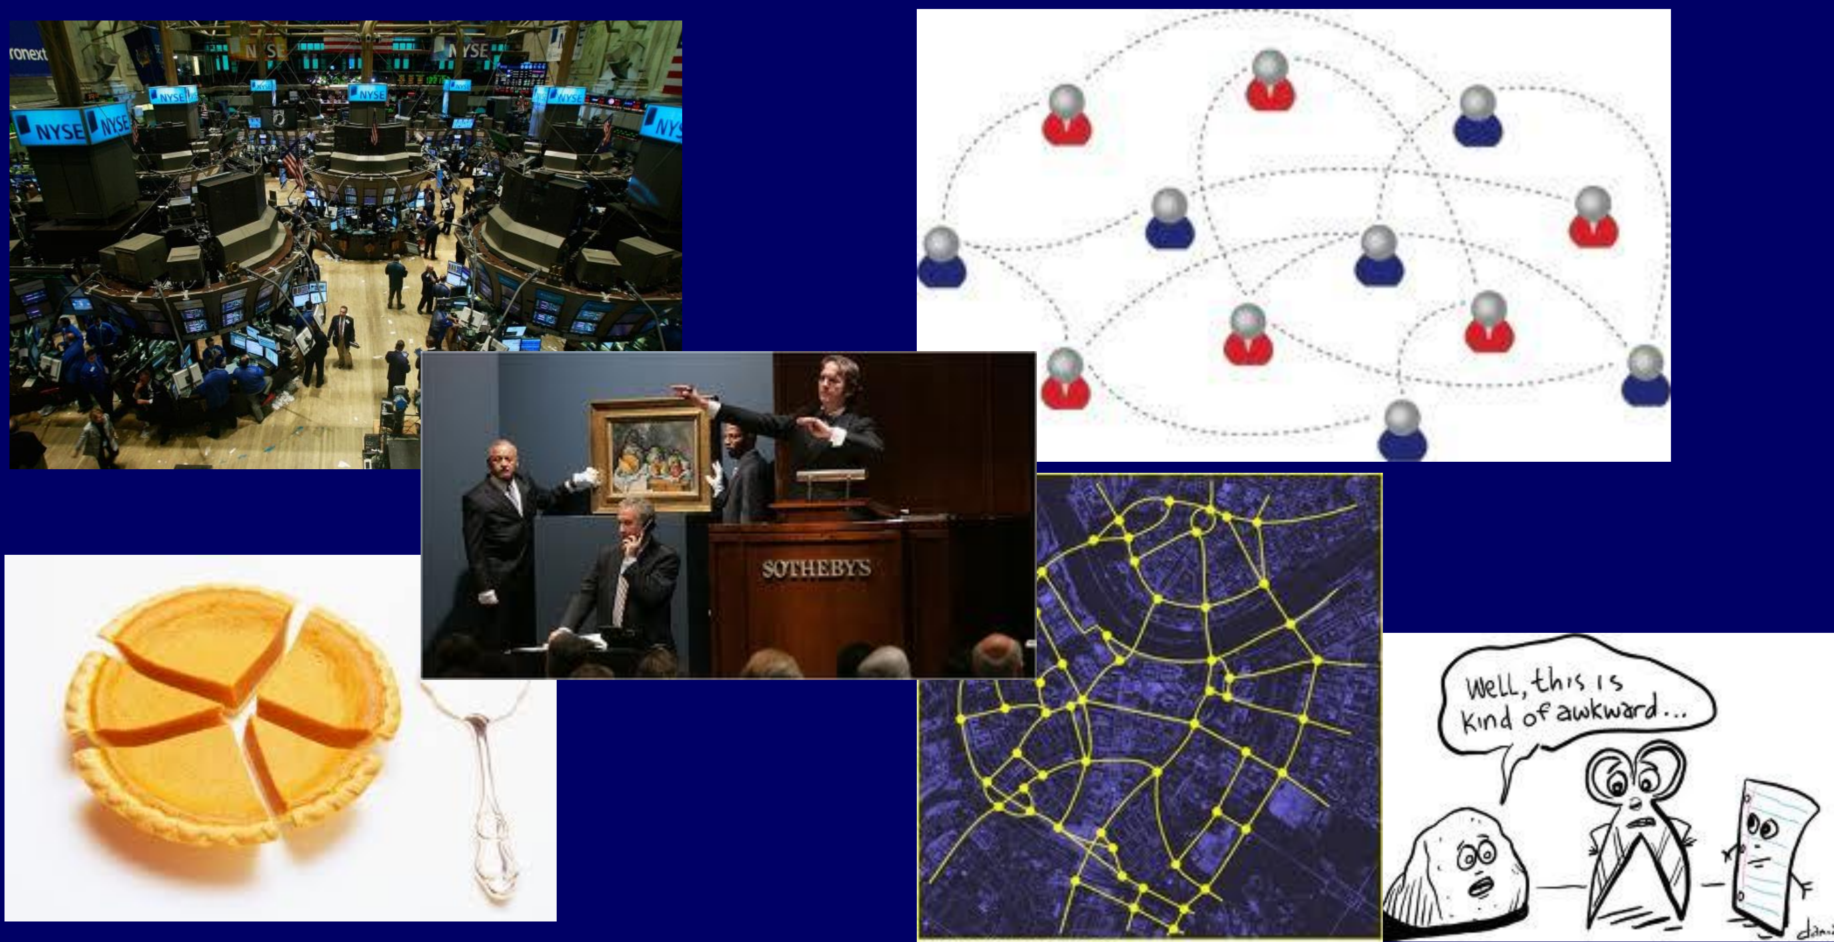

Topics (subject to change)

- Basics of game theory, equilibria
- Connections between learning and equilibria
- Quality of equilibria: price of anarchy
- Social choice: voting, manipulation
- Mechanism design: designing rules of the game to achieve desired outcome, auctions.
- Fair division
- Matching markets

Admin

Free online book: “Algorithmic game theory”

Course requirements:

- 4 homeworks [60%]
- Course project [20%]
  - Could be reading and explaining a recent paper related to class topics
  - Could be theoretically investigating a question related to class topics
  - Could be conducting an experimental investigation
- Class participation including helping to grade one homework [20%]

03/18/24

A Basic Introduction to Game Theory

[Readings: Ch. 1.1-1.3 of AGT book]

Game theory

- Field developed by economists to study social & economic interactions.
  - Wanted to understand why people behave the way they do in different economic situations. Effects of incentives. Rational explanation of behavior.
- "Game" = interaction between parties with their own interests. Could be called "interaction theory".
- Big in CS for understanding large systems:
  - Internet routing, social networks, e-commerce
  - Problems like spam etc.

Setting

- Have a collection of participants, or *players*.
- Each has a set of choices, or *strategies* for how to play/behave.
- Combined behavior results in *payoffs* (satisfaction level) for each player.

Most examples today will involve just 2 players (which will make them easier to picture, as will become clear in a moment…)

Example: walking on the sidewalk

- What side of ~~sidewalk should I walk on~~ street to drive on?
- Two options for you (left or right). Same for person walking towards you.
- Can describe payoffs in matrix:

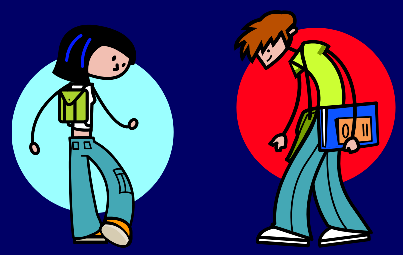

you

person walking towards you

	Left	Right
Left	(1,1)	(-1,-1)
Right	(-1,-1)	(1,1)

Your payoff for RR

His payoff for RR

Key notion

Could be randomized

- "Nash Equilibrium": pair of strategies such that each player is playing a best-response to the other. Neither has an incentive to change.

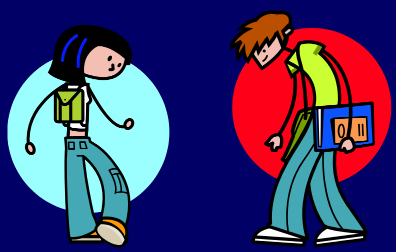

you

	Left	Right
Left	(1,1)	(-1,-1)
Right	(-1,-1)	(1,1)

person walking towards you

Your payoff for RR

His payoff for RR

Example: prisoner's dilemma

- Consider two companies deciding whether to install pollution controls.
- Imagine pollution controls cost $4 but improve everyone's environment by $3

	control	don't control
control	(2,2)	(-1,3)
don't control	(3,-1)	(0,0)

For both, defecting is dominant strategy

What do equilibria look like here?

Example: prisoner's dilemma

- Consider two companies deciding whether to install pollution controls.
- Imagine pollution controls cost $4 but improve everyone's environment by $3

	control	don't control
control	(2,2)	(-1,3)
don't control	(3,-1)	(0,0)

For both, defecting is dominant strategy

Need to add extra incentives to get good overall behavior.

Example: matching pennies / penalty shot

- Shooter can choose to shoot left or shoot right.
- Goalie can choose to dive left or dive right.
- If goalie guesses correctly, (s)he saves the day. If not, it's a goooooaaaaall! Vice-versa for shooter.

goalie

shooter

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

GOAALLL!!!

No deterministic equilibrium

Example: matching pennies / penalty shot

- Shooter can choose to shoot left or shoot right.
- Goalie can choose to dive left or dive right.
- If goalie guesses correctly, (s)he saves the day. If not, it's a goooooaaaaall! Vice-versa for shooter.

goalie

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

shooter

GOAALLL!!!

Each playing 50/50 is a Nash equilibrium

Nash (1950)

- Proved that if you allow *randomized* (mixed) strategies then every game has at least one equilibrium.
- I.e., a pair of (randomized) strategies that is stable in the sense that each is a best response to the other in terms of expected payoff.
- For this, and its implications, Nash received the Nobel prize.

Game theory terminology

- Rows and columns called pure strategies.
- Randomized algs called mixed strategies.

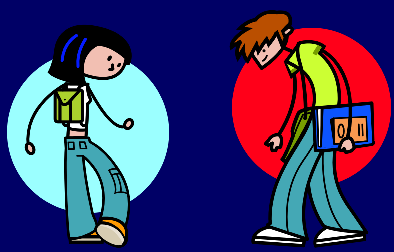

	Left	Right
Left	(1,1)	(-1,-1)
Right	(-1,-1)	(1,1)

you

person walking towards you

Your payoff for RR

His payoff for RR

Game theory terminology

- Rows and columns called pure strategies.
- Randomized algs called mixed strategies.
- Often describe in terms of 2 matrices R, C.

R

1	-1
-1	1

C

1	-1
-1	1

(p,q) is Nash equilib if $p^T R q \geq e_i^T R q$ $\forall i$ and $p^T C q \geq p^T C e_j$ $\forall j$.

Basic facts

- (p,q) is NashEq if $p^T R q \geq e_i^T R q \ \forall i$, $p^T C q \geq p^T C e_j \ \forall j$.
- $\Rightarrow$ for all i s.t. $p_i > 0$ we have $e_i^T R q = \max_{i'} e_{i'}^T R q$
- $\Rightarrow$ for all j s.t. $q_j > 0$ we have $p^T C e_j = \max_{j'} p^T C e_{j'}$

R

1	-1
-1	1

C

1	-1
-1	1

NE can do strange things

- Braess paradox:
  - Road network, traffic going from s to t.
  - travel time as function of fraction x of traffic on a given edge.

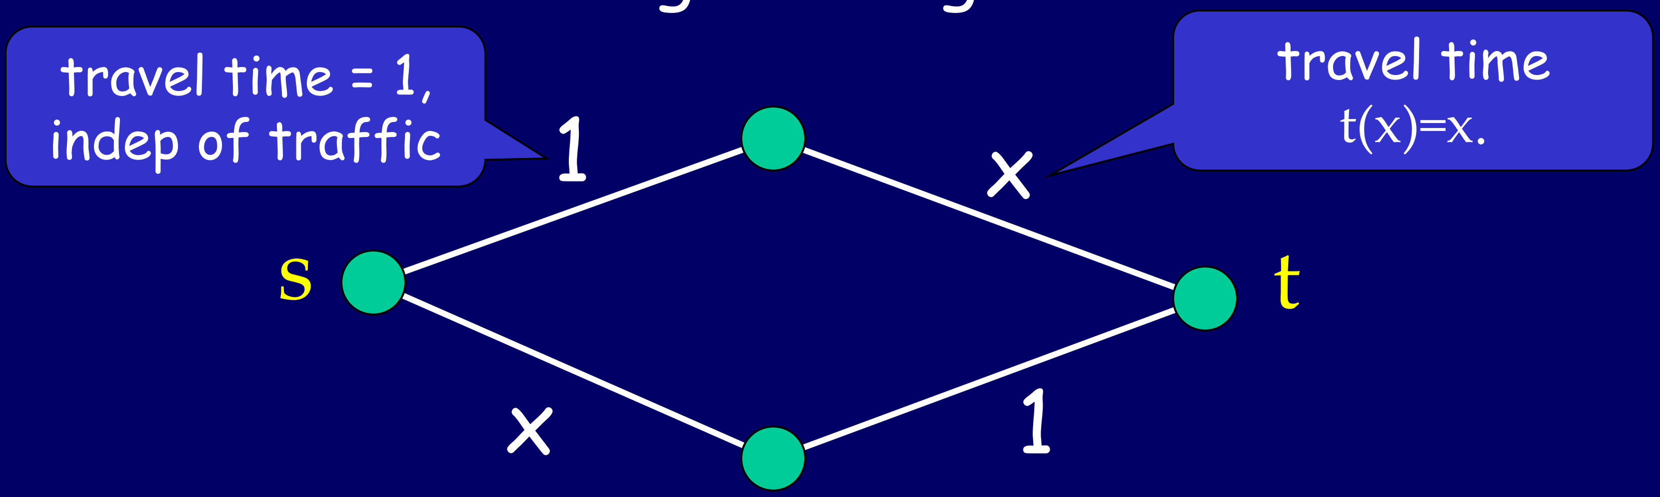

Fine. NE is 50/50. Travel time = 1.5

NE can do strange things

- Braess paradox:
  - Road network, traffic going from s to t.
  - travel time as function of fraction x of traffic on a given edge.

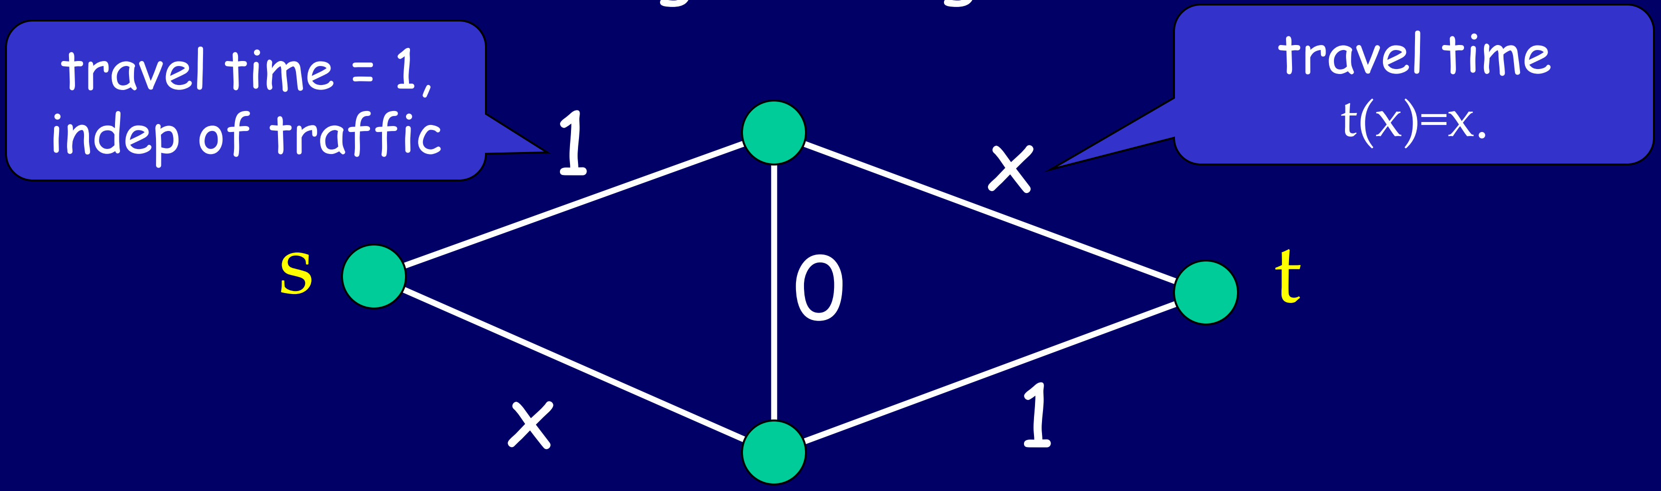

Add new superhighway. NE: everyone uses zig-zag path. Travel time = 2.

2-Player Zero-Sum games

- Zero-sum games are the special case of purely-competitive 2-player games.
  - Recall: an entry (x,y) means: x = payoff to row player, y = payoff to column player. "Zero sum" means that y = -x.
- E.g., matching pennies / penalty shot:

goalie

shooter

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

GOAALLL!!!

Minimax-optimal strategies

- Minimax optimal strategy is a (randomized) strategy that has the best guarantee on its expected gain, over choices of the opponent. [maximizes the minimum]
- I.e., the thing to play if your opponent knows you well.

goalie

shooter \ goalie	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

shooter

GOAALLL!!!

Minimax optimal for both players is 50/50. Gives expected gain of ½ for shooter, -½ for goalie. Any other is worse.

Minimax-optimal strategies

- How about penalty shot with goalie who's weaker on the left?

  Say shooter uses (p,1-p).
  - If goalie dives left, shooter gets p/2 + 1-p = 1 – p/2.
  - If goalie dives right, shooter gets p.
  - Maximize minimum by setting equal.
  - Gives p = 2/3.

	Left	Right
Left	$(\frac{1}{2},-\frac{1}{2})$	(1,-1)
Right	(1,-1)	(0,0)

goalie

shooter

GOAALLL!!!

50/50

Minimax-optimal strategies

- How about penalty shot with goalie who's weaker on the left?

  Minimax optimal for shooter is (2/3,1/3).
  Guarantees expected gain at least 2/3.
  Minimax optimal for goalie is also (2/3,1/3).
  Guarantees expected loss at most 2/3.

goalie

shooter \ goalie	Left	Right
Left	$(\frac{1}{2},-\frac{1}{2})$	(1,-1)
Right	(1,-1)	(0,0)

GOAALLL!!!

50/50

Minimax-optimal strategies

- Can solve for minimax optimal strategy using Linear Programming:

  Variables **p**, v.

  Maximize v subject to:

  - **p** $\cdot$ $R_j \geq$ v, for all j. (R_j is jth column of R)
  - **p** is legal prob dist ($p_i \geq 0$, $\sum_i p_i = 1$).

		Left	Right
shooter	Left	$(\frac{1}{2},-\frac{1}{2})$	(1,-1)
	Right	(1,-1)	(0,0)

goalie

GOAALLL!!!

50/50

Minimax Theorem (von Neumann 1928)

- Every 2-player zero-sum game has a unique value V.
- Minimax optimal strategy for Row guarantees Row's expected gain at least V.
- Minimax optimal strategy for Col guarantees Col's expected loss at most V.

Counterintuitive: Means it doesn't hurt to publish your strategy if both players are optimal. (Borel had proved for symmetric 5x5 but thought was false for larger games)

Nash $\Rightarrow$ Minimax

- Nash's theorem actually gives minimax thm as a corollary.
  - Pick some NE and let V = value to row player in that equilibrium.
  - Since it's a NE, neither player can do better even knowing the (randomized) strategy their opponent is playing.
  - So, they're each playing minimax optimal.

Nash $\Rightarrow$ Minimax

- On the other hand, for minimax, also have very constructive, algorithmic arguments:
  - Can solve for minimax optimum using linear programming in time poly(n) (n = size of game)
  - Have adaptive procedures that in repeated play guarantee to approach/beat best fixed strategy in hindsight
- But for Nash, no efficient procedures to find: NP-hard to find equilib with special properties, PPAD-hard just to find one.

Can use notion of minimax optimality to explain bluffing in poker

Simplified Poker (Kuhn 1950)

- Two players A and B.
- Deck of 3 cards: 1,2,3.
- Players ante \$1.
- Each player gets one card.
- A goes first. Can bet \$1 or pass.
  - If A bets, B can call or fold.
  - If A passes, B can bet \$1 or pass.
    - If B bets, A can call or fold.
- High card wins (if no folding). Max pot \$2.

- Two players A and B. 3 cards: 1,2,3.
- Players ante $1. Each player gets one card.
- A goes first. Can bet $1 or pass.
  - If A bets, B can call or fold.
  - If A passes, B can bet $1 or pass.
    - If B bets, A can call or fold.

Writing as a Matrix Game

- For a given card, A can decide to
  - Pass but fold if B bets. [PassFold]
  - Pass but call if B bets. [PassCall]
  - Bet. [Bet]
- Similar set of choices for B.

Can look at all strategies as a big matrix…

	[FP,FP,CB]	[FP,CP,CB]	[FB,FP,CB]	[FB,CP,CB]
[PF,PF,PC]	0	0	-1/6	-1/6
[PF,PF,B]	0	1/6	-1/3	-1/6
[PF,PC,PC]	-1/6	0	0	1/6
[PF,PC,B]	-1/6	-1/6	1/6	1/6
[B,PF,PC]	-1/6	0	0	1/6
[B,PF,B]	1/6	-1/3	0	-1/2
[B,PC,PC]	1/6	-1/6	-1/6	-1/2
[B,PC,B]	0	-1/2	1/3	-1/6
	0	-1/3	1/6	-1/6

And the minimax optimal strategies are…

- A:
  - If hold 1, then 5/6 PassFold and 1/6 Bet.
  - If hold 2, then $\frac{1}{2}$ PassFold and $\frac{1}{2}$ PassCall.
  - If hold 3, then $\frac{1}{2}$ PassCall and $\frac{1}{2}$ Bet.

  Has both bluffing and underbidding…
- B:
  - If hold 1, then 2/3 FoldPass and 1/3 FoldBet.
  - If hold 2, then 2/3 FoldPass and 1/3 CallPass.
  - If hold 3, then CallBet

Minimax value of game is –1/18 to A.

How to prove existence of NE

- Proof will be non-constructive.
- Notation:
  - Assume an nxn matrix.
  - Use $(p_1,\ldots,p_n)$ to denote mixed strategy for row player, and $(q_1,\ldots,q_n)$ to denote mixed strategy for column player.

Proof

- We'll start with Brouwer's fixed point theorem.
  - Let S be a bounded convex region in R^n and let $f:S \to S$ be a continuous function.
  - Then there must exist $x \in S$ such that $f(x)=x$.
  - x is called a "fixed point" of f.
- Simple case: S is the interval [0,1].
- We will care about:
  - S = {(p,q): p,q are legal probability distributions on 1,...,n}. I.e., $S = \text{simplex}_n \times \text{simplex}_n$

Proof (cont)

- $S = \{(p,q)$: p,q are mixed strategies$\}$.
- Want to define $f(p,q) = (p',q')$ such that:
  - f is continuous. This means that changing p or q a little bit shouldn't cause p' or q' to change a lot.
  - Any fixed point of f is a Nash Equilibrium.
- Then Brouwer will imply existence of NE.

Try #1

- What about $f(p,q) = (p',q')$ where p' is best response to q, and q' is best response to p?
- Problem: not continuous:
  - E.g., penalty shot: If $p = (0.51, 0.49)$ then $q' = (1,0)$. If $p = (0.49,0.51)$ then $q' = (0,1)$.

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

Try #1

- What about $f(p,q) = (p',q')$ where p' is best response to q, and q' is best response to p?
- Problem: not continuous:
  - E.g., penalty shot: If $p = (0.51, 0.49)$ then $q' = (1,0)$. If $p = (0.49,0.51)$ then $q' = (0,1)$.

$$R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \qquad C = \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}$$

Try #1

- What about $f(p,q) = (p',q')$ where p' is best response to q, and q' is best response to p?
- Problem: also not necessarily well-defined:
  - E.g., if $p = (0.5,0.5)$ then q' could be anything.

R =

0	1
1	0

C =

0	-1
-1	0

Instead we will use…

$p^T C q'$

- $f(p,q) = (p',q')$ such that:
  - q' maximizes [(expected gain wrt p) - $||q-q'||^2$]
  - p' maximizes [(expected gain wrt q) - $||p-p'||^2$]

$p'^T R q$

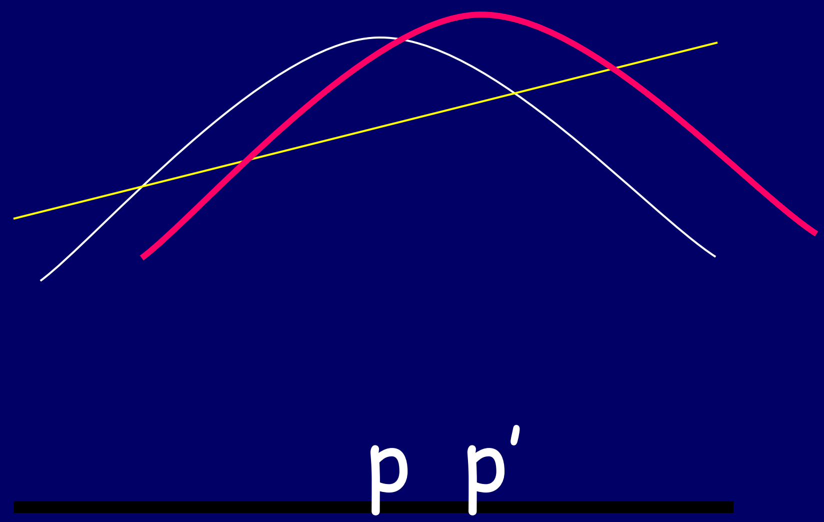

Note: quadratic + linear = quadratic.

Instead we will use…

$p^T C q'$

- $f(p,q) = (p',q')$ such that:
  - q' maximizes [(expected gain wrt p) - $||q-q'||^2$]
  - p' maximizes [(expected gain wrt q) - $||p-p'||^2$]

$p'^T R q$

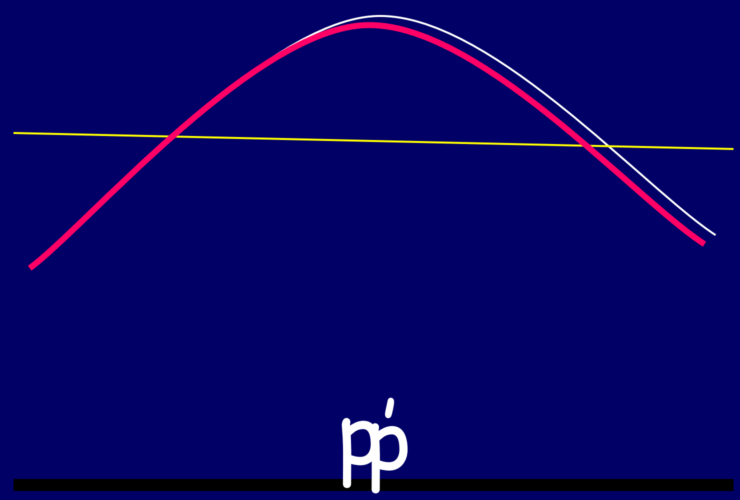

Note: quadratic + linear = quadratic.

Instead we will use…

- $f(p,q) = (p',q')$ such that:
  - q' maximizes [(expected gain wrt p) - $||q-q'||^2$]
  - p' maximizes [(expected gain wrt q) - $||p-p'||^2$]

- f is well-defined and continuous since quadratic has unique maximum and small change to p,q only moves this a little.
- Also fixed point = NE. (even if tiny incentive to move, will move little bit).
- So, that's it!

Algorithmic Game Theory

Algorithmic issues in game theory:

- Computing equilibria / approximate equilibria in different kinds of games
- Understanding quality of equilibria in load-balancing, network-design, routing, machine scheduling…
- Analyzing dynamics of simple behaviors or adaptive (learning) algorithms: quality guarantees, convergence,…
- Design issues: constructing rules so that game will (ideally) have dominant-strategy equilibria with good properties.

End of Game Theory Intro