Two Applications of Learning Theory in Algorithmic Game Theory

Tim Roughgarden (Stanford)

Two Case Studies

Case Study #1: approximation guarantees for game-theoretic equilibria ("price of anarchy").

- Issue: Non-existence, intractability of Nash equilibria.
- Solution: prove guarantees for no-regret learners.

Case Study #2: optimal auction design.

- Issue: traditional reliance on a common prior.
- Solution: learn a near-optimal auction from samples.

Two Case Studies

Case Study #1: approximation guarantees for game-theoretic equilibria ("price of anarchy").

- *Issue: Non-existence, intractability of Nash equilibria.*
- *Solution: prove guarantees for no-regret learners.*

Case Study #2: optimal auction design.

- Issue: traditional reliance on a common prior.
- Solution: learn a near-optimal auction from samples.


FIRE DEPT.
39
FIRE DEPT.
THE PRICE OF ANARCHY
IS IT OKAY TO BE SELFISH?

The Price of Anarchy

Network with 2 players:

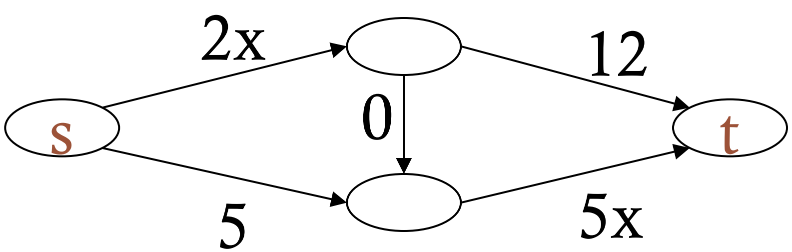

The Price of Anarchy

Nash Equilibrium:

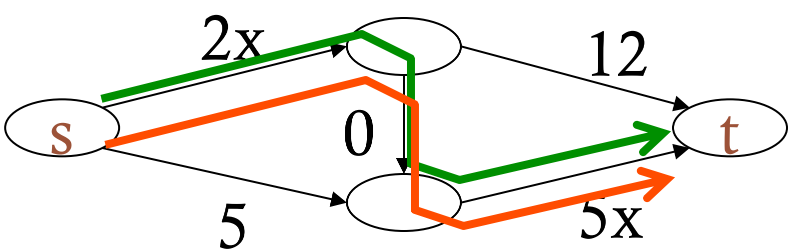

cost = 14+14 = 28

The Price of Anarchy

Nash Equilibrium:

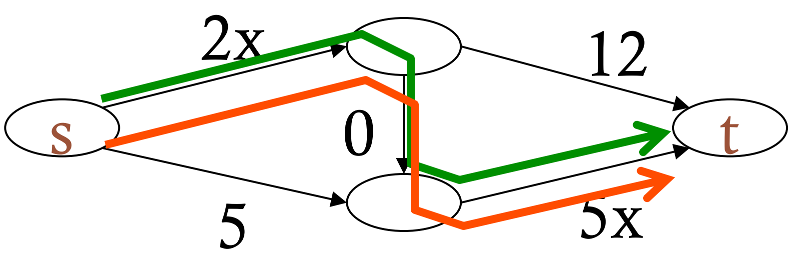

To Minimize Cost:

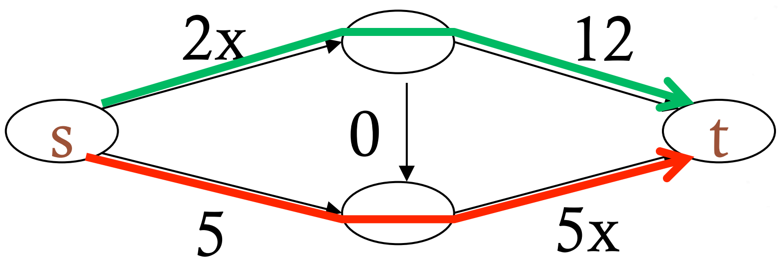

cost = 14+14 = 28

cost = 14+10 = 24

Price of anarchy (POA) = 28/24 = 7/6.

- if multiple equilibria exist, look at the *worst* one

The Price of Anarchy of Health Care

Dr. Vincent Knight, Professor Paul Harper
Cardiff University, Cardiff, UK
knightva@cf.ac.uk

12 July 2011

Health Modelling Centre Cymru

basketball.pdf - Adobe Reader
File Edit View Document Tools Window Help
100% 1 / 9 Find

The price of anarchy in basketball

Brian Skinner
School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455

(Dated: January 18, 2010)

Optimizing the performance of a basketball offense may be viewed as a network problem, wherein each play represents a "pathway" through which the ball and players may move from origin (the in-bounds pass) to goal (the basket). Effective field goal percentages from the resulting shot attempts can be used to characterize the efficiency of each pathway. Inspired by recent discussions of the "price of anarchy" in traffic networks, this paper makes a formal analogy between a basketball offense and a simplified traffic network. The analysis suggests that there may be a significant difference between taking the highest-percentage shot each time down the court and playing the most efficient possible game. There may also be an analogue of Braess's Paradox in basketball, such that removing a key player from a team can result in the improvement of the team's offensive efficiency.

-ph] 18 Jan 2010

I. INTRODUCTION

In its essence, basketball is a network problem. Each possession has a definite starting point (the sideline or baseline in-bounds pass) and a definite goal (putting the ball in the basket). Further, each possession takes place through a particular "pathway": the sequence of player movements and passes leading up to the shot attempt. When a coach diagrams a play for his/her players, he/she is essentially instructing them to move the ball through a particular pathway in order to reach the goal. If we think of a basketball offense as a network of possibilities for moving from

6:50 AM
9/23/2011

POA Bounds Without Convergence

Meaning of a POA bound: *if* the game is at an equilibrium, *then* outcome is near-optimal.

Problem: what if can't reach an equilibrium?

- non-existence (pure Nash equilibria)
- intractability (mixed Nash equilibria) [Daskalakis/Goldberg/Papadimitriou 06], [Chen/Deng/Teng 06],[Etessami/Yannakakis 07]

Worry: are our POA bounds "meaningless"?

Robust POA Bounds

High-Level Goal: worst-case bounds that apply *even to non-Nash equilibrium outcomes!*

- best-response dynamics, pre-convergence
  - [Mirrokni/Vetta 04], [Goemans/Mirrokni/Vetta 05], [Awerbuch/Azar/Epstein/Mirrokni/Skopalik 08]
- correlated equilibria
  - [Christodoulou/Koutsoupias 05]
- no-regret learners (“coarse correlated equilibria”)
  - [Blum/Even-Dar/Ligett 06], [Blum/Hajiaghayi/Ligett/Roth 08]

POA Bounds Without Convergence

Theorem: [Roughgarden 09] most known POA bounds hold for all no-regret sequences (not just for Nash equilibria).

- eludes non-existence/intractability critiques.

POA Bounds Without Convergence

Theorem: [Roughgarden 09] most known POA bounds hold for all no-regret sequences (not just for Nash equilibria).

- eludes non-existence/intractability critiques.

Part I: [extension theorem] every POA bound proved for pure Nash equilibria *in a prescribed way* extends automatically to all no-regret sequences.

Part II: most known POA bounds can be proved in this way (so extension theorem applies).

Extension Theorems

permissive equilibrium concept (e.g., no-regret outcomes)

what we care about

Extension Theorems

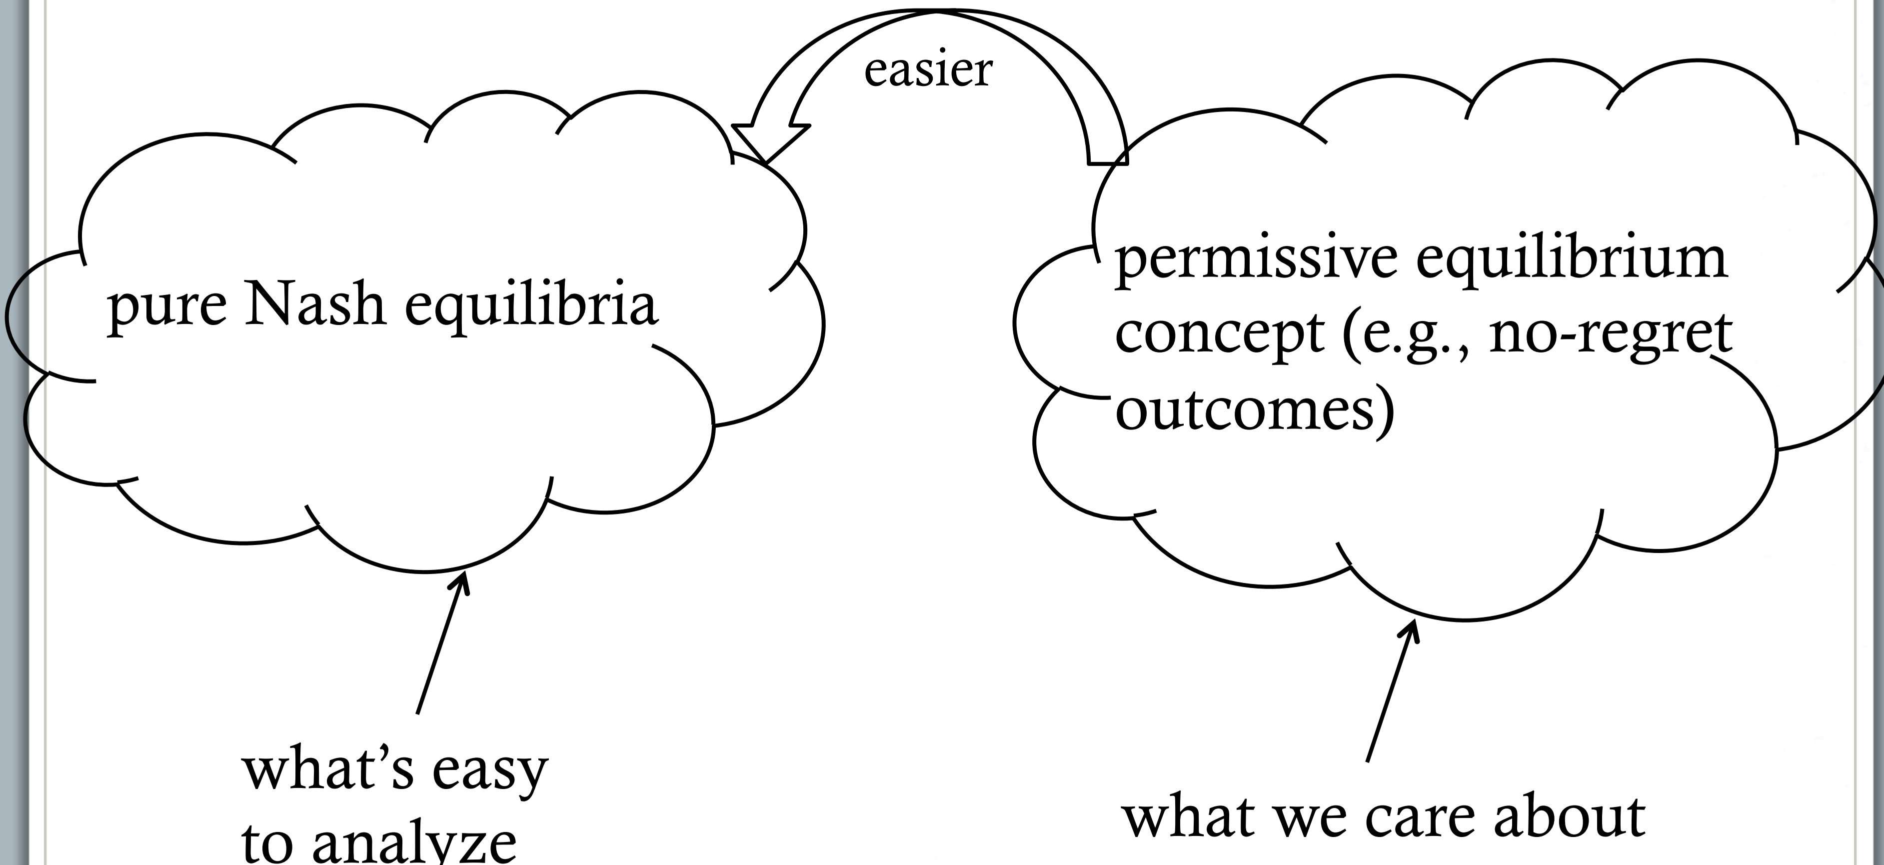

Extension Theorems

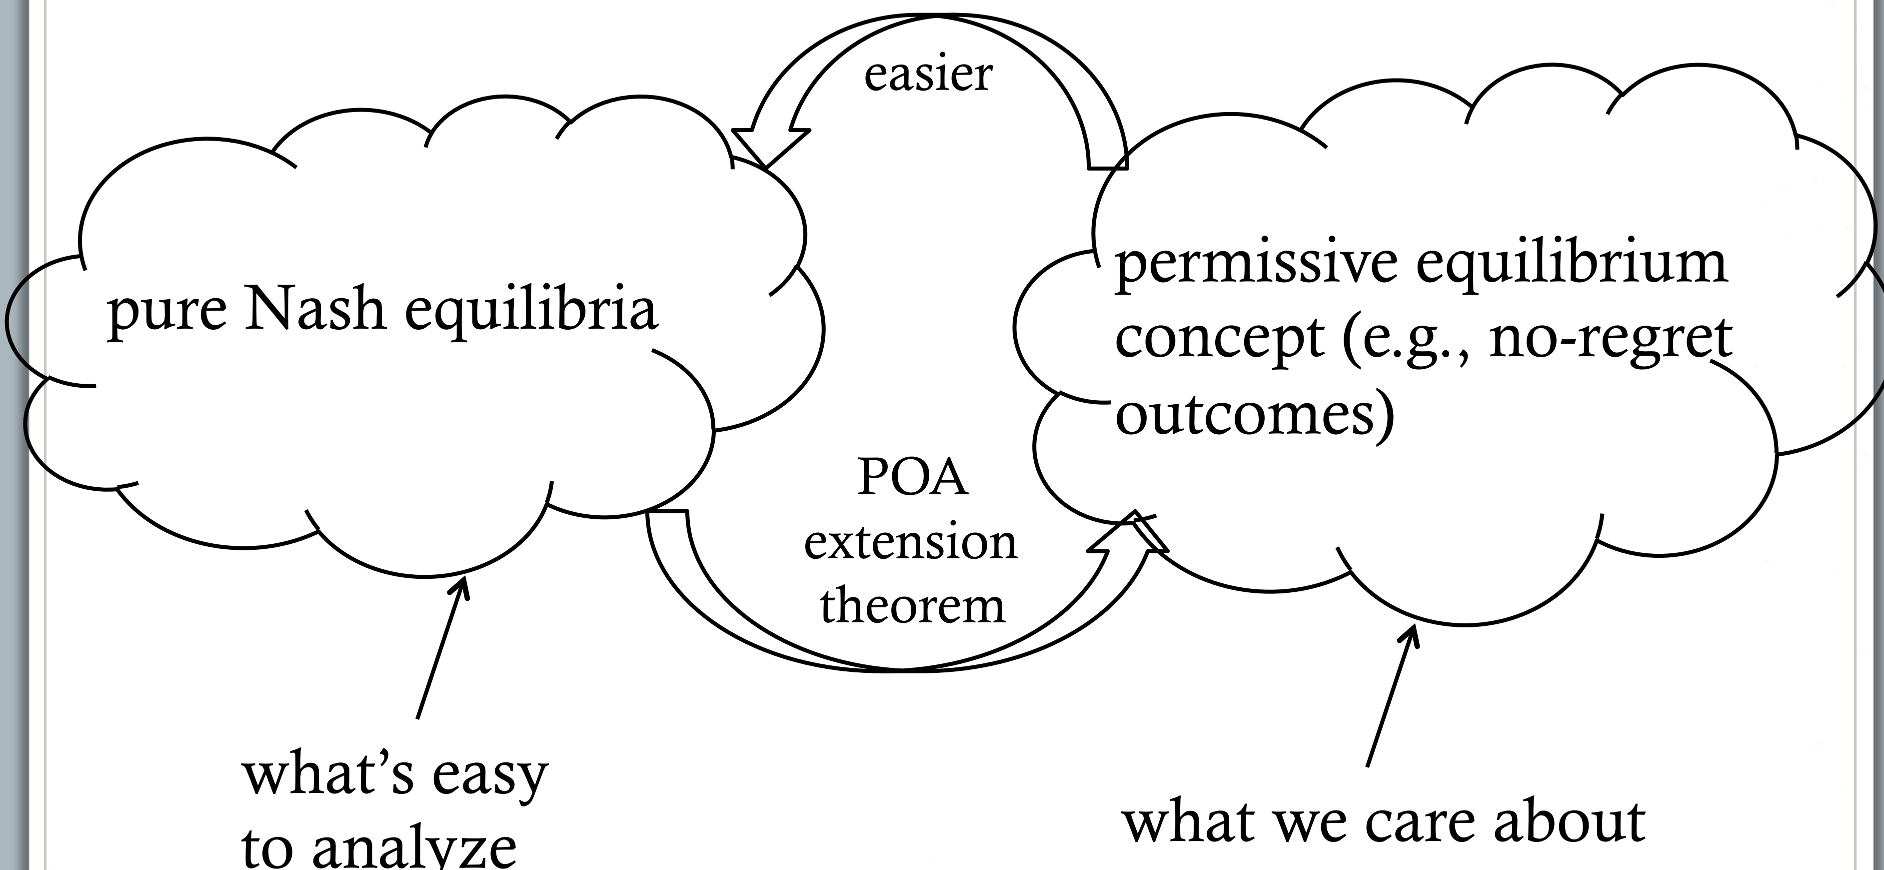

The Math

- n players, each picks a strategy s_i
- player i incurs a cost $C_i(\mathbf{s})$

Assume: objective function is $\mathrm{cost}(\mathbf{s}) := \sum_i C_i(\mathbf{s})$

The Math

- n players, each picks a strategy s_i
- player i incurs a cost $C_i(\mathbf{s})$

Assume: objective function is $\mathrm{cost}(\mathbf{s}) := \Sigma_i C_i(\mathbf{s})$

To Bound POA: (let $\mathbf{s}$ =a Nash eq; $\mathbf{s}^*$ =optimal)

$$
\begin{aligned}
\mathrm{cost}(\mathbf{s}) &= \Sigma_i C_i(\mathbf{s}) && \text{[defn of cost]} \\
&\leq \Sigma_i C_i(s^*_i, s_{-i}) && \text{[s a Nash eq]}
\end{aligned}
$$

Smooth Games

Key Definition: A game is *(λ, μ)-smooth* if, for every pair $\mathbf{s},\mathbf{s}^*$ of outcomes ($\lambda > 0$; $\mu < 1$):

$$\Sigma_i C_i(s^*_i, s_{-i}) \leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s}) \quad [(*)]$$

Implies: $\text{cost}(\mathbf{s}) \leq \Sigma_i C_i(s^*_i, s_{-i})$ [**s** a Nash eq]

$\leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s})$ [(*)]

Smooth Games

Key Definition: A game is *(λ, μ)-smooth* if, for every pair $\mathbf{s},\mathbf{s}^*$ of outcomes ($\lambda > 0$; $\mu < 1$):

$$\Sigma_i C_i(s^*_i, s_{-i}) \leq \lambda \bullet \mathrm{cost}(\mathbf{s}^*) + \mu \bullet \mathrm{cost}(\mathbf{s}) \quad [(*)]$$

Implies: $\mathrm{cost}(\mathbf{s}) \leq \Sigma_i C_i(s^*_i, s_{-i})$ [s a Nash eq]

$\leq \lambda \bullet \mathrm{cost}(\mathbf{s}^*) + \mu \bullet \mathrm{cost}(\mathbf{s})$ [(*)]

So: POA (of pure Nash equilibria) $\leq \lambda/(1-\mu)$.

Note: only needed (*) to hold in special case where $\mathbf{s}$ = a Nash equilibrium and $\mathbf{s}^*$ = optimal outcome.

Some Smoothness Bounds

- selfish routing + related models [Roughgarden/Tardos 00], [Perakis 04], [Correa/Schulz/Stier Moses 05], [Awerbuch/Azar/Epstein 05], [Christodoulou/Koutsoupias 05], [Aland/Dumrauf/Gairing/Monien/Schoppmann 06], [Roughgarden 09], [Bhawalkar/Gairing/Roughgarden 10]
- submodular maximization games [Vetta 02], [Marden/Roughgarden 10]
- coordination mechanisms [Cole/Gkatzelis/Mirrokni 10]
- auctions [Christodoulou/Kovacs/Schapira 08], [Lucier/Borodin 10], [Bhawalkar/Roughgarden 11], [Caragiannis/Kaklamanis/Kanellopolous/Kyropoulou/Lucier/Paes Leme/Tardos 12], [Lucier/Singer/Syrgkanis/Tardos 11], [Markakis/Telelis 12], [Paes Leme/Syrgkanis/Tardos 12], [Bhawalkar/Roughgarden 12], [Feldman/Fu/Gravin/Lucier 13], [Syrgkanis/Tardos 13], [de Keijzer/Markakis/Schaefer/Telelis 13], [Duetting/Henzinger/Starnberger 13], [Babaioff/Lucier/Nisan/Paes Leme 13], [Devanur/Morgenstern/Syrgkanis 15], …

Canonical Example

Claim [Christodoulou/Koutsoupias 05]: routing games with affine cost functions are (5/3,1/3)-smooth.

- for all integers y,z: $y(z+1) \leq (5/3)y^2 + (1/3)z^2$
- so: $ay(z+1) + by \leq (5/3)[ay^2 + by] + (1/3)[az^2 + bz]$
  - for all integers y,z and $a,b \geq 0$
- so: $\Sigma_e [a_e(x_e+1) + b_e)x_e^*] \leq (5/3)\ \Sigma_e [(a_e x_e^* + b_e)x_e^*]$ $+ (1/3)\ \Sigma_e [(a_e x_e + b_e)x_e]$
- so: $\Sigma_i C_i(s^*_i, s_{-i}) \leq (5/3)\bullet \text{cost}(\mathbf{s}^*) + (1/3)\bullet \text{cost}(\mathbf{s})$

An Out-of-Equilibrium Bound

Theorem: [Roughgarden 09] in a (λ, μ)-smooth game, the average cost of every no-regret sequence is at most

$$[\lambda/(1-\mu)] \bullet \text{cost of optimal outcome.}$$

(the same bound as for pure Nash equilibria!)

No-Regret Sequences

Definition: a sequence $s^1, s^2, \ldots, s^T$ of outcomes of a game is *no-regret* if:

- for each i, each (time-invariant) deviation q_i:

$$(1/T)\ \Sigma_t C_i(s^t) \leq (1/T)\ \Sigma_t C_i(q_i, s^t_{-i})\ \ [+\ o(1)]$$

Fact: simple hedging strategies can be used by players to enforce this (as T grows large).

- [Blackwell 56], [Hannan 57], …, [Freund/Schapire 99], …

Smooth => No-Regret Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ , μ)-smooth:

$\Sigma_t \, \text{cost}(\mathbf{s}^t) \; = \; \Sigma_t \, \Sigma_i \, C_i(\mathbf{s}^t)$ [defn of cost]

Smooth => No-Regret Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ , μ)-smooth:

$\Sigma_t \,\mathrm{cost}(\mathbf{s}^t) = \Sigma_t \Sigma_i C_i(\mathbf{s}^t)$ [defn of cost]

$= \Sigma_t \Sigma_i [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}]$ [$\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})$]

Smooth => No-Regret Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ , μ)-smooth:

$\Sigma_t \, \mathrm{cost}(\mathbf{s}^t) \;=\; \Sigma_t \, \Sigma_i \, C_i(\mathbf{s}^t)$ [defn of cost]

$= \; \Sigma_t \, \Sigma_i \, [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}]$ $[\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})]$

$\leq \; \Sigma_t \, [\, \lambda \bullet \mathrm{cost}(\mathbf{s}^*) + \mu \bullet \mathrm{cost}(\mathbf{s}^t)] + \Sigma_i \, \Sigma_t \, \Delta_{i,t}$ [(*)]

Smooth => No-Regret Bound

- notation: $\mathbf{s}^1, \mathbf{s}^2, \ldots, \mathbf{s}^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ , μ)-smooth:

$$\Sigma_t \operatorname{cost}(\mathbf{s}^t) = \Sigma_t \Sigma_i C_i(\mathbf{s}^t) \qquad \text{[defn of cost]}$$

$$= \Sigma_t \Sigma_i [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}] \qquad [\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})]$$

$$\leq \Sigma_t [\lambda \bullet \operatorname{cost}(\mathbf{s}^*) + \mu \bullet \operatorname{cost}(\mathbf{s}^t)] + \Sigma_i \Sigma_t \Delta_{i,t} \qquad [(*)]$$

No regret: $\Sigma_t \Delta_{i,t} \leq 0$ for each i.

To finish proof: divide through by T.

Two Case Studies

Case Study #1: approximation guarantees for game-theoretic equilibria ("price of anarchy").

- Issue: Non-existence, intractability of Nash equilibria.
- Solution: prove guarantees for no-regret learners.

Case Study #2: optimal auction design.

- *Issue: traditional reliance on a common prior.*
- *Solution: learn a near-optimal auction from samples.*

Myerson's Auction (i.i.d.)

- one seller with one item
- n bidders, bidder i has private valuation v_i
- valuations v_i drawn i.i.d. from known prior F
- goal: maximize seller's expected revenue

Myerson's Auction (i.i.d.)

- one seller with one item
- n bidders, bidder i has private valuation v_i
- valuations v_i drawn i.i.d. from known prior F
- goal: maximize seller's expected revenue
- [Myerson 81] solution = 2nd-price auction + reserve
  - reserve price r = monopoly price for F [i.e., $\text{argmax}_p\, p(1-F(p))$]
  - winner = highest bidder above r (if any)
  - price = maximum of r and 2nd-highest bid

Optimal Single-Item Auctions

[Myerson 81]: characterized the optimal auction, as a function of the prior distributions $F_1,\ldots,F_n$.

- Step 1: transform bids to virtual bids: $b_i \to \varphi_i(b_i)$
  - formula depends on distribution: $\varphi_i(b_i) = b_i - [1 - F_i(b_i)] / f_i(b_i)$
- Step 2: winner = highest positive virtual bid (if any)
- Step 3: price = lowest bid that still would've won

I.i.d. case: 2nd-price auction with monopoly reserve price.

General case: requires full knowledge of $F_1,\ldots,F_n$.

Motivating Question

Question: Does a near-optimal single-item auction require detailed distributional knowledge?

Motivating Question

Question: Does a near-optimal single-item auction require detailed distributional knowledge?

Reformulation: How much data is necessary and sufficient to justify revenue-optimal auction theory?

- “data” = samples from unknown $F_1,\ldots,F_n$
  - inspired by PAC/statistical learning [Vapnik, Valiant, …]
  - Yahoo! example: [Ostrovsky/Schwarz 09]
- benchmark: Myerson’s optimal auction for $F_1,\ldots,F_n$
  - want expected revenue at least $(1-\varepsilon)$ times benchmark

Formalism: Single Buyer

Step 1: seller gets s samples $v_1,\ldots,v_s$ from unknown F

Step 2: seller picks a price $p = p(v_1,\ldots,v_s)$

Step 3: price p applied to a fresh sample v_{s+1} from F

m samples $v_1,\ldots,v_s$ → price $p(v_1,\ldots,v_s)$ → revenue of p on v_{s+1}

valuation v_{s+1}

Goal: design p so that $E_{v_1,\ldots,v_s}[p(v_1,\ldots,v_s)\bullet(1-F(p(v_1,\ldots,v_s))]$ is close to $\max_p[p\bullet(1-F(p)]$ (no matter what F is)

Results for a Single Buyer

1. no assumption on *F*: no finite number of samples yields non-trivial revenue guarantee (for every F)

Results for a Single Buyer

1. no assumption on *F*: no finite number of samples yields non-trivial revenue guarantee (for every F)
2. if *F* is "regular": with s=1, setting $p(v_1) = v_1$ yields a ½-approximation (consequence of [Bulow/Klemperer 96])

Results for a Single Buyer

1. no assumption on F: no finite number of samples yields non-trivial revenue guarantee (for every F)
2. if F is "regular": with s=1, setting $p(v_1) = v_1$ yields a ½-approximation (consequence of [Bulow/Klemperer 96])
3. for regular F, arbitrary ε :
   $\approx (1/\varepsilon)^3$ samples necessary and sufficient for $(1-\varepsilon)$-approximation [Dhangwatnotai/Roughgarden/Yan 10], [Huang/Mansour/Roughgarden 14]
4. for F with a montone hazard rate, arbitrary ε :
   $\approx (1/\varepsilon)^{3/2}$ samples necessary and sufficient for $(1-\varepsilon)$-approximation [Huang/Mansour/Roughgarden 14]

Formalism: Multiple Buyers

Step 1: seller gets s samples $\mathbf{v_1},\ldots,\mathbf{v_s}$ from $\boldsymbol{F} = F_1 \times \cdots \times F_n$

- each $\mathbf{v_i}$ an n-vector (one valuation per bidder)

Step 2: seller picks single-item auction $A = A(\mathbf{v_1},\ldots,\mathbf{v_s})$

Step 3: auction A is run on a fresh sample $\mathbf{v_{s+1}}$ from $\boldsymbol{F}$

m samples $\mathbf{v_1},\ldots,\mathbf{v_s}$ → $\boxed{\text{auction } A(\mathbf{v_1},\ldots,\mathbf{v_s})}$ → revenue of A on $\mathbf{v_{s+1}}$

↑ valuation profile $\mathbf{v_{s+1}}$

Goal: design A so $E_{v_1,\ldots,v_s}[E_{v_{s+1}}[\mathrm{Rev}(A(v_1,\ldots,v_s)(v_{s+1}))]]$ close to OPT

Positive Results

One sample (s=1) still suffices for ¼-approximation

- 2^{nd}-price auction with reserves = samples
- consequence of [Hartline/Roughgarden 09]

Polynomial (in ε^{-1} only) samples still suffice for $(1-\varepsilon)$-approximation if bidders are i.i.d.

- only need to learn the optimal reserve price

Take-away: for these cases,

- modest amount of data (independent of n) suffices
- modest distributional dependence suffices

Negative Results

Theorem: [Cole/Roughgarden 14] at least $\approx n/\sqrt{\varepsilon}$ samples are necessary for (1- ε)-approximation.

- for every sufficiently small constant ε
- even when distributions are truncated exponential distributions

Corollary (of proof): near-optimal auctions require detailed knowledge of the valuation distributions.

Recent Developments

[Morgenstern/Roughgarden 15]

- view a set C of auctions = real-valued functions
  - (from valuation profiles to revenue)
  - see also [Medina/Mohri ICML 14]
- sample complexity bounds reduce to pseudo-dimension bounds

Main Result: in all "single-parameter settings", can learn a (1- ε)-approximate auction from poly samples.

- can choose C to simultaneously have small representation error and small pseudo-dimension

Open Directions

- computationally efficient algorithms for learning near-optimal auctions
  - [Morgenstern/Roughgarden 15] only bounds sample complexity
  - [Cole/Roughgarden 14] efficient for the single-item case

- partial feedback settings [Awerbuch/Kleinberg 03], [Cesa-Bianchi/Gentile/Mansour 13], …

- strategic data providers (e.g. [Cai/Daskalakis/Papadimitriou COLT 15])