

Approximation in Algorithmic Game Theory

Robust Approximation Bounds for Equilibria and Auctions

Tim Roughgarden
Stanford University

Motivation

Clearly: many modern applications in CS involve autonomous, self-interested agents

- motivates noncooperative games as modeling tool

Unsurprising fact: this often makes full optimality hard/impossible.

- equilibria (e.g., Nash) of noncooperative games are typically suboptimal
- auctions lose revenue from strategic behavior
- incentive constraints can make poly-time approximation of NP-hard problems even harder

Approximation in AGT

- The Price of Anarchy (etc.)
 - worst-case approximation guarantees for equilibria
- Revenue Maximization
 - guarantees for auctions in non-Bayesian settings (information-theoretic)
- Algorithm Mechanism Design
 - approximation algorithms robust to selfish behavior (computational)
- Computing Approximate Equilibria
 - e.g., is there a PTAS for computing an approximate Nash equilibrium?

} this talk

} FOCs 2010 tutorial



Price of Anarchy

Price of anarchy: [Koutsoupias/Papadimitriou 99]
quantify inefficiency w.r.t some objective function.

- e.g., *Nash equilibrium*: an outcome such that no player better off by switching strategies

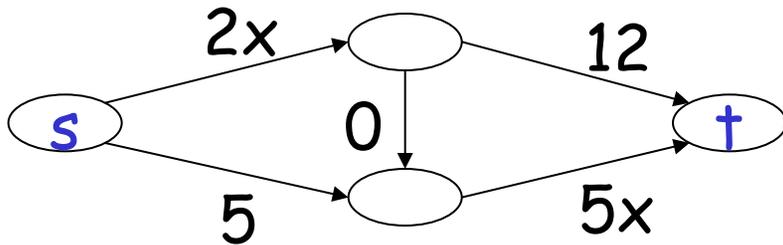
Definition: *price of anarchy (POA)* of a game (w.r.t. some objective function):

$$\frac{\text{equilibrium objective fn value}}{\text{optimal obj fn value}}$$

the closer to 1
the better

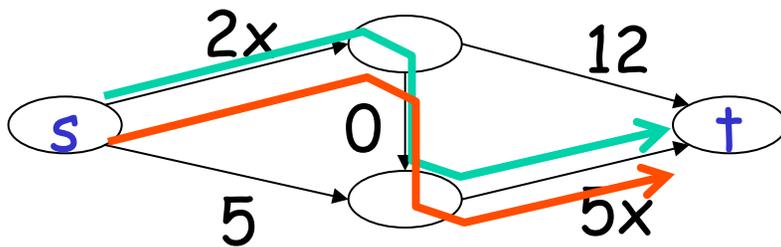
The Price of Anarchy

Network w/2 players:



The Price of Anarchy

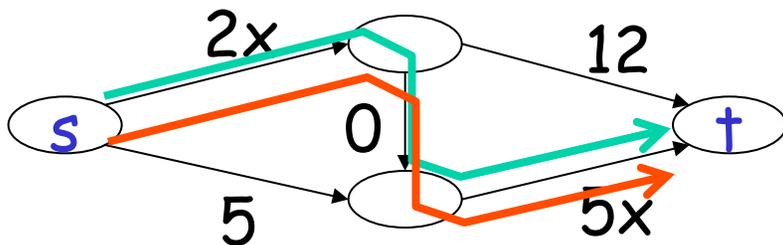
Nash Equilibrium:



$$\text{cost} = 14 + 14 = 28$$

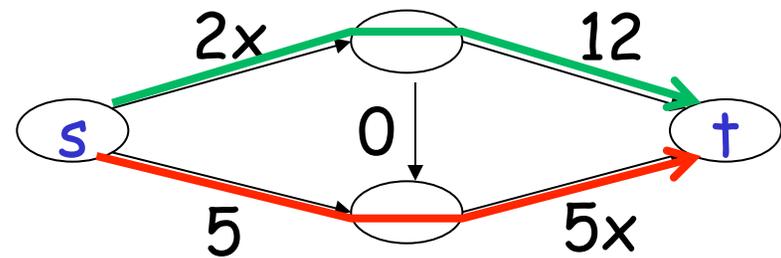
The Price of Anarchy

Nash Equilibrium:



$$\text{cost} = 14 + 14 = 28$$

To Minimize Cost:



$$\text{cost} = 14 + 10 = 24$$

Price of anarchy = $28/24 = 7/6$.

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

The Need for Robustness

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

The Need for Robustness

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

Problem: what if can't reach equilibrium?

- (pure) equilibrium might not exist
- might be hard to compute, even centrally
 - [Fabrikant/Papadimitriou/Talwar], [Daskalakis/ Goldbeg/ Papadimitriou], [Chen/Deng/Teng], etc.
- might be hard to learn in a distributed way

Worry: are our POA bounds "meaningless"?

Robust POA Bounds

- High-Level Goal:** worst-case bounds that apply *even to non-equilibrium outcomes!*
- best-response dynamics, pre-convergence
 - [Mirrokni/Vetta 04], [Goemans/Mirrokn/Vetta 05], [Awerbuch/Azar/Epstein/Mirrokn/Skopalik 08]
 - correlated equilibria
 - [Christodoulou/Koutsoupias 05]
 - coarse correlated equilibria aka "price of total anarchy" aka "no-regret players"
 - [Blum/Even-Dar/Ligett 06], [Blum/Hajiaghayi/Ligett/Roth 08]

Abstract Setup

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

Important Assumption: objective function is
 $\text{cost}(s) := \sum_i C_i(s)$

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

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

Smooth \Rightarrow POA Bound

Next: "canonical" way to upper bound POA
(via a smoothness argument).

- notation: \mathbf{s} = a Nash eq; \mathbf{s}^* = optimal

Assuming (λ, μ) -smooth:

$$\begin{aligned} \text{cost}(\mathbf{s}) &= \sum_i C_i(\mathbf{s}) && \text{[defn of cost]} \\ &\leq \sum_i C_i(\mathbf{s}^*_i, \mathbf{s}_{-i}) && \text{[s a Nash eq]} \\ &\leq \lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \text{cost}(\mathbf{s}) && \text{[(*)]} \end{aligned}$$

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

Why Is Smoothness Stronger?

Key point: to derive POA bound, only needed

$$\sum_i C_i(s_i^*, s_{-i}) \leq \lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s) \quad [(*)]$$

to hold in special case where s = a Nash eq
and s^* = optimal.

Smoothness: requires (*) for *every* pair s, s^*
outcomes.

- even if s is *not* a pure Nash equilibrium

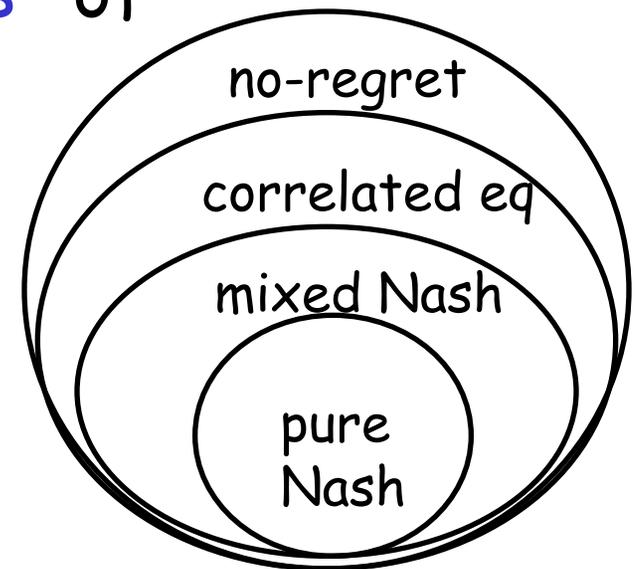
Some Smoothness Bounds

- atomic (unweighted) selfish routing [Awerbuch/Azar/Epstein 05], [Christodoulou/Koutsoupias 05], [Aland/Dumrauf/Gairing/Monien/Schoppmann 06], [Roughgarden 09]
- nonatomic selfish routing [Roughgarden/Tardos 00],[Perakis 04] [Correa/Schulz/Stier Moses 05]
- weighted congestion games [Aland/Dumrauf/Gairing/Monien/Schoppmann 06], [Bhawalkar/Gairing/Roughgarden 10]
- submodular maximization games [Vetta 02], [Marden/Roughgarden 10]
- coordination mechanisms [Cole/Gkatzelis/Mirroknis 10]

Beyond Nash Equilibria

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

- for each player i , each fixed action q_i :
 - average cost player i incurs over sequence no worse than playing action q_i every time
 - if every player uses e.g. "multiplicative weights" then get $o(1)$ regret in poly-time
 - empirical distribution = "*coarse correlated eq*"



An Out-of-Equilibrium Bound

Theorem: [Roughgarden STOC 09]

in a (λ, μ) -smooth game, average cost of every no-regret sequence at most

$[\lambda/(1-\mu)] \times$ cost of optimal outcome.

(the same bound we proved for pure Nash equilibria)

Smooth \Rightarrow No-Regret Bound

- notation: s^1, s^2, \dots, s^T = no regret; s^* = optimal

Assuming (λ, μ) -smooth:

$$\sum_t \text{cost}(s^t) = \sum_t \sum_i C_i(s^t) \quad [\text{defn of cost}]$$

Smooth \Rightarrow No-Regret Bound

- notation: s^1, s^2, \dots, s^T = no regret; s^* = optimal

Assuming (λ, μ) -smooth:

$$\sum_t \text{cost}(s^t) = \sum_t \sum_i C_i(s^t) \quad [\text{defn of cost}]$$

$$= \sum_t \sum_i [C_i(s^*_{-i}, s^t_{-i}) + \Delta_{i,t}] \quad [\Delta_{i,t} := C_i(s^t) - C_i(s^*_{-i}, s^t_{-i})]$$

Smooth \Rightarrow No-Regret Bound

- notation: s^1, s^2, \dots, s^T = no regret; s^* = optimal

Assuming (λ, μ) -smooth:

$$\begin{aligned}\sum_t \text{cost}(s^t) &= \sum_t \sum_i C_i(s^t) && \text{[defn of cost]} \\ &= \sum_t \sum_i [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}] && [\Delta_{i,t} := C_i(s^t) - C_i(s^*_i, s^t_{-i})] \\ &\leq \sum_t [\lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s^t)] + \sum_i \sum_t \Delta_{i,t} && [(\ast)]\end{aligned}$$

Smooth \Rightarrow No-Regret Bound

- notation: s^1, s^2, \dots, s^T = no regret; s^* = optimal

Assuming (λ, μ) -smooth:

$$\begin{aligned}\sum_t \text{cost}(s^t) &= \sum_t \sum_i C_i(s^t) && \text{[defn of cost]} \\ &= \sum_t \sum_i [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}] && [\Delta_{i,t} := C_i(s^t) - C_i(s^*_i, s^t_{-i})] \\ &\leq \sum_t [\lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s^t)] + \sum_i \sum_t \Delta_{i,t} && [(\ast)]\end{aligned}$$

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

To finish proof: divide through by T .

Intrinsic Robustness

Theorem: [Roughgarden STOC 09] for every set C , unweighted congestion games with cost functions restricted to C are *tight*:

$$\text{maximum [pure POA]} = \text{minimum } [\lambda/(1-\mu)]$$

congestion games
w/cost functions in C

(λ, μ) : all such games
are (λ, μ) -smooth

Intrinsic Robustness

Theorem: [Roughgarden STOC 09] for every set C , unweighted congestion games with cost functions restricted to C are *tight*:

$$\begin{array}{l} \text{maximum [pure POA]} = \text{minimum } [\lambda/(1-\mu)] \\ \text{congestion games} \qquad \qquad \qquad (\lambda, \mu): \text{ all such games} \\ \text{w/cost functions in } C \qquad \qquad \qquad \text{are } (\lambda, \mu)\text{-smooth} \end{array}$$

- *weighted* congestion games [Bhawalkar/ Gairing/ Roughgarden ESA 10] and submodular maximization games [Marden/Roughgarden CDC 10] are also tight in this sense

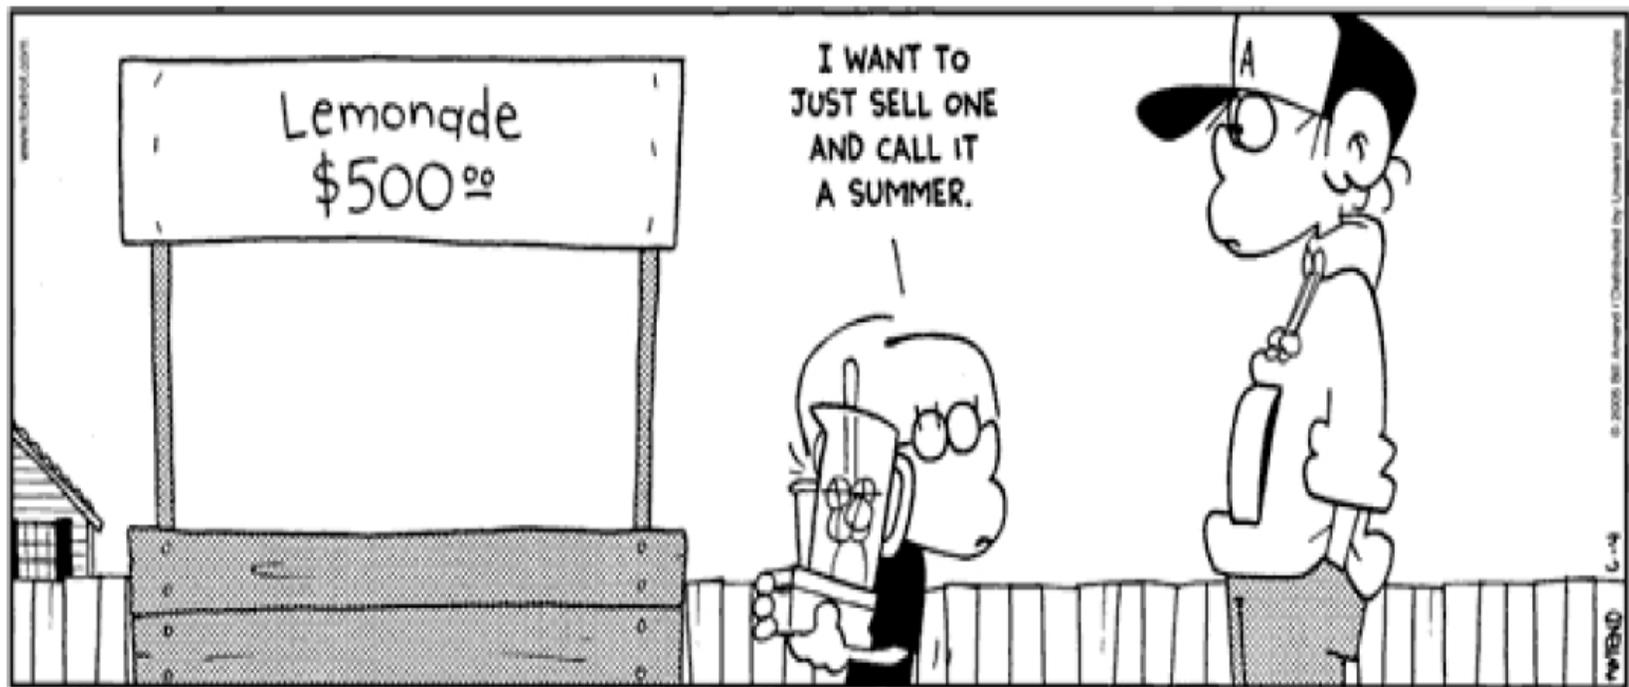
What's Next?

- **beating worst-case POA bounds:** want to reach a non-worst-case equilibrium
 - because of learning dynamics [Charikar/Karloff/Mathieu/Naor/Saks 08], [Kleinberg/Pilouras/Tardos 09], etc.
 - from modest intervention [Balcan/Blum/Mansour], etc.
- **POA bounds for auctions**
 - practical auctions often lack "dominant strategies" (sponsored search, combinatorial auctions, etc.)
 - want bounds on their (Bayes-Nash) equilibria [Christodoulou et al 08], [Paes Leme/Tardos 10], [Bhawalkar/Roughgarden 11], [Hassadim et al 11]

Key Points

- **smoothness:** a “canonical way” to bound the price of anarchy (for pure equilibria)
- **robust POA bounds:** smoothness bounds extend automatically beyond Nash equilibria
- **tightness:** smoothness bounds provably give optimal POA bounds in fundamental cases
- **extensions:** approximate equilibria; best-response dynamics; local smoothness for correlated equilibria; also Bayes-Nash eq

Reasoning About Auctions



Competitive Analysis Fails

Observation: which auction (e.g., opening bid) is best depends on the (unknown) input.

- e.g., opening bid of \$0.01 or \$10 better?

Competitive analysis: compare your revenue to that obtained by an omniscient opponent.

Problem: fails miserably in this context.

- predicts that all auctions are equally terrible
- novel analysis framework needed

A New Analysis Framework

Prior-independent analysis framework: [Hartline/Roughgarden STOC 08, EC 09] compare revenue to that of opponent with *statistical information* about input.

Goal: design a distribution-independent auction that is always near-optimal for the underlying distribution (no matter what the distribution is).

- distribution over inputs not used in the *design* of the auction, only in its *analysis*

Bulow-Klemperer ('96)

Setup: single-item auction. Let F be a *known* valuation distribution. [Needs to be "regular".]

Theorem: [Bulow-Klemperer 96]: for every n :

Vickrey's revenue

OPT's revenue

Bulow-Klemperer ('96)

Setup: single-item auction. Let F be a *known* valuation distribution. [Needs to be "regular".]

Theorem: [Bulow-Klemperer 96]: for every n :

$$\begin{array}{ccc} \text{Vickrey's revenue} & \geq & \text{OPT's revenue} \\ \text{[with } (n+1) \text{ i.i.d. bidders]} & & \text{[with } n \text{ i.i.d. bidders]} \end{array}$$

Bulow-Klemperer ('96)

Setup: single-item auction. Let F be a *known* valuation distribution. [Needs to be "regular".]

Theorem: [Bulow-Klemperer 96]: for every n :

$$\begin{array}{ccc} \text{Vickrey's revenue} & \geq & \text{OPT's revenue} \\ \text{[with } (n+1) \text{ i.i.d. bidders]} & & \text{[with } n \text{ i.i.d. bidders]} \end{array}$$

Interpretation: small increase in competition more important than running optimal auction.

- a "bicriteria bound"!

Bayesian Profit Maximization

Example: 1 bidder, 1 item, $v \sim$ known distribution F

- want to choose optimal reserve price p
- expected revenue of p : $p(1-F(p))$
 - given F , can solve for optimal p^*
 - e.g., $p^* = 1/2$ for $v \sim \text{uniform}[0,1]$
- but: what about $n > 1$ (with i.i.d. v_i 's)?

Bayesian Profit Maximization

Example: 1 bidder, 1 item, $v \sim$ known distribution F

- want to choose optimal reserve price p
- expected revenue of p : $p(1-F(p))$
 - given F , can solve for optimal p^*
 - e.g., $p^* = 1/2$ for $v \sim$ uniform $[0,1]$
- but: what about $n > 1$ (with i.i.d. v_i 's)?

need minor
technical
conditions
on F

Theorem: [Myerson 81] auction with max expected revenue is second-price with above reserve p^* .

- note p^* is *independent of n*

Reformulation of BK Theorem

Theorem: [Bulow-Klemperer 96]: for every n :

$$\begin{array}{ccc} \text{Vickrey's revenue} & \geq & \text{OPT's revenue} \\ \text{[with } (n+1) \text{ i.i.d. bidders]} & & \text{[with } n \text{ i.i.d. bidders]} \end{array}$$

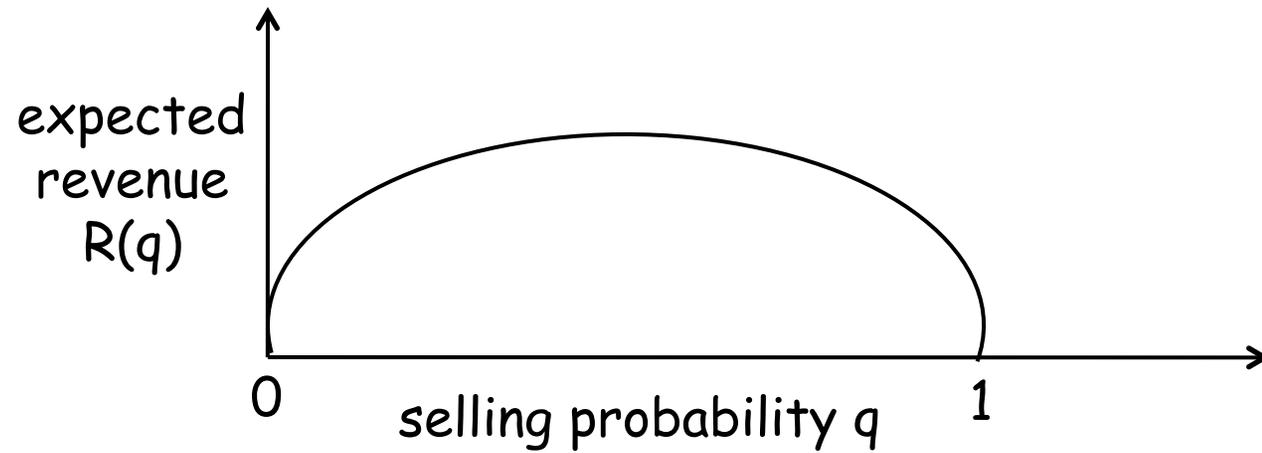
Lemma: if true for $n=1$, then true for all n .

- relevance of OPT reserve price decreases with n

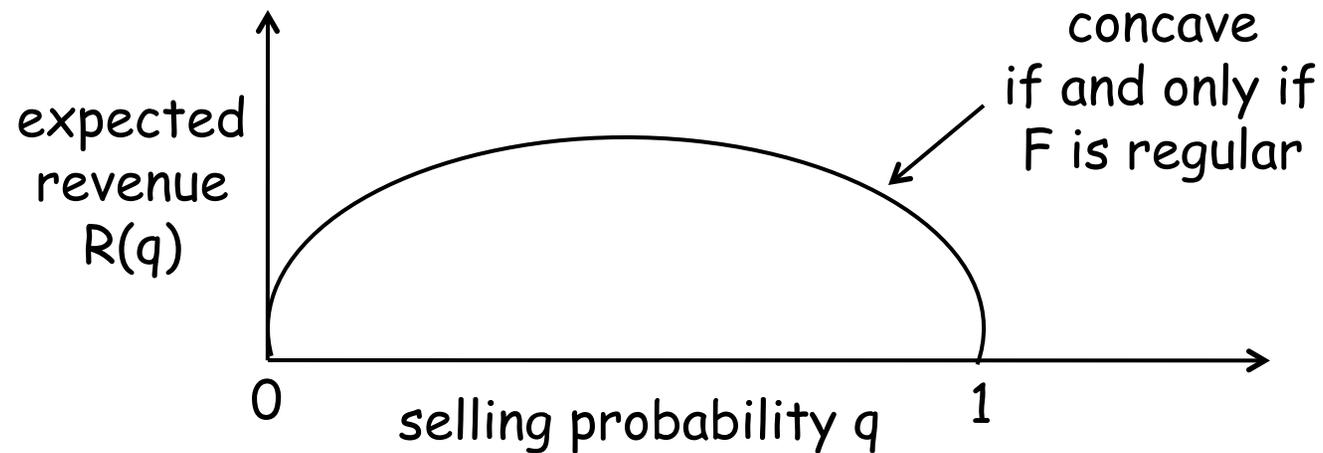
Reformulation for $n=1$ case:

$$\begin{array}{ccc} 2 \times \text{Vickrey's revenue} & & \text{Vickrey's revenue} \\ \text{with } n=1 \text{ and random} & \geq & \text{with } n=1 \text{ and opt} \\ \text{reserve [drawn from } F] & & \text{reserve } r^* \end{array}$$

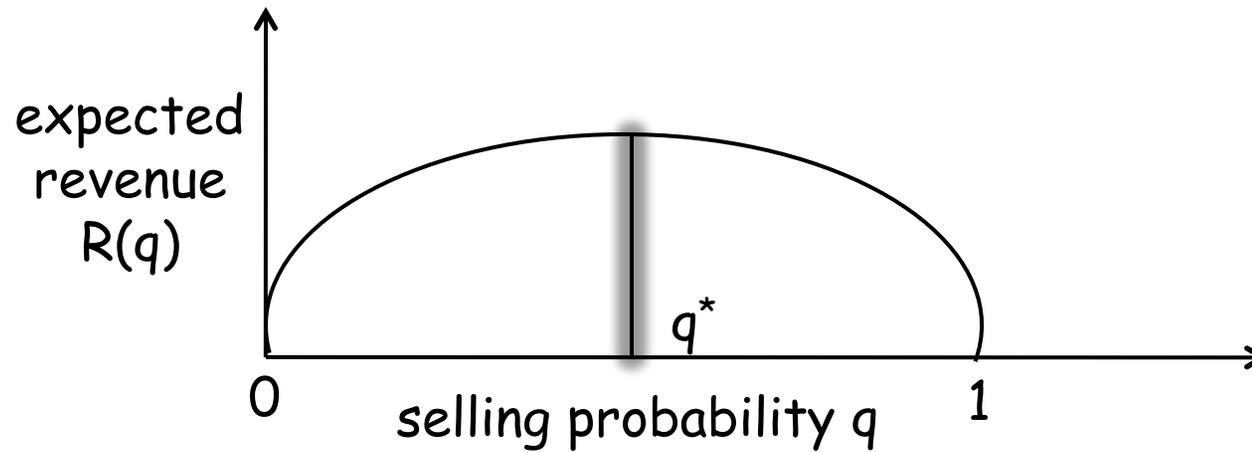
Proof of BK Theorem



Proof of BK Theorem

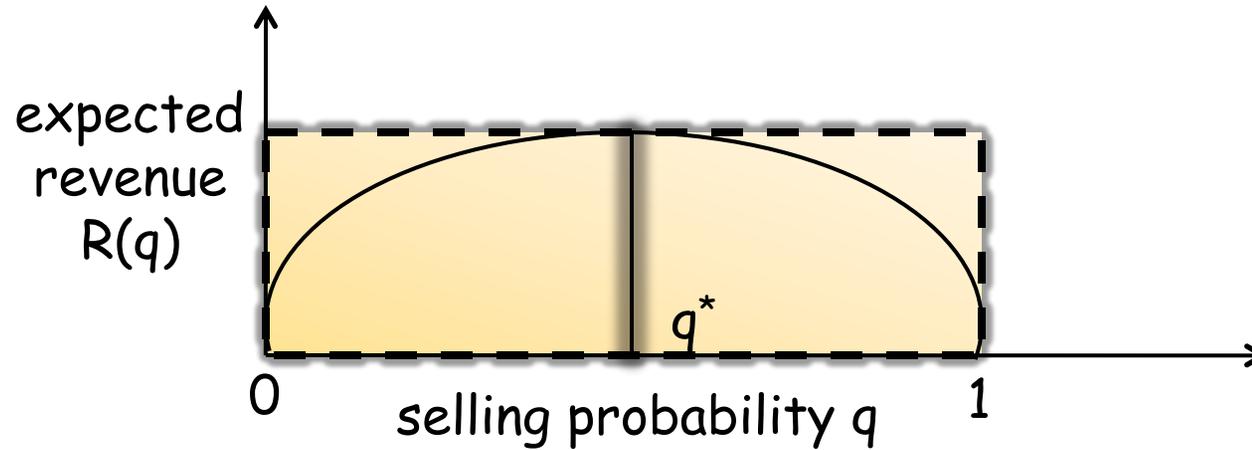


Proof of BK Theorem



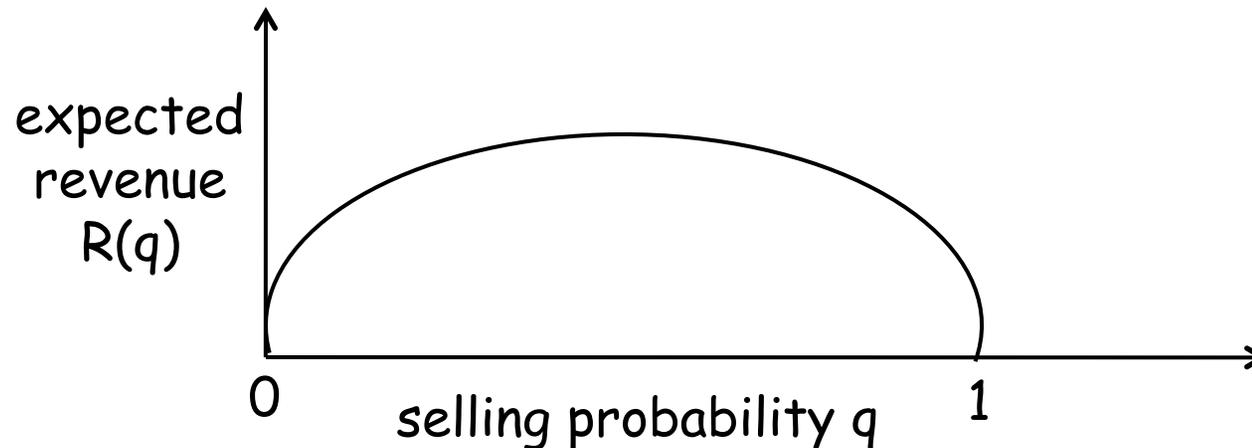
- opt revenue = $R(q^*)$

Proof of BK Theorem



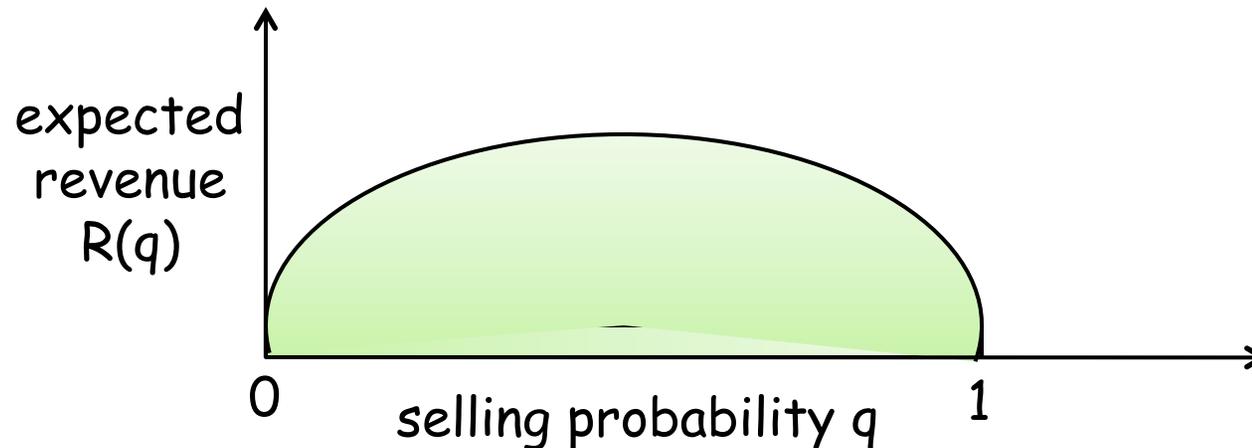
- opt revenue = $R(q^*)$

Proof of BK Theorem



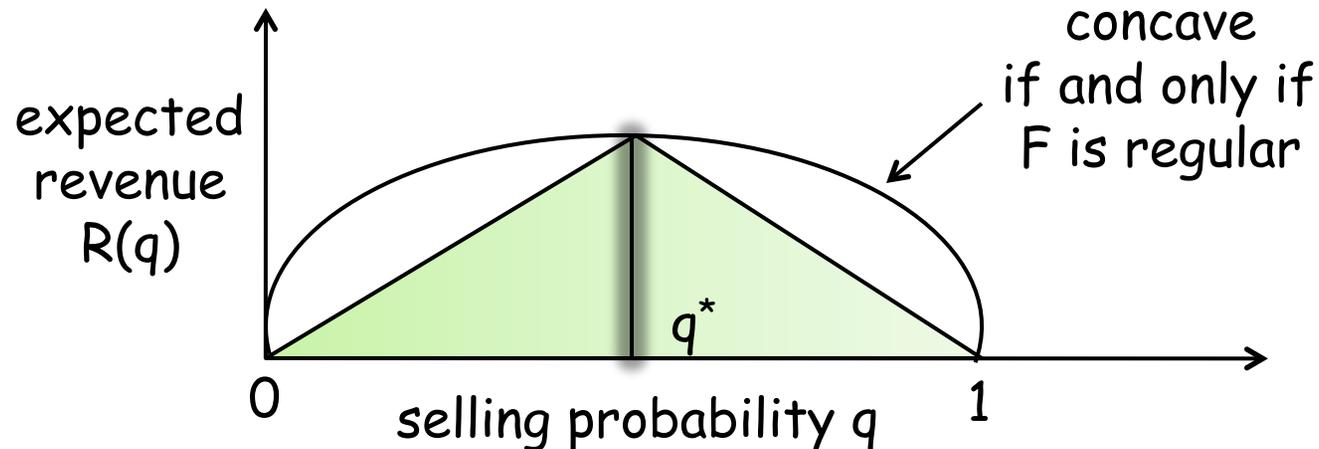
- opt revenue = $R(q^*)$
- revenue of random reserve r (from F) = expected value of $R(q)$ for q uniform in $[0,1]$ = area under revenue curve

Proof of BK Theorem



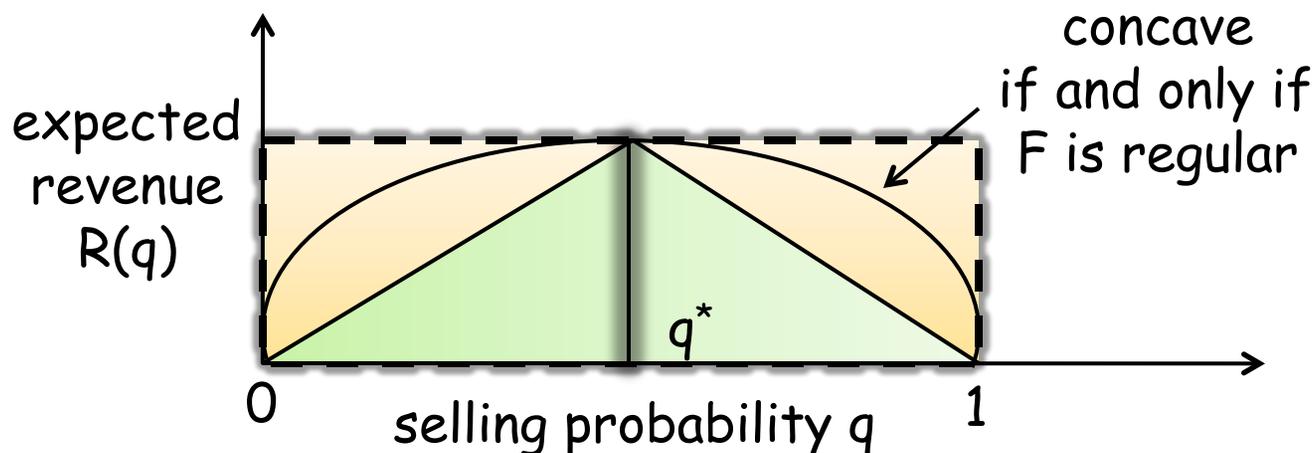
- opt revenue = $R(q^*)$
- revenue of random reserve r (from F) = expected value of $R(q)$ for q uniform in $[0,1]$ = area under revenue curve

Proof of BK Theorem



- opt revenue = $R(q^*)$
- revenue of random reserve r (from F) = expected value of $R(q)$ for q uniform in $[0,1]$ = area under revenue curve

Proof of BK Theorem



- opt revenue = $R(q^*)$
- revenue of random reserve r (from F) = expected value of $R(q)$ for q uniform in $[0, 1]$ = area under revenue curve $\geq \frac{1}{2} \circ R(q^*)$

Recent Progress

BK theorem: the "prior-free" Vickrey auction with extra bidder as good as optimal (w.r.t. F) mechanism, no matter what F is.

More general "bicriteria bounds": [Hartline/Roughgarden EC 09], [Dughmi/Roughgarden/Sundararajan EC 09]

Prior-independent approximations: [Devanur/Hartline EC 09], [Dhangwotnotai/Roughgarden/Yan EC 10], [Hartline/Yan EC 11]

What's Next?

Take-home points:

- standard competitive analysis useless for worst-case revenue maximization
- but can get *simultaneous* competitive guarantee with all Bayesian-optimal auctions

Future Directions:

- thoroughly understand “single-parameter” problems, include non "downward-closed" ones
- non-i.i.d. settings
- multi-parameter? (e.g., combinatorial auctions)

Approximation in AGT

- The Price of Anarchy (etc.)
 - worst-case approximation guarantees for equilibria
- Revenue Maximization
 - guarantees for auctions in non-Bayesian settings (information-theoretic)
- Algorithm Mechanism Design
 - approximation algorithms robust to selfish behavior (computational)
- Computing Approximate Equilibria
 - e.g., is there a PTAS for computing an approximate Nash equilibrium?

} this talk

} FOCs 2010 tutorial

Epilogue

Higher-Level Moral: worst-case approximation guarantees as powerful "intellectual export" to other fields (e.g., game theory).

- many reasons for approximation (not just computational complexity)

Epilogue

Higher-Level Moral: worst-case approximation guarantees as powerful "intellectual export" to other fields (e.g., game theory).

- many reasons for approximation (not just computational complexity)

THANKS!

