Near-Optimal Equilibria

Tim Roughgarden (Stanford)

Simons Institute boot camp on economics and computation

A Representative Result

Example Theorem: [Syrgkanis/Tardos 13] (improving [Hassidim/Kaplan/Nisan/Mansour 11]) Suppose m items are sold simultaneously via first-price single-item auctions:

- for every product distribution over submodular bidder valuations (independent, not necessarily identical), and
- for every (mixed) Bayes-Nash equilibrium,

expected welfare of the equilibrium is within 63% of the maximum possible.

(matches best-possible algorithms!)

Outline

1. *Smooth Games, Extension Theorems, and Robust POA Bounds*

2. Smooth Mechanisms and Bayes-Nash POA Bounds

3. Reducing Complex Mechanisms to Simple Mechanisms Using Composition Theorems

4. Complexity-Based POA Lower Bounds


The Price of Anarchy
Is it okay to be selfish?
FIRE DEPT.
39
FIRE DEPT.

The Price of Anarchy

Network with 2 players:

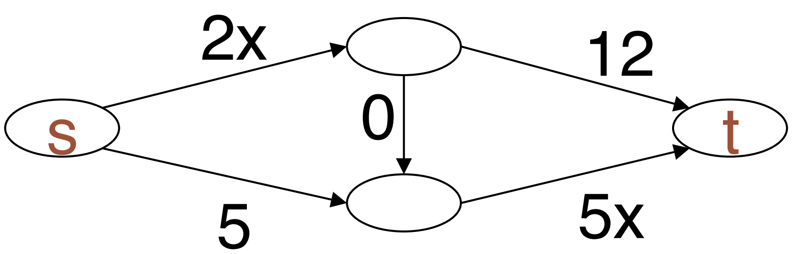

The Price of Anarchy

Nash Equilibrium:

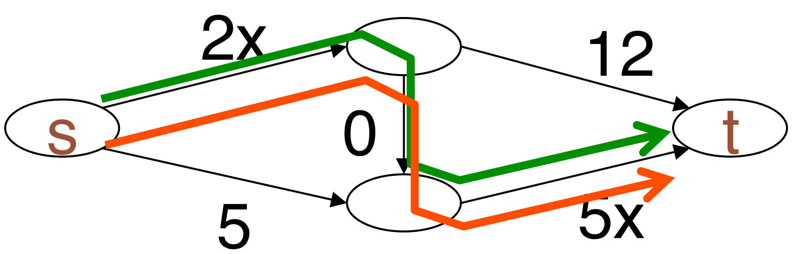

cost = 14+14 = 28

The Price of Anarchy

Nash Equilibrium:

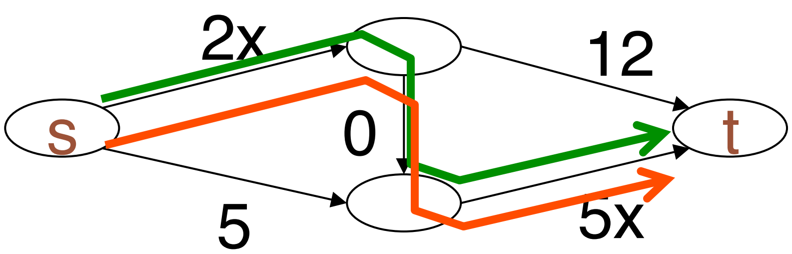

cost = 14+14 = 28

To Minimize Cost:

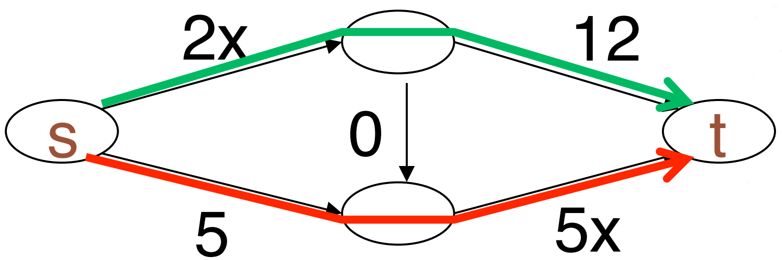

cost = 14+10 = 24

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

- if multiple equilibria exist, look at the *worst* one
  - [Koutsoupias/Papadimitriou 99]

What Do POA Bounds Look Like?

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

Objective function: $cost(\mathbf{s}) := \Sigma_i C_i(\mathbf{s})$

What Do POA Bounds Look Like?

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

Objective function: $\text{cost}(\mathbf{s}) := \Sigma_i C_i(\mathbf{s})$

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

$\text{cost}(\mathbf{s}) \quad = \quad \Sigma_i C_i(\mathbf{s})$ [defn of cost]

What Do POA Bounds Look Like?

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

Objective function: $\text{cost}(\mathbf{s}) := \Sigma_i C_i(\mathbf{s})$

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

$\text{cost}(\mathbf{s}) = \Sigma_i C_i(\mathbf{s})$ [defn of cost]

$\leq \Sigma_i C_i(s^*_i, \mathbf{s}_{-i})$ [**s** a Nash eq]

"baseline" strategies → s^*_i

What Do POA Bounds Look Like?

Suppose: we prove that (for $\lambda > 0$; $\mu < 1$)

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

What Do POA Bounds Look Like?

Suppose: we prove that (for $\lambda > 0$; $\mu < 1$)

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

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

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

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

Canonical Example

Claim [Christodoulou/Koutsoupias 05] (see also [Awerbuch/Azar Epstein 05]) worst-case POA in routing games with affine cost functions is 5/2.

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

Smooth Games

Definition: [Roughgarden 09] A game is *(λ,μ)-smooth w.r.t. baselines $\mathbf{s}^*$* if, for every outcome $\mathbf{s}$ ($\lambda > 0$; $\mu < 1$):

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

Smooth Games

Definition: [Roughgarden 09] A game is *(λ,μ)-smooth w.r.t. baselines* $\mathbf{s}^*$ if, for every outcome **s** (λ > 0; μ < 1):

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

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

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

So: if (λ,μ)-smooth w.r.t. optimal outcome, then POA (of pure Nash equilibria) is at most λ/(1-μ).

(using (*) only in the special case where **s** = equilibrium)

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: fail to converge, POA bound won't apply.

Learnable Equilibria

Fact: simple strategies converge quickly to more permissive equilibrium sets.

- correlated equilibria: [Foster/Vohra 97], [Fudenberg/Levine 99], [Hart/Mas-Colell 00], ...
- coarse/weak correlated equilibria (of [Moulin/Vial 78]): [Hannan 57], [Littlestone/Warmuth 94], ...

Question: are there good "robust" POA bounds, which hold more generally for such "easily learned" equilibria?

[Mirrokni/Vetta 04], [Goemans/Mirrokni/Vetta 05], [Awerbuch/Azar/Epstein/Mirrokni/Skopalik 08], [Christodoulou/Koutsoupias 05], [Blum/Even-Dar/Ligett 06], [Blum/Hajiaghayi/Ligett/Roth 08]

A Hierarchy of Equilibria

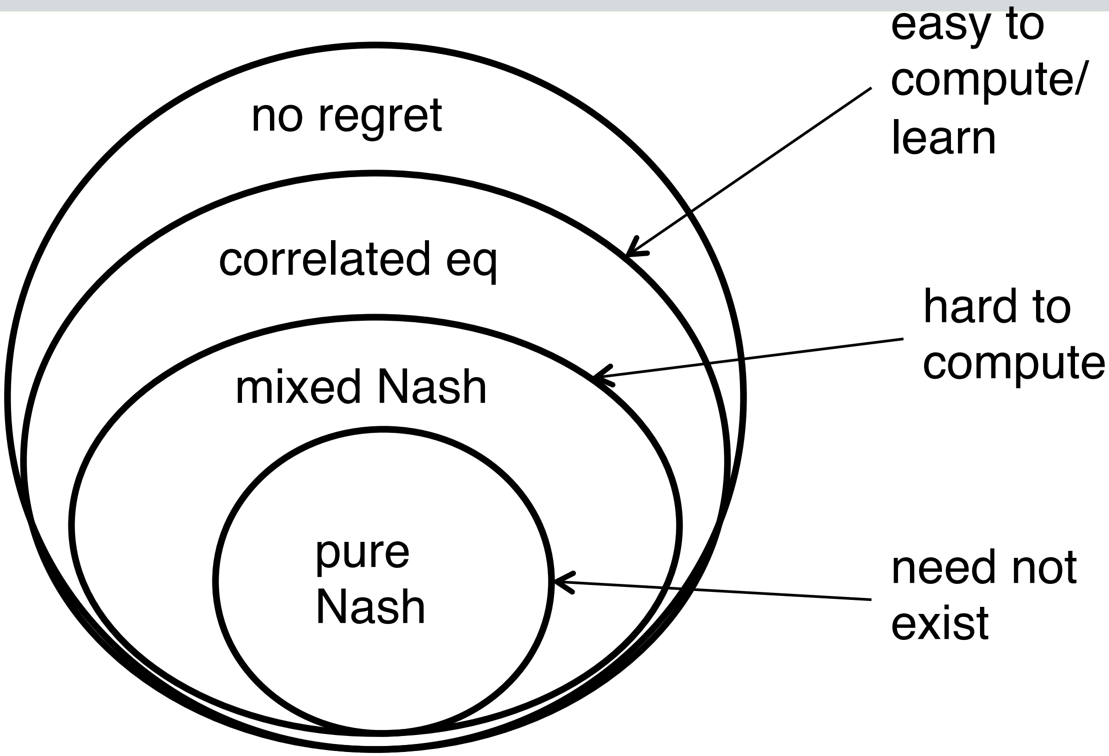

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

An Out-of-Equilibrium Bound

Theorem: [Roughgarden 09] if game is (λ,μ)-smooth w.r.t. an optimal outcome, then the average cost of every no-regret sequence is at most

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

(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:

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

(will ignore the "o(1)" term)

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]

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]

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

Smooth => No-Regret Bound

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

Assuming (λ,μ)-smooth:

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

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

$\leq \Sigma_t [\lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \text{cost}(\mathbf{s}^t)] + \Sigma_i \Sigma_t \Delta_{i,t}$ [smooth]

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]

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

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

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

To finish proof: divide through by T.

Extension Theorems

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

what we care about

Extension Theorems

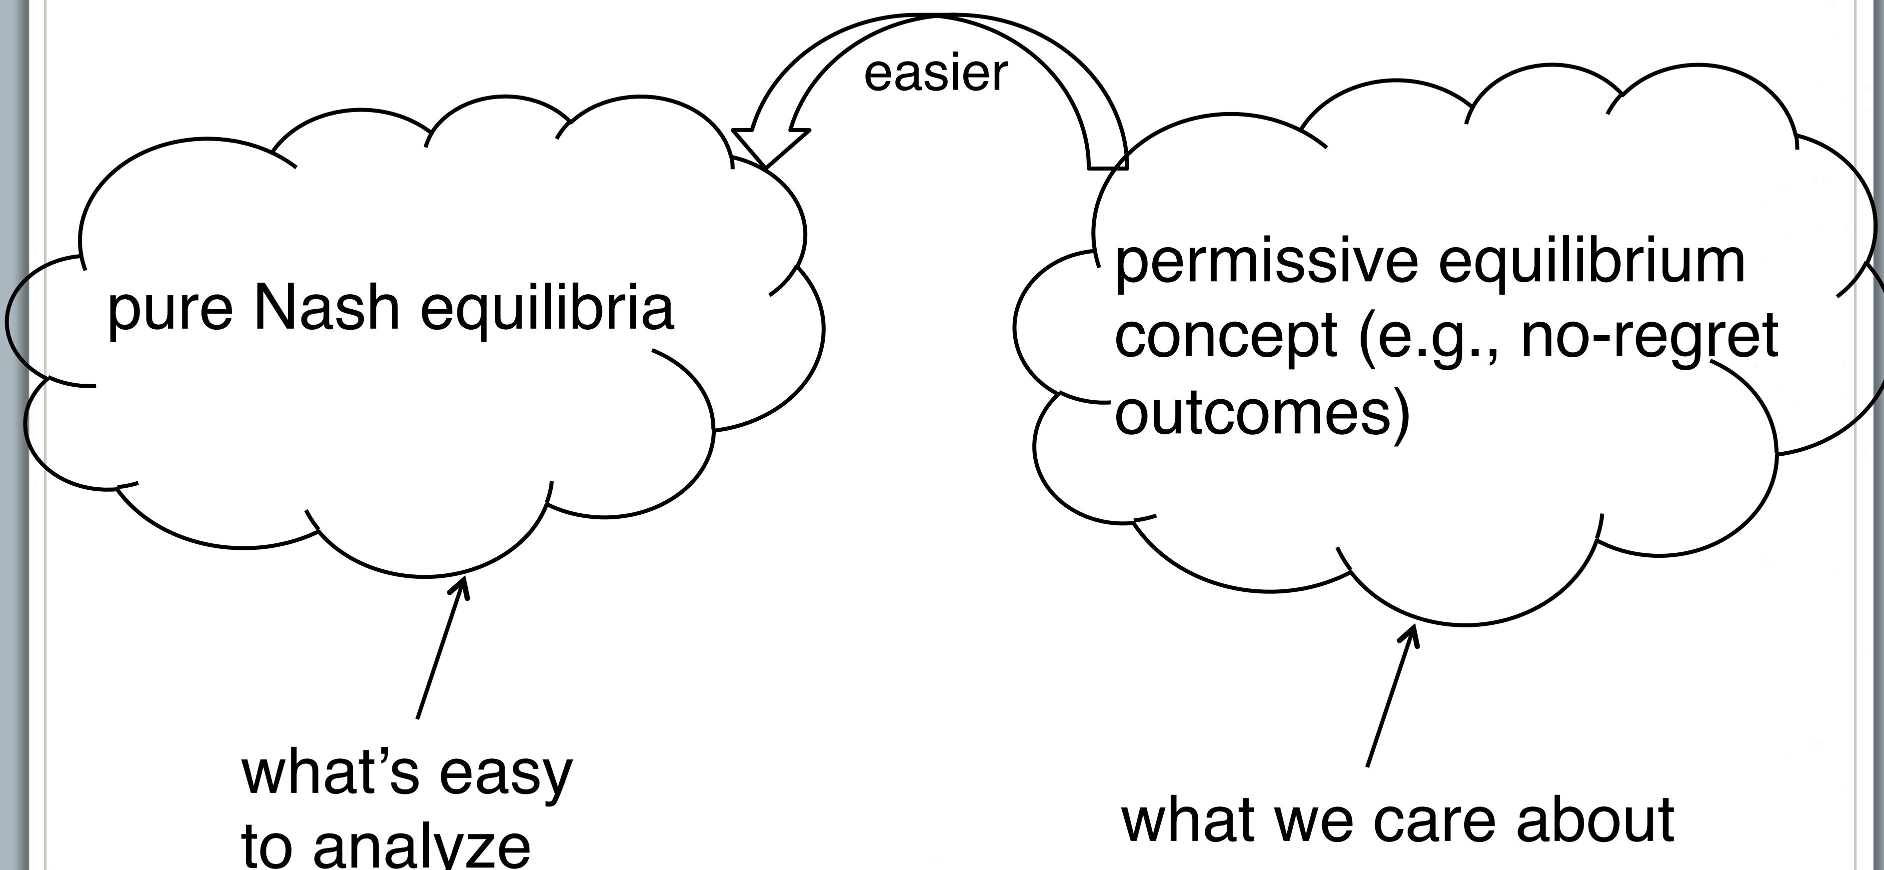

Extension Theorems

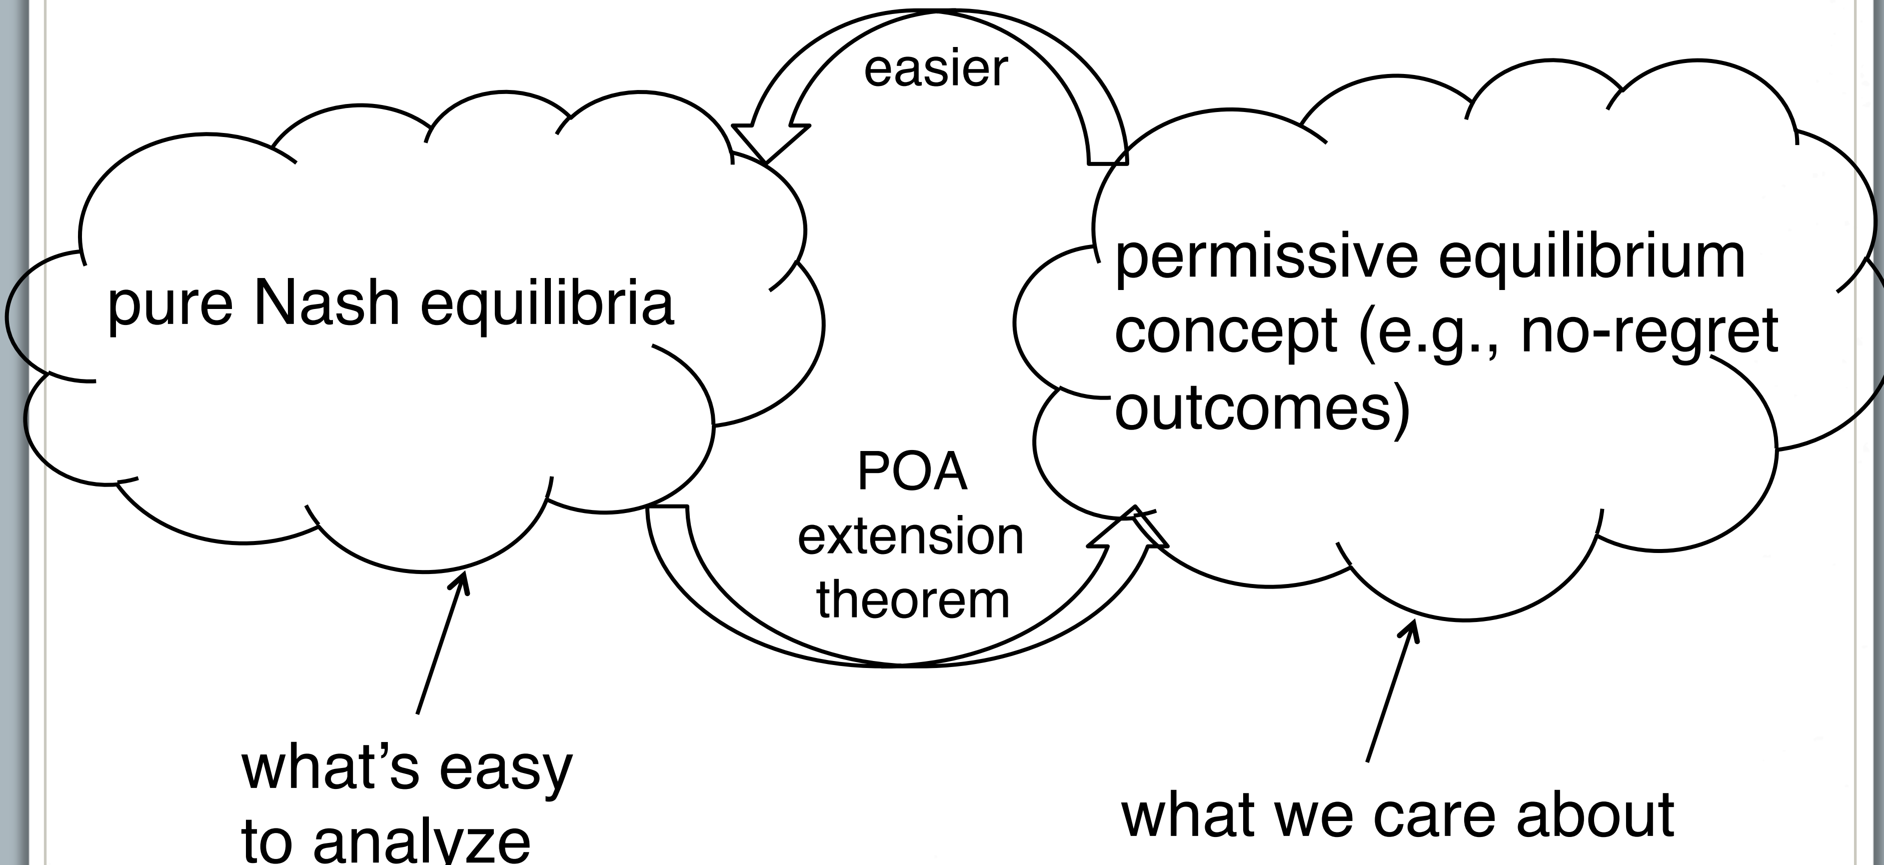

Bells and Whistles

- can allow baseline s^*_i to depend on s_i, but not $\mathbf{s}_{-i}$
- POA bound extends to correlated equilibria
- but *not* to no-regret sequences
- applications include:
  - splittable routing games [Roughgarden/Schoppman 11]
  - opinion formation games [Bhawalkar/Gollapudi/Munagala 13]
  - sequential composition of auctions [Syrgkanis/Tardos 13]

Outline

1. Smooth Games, Extension Theorems, and Robust POA Bounds

2. *Smooth Mechanisms and Bayes-Nash POA Bounds*

3. Reducing Complex Mechanisms to Simple Mechanisms Using Composition Theorems

4. Complexity-Based POA Lower Bounds

Incomplete-Information Games

Game of incomplete information: [Harsanyi 67,68] specified by players, types, actions, payoffs.

- e.g., type = private valuation for a good
- player payoff depends on outcome *and type*
- strategy: function from types to actions
  - semantics: "if my type is t, then I will play action a"

Common Prior Assumption: types drawn from a distribution known to all players (independent, or not)

- realization of type i known only to player i

Example: First-Price Auction

Bayes-Nash Equilibrium: every player picks expected utility-maximizing action, given its knowledge.

Exercise: with n bidders, valuations drawn i.i.d. from U[0,1], the following is a Bayes-Nash equilibrium: all bidders use the strategy $v_i \longrightarrow [(n-1)/n] \bullet v_i$.

- highest-valuation player wins (maximizes welfare)

Example: First-Price Auction

Bayes-Nash Equilibrium: every player picks expected utility-maximizing action, given its knowledge.

Exercise: with n bidders, valuations drawn i.i.d. from U[0,1], the following is a Bayes-Nash equilibrium: all bidders use the strategy $v_i \longrightarrow [(n-1)/n] \bullet v_i$.

- highest-valuation player wins (maximizes welfare)

Exercise: with 2 bidders, valuations from U[0,1] and U[0,2], no Bayes-Nash equilibrium maximizes expected welfare. (Second bidder shades bid more.)

POA with Incomplete Information: The Best-Case Scenario

Ideal: POA bounds w.r.t an *arbitrary* prior distribution.

(or maybe assuming only independence)

Observation: point mass prior distribution ⇔ game of full-information (Bayes-Nash equilibria ⇔ Nash eq).

POA with Incomplete Information: The Best-Case Scenario

Ideal: POA bounds w.r.t an *arbitrary* prior distribution.

(or maybe assuming only independence)

Observation: point mass prior distribution ⇔ game of full-information (Bayes-Nash equilibria ⇔ Nash eq).

Coolest Statement That Could Be True: POA of Bayes-Nash equilibria (for worst-case prior distribution) same as that of Nash equilibria in worst induced full-info game. (Observation above => can only be worse)

Ideal Extension Theorem

Hypothesis: in every induced full-information game, a smoothness-type proof shows that the POA of (pure) Nash equilibria is α or better.

- induced full-info game ⇔ specific type profile
- ex: first-price auction with known valuations

Conclusion: for every common prior distribution, the POA of (mixed) Bayes-Nash equilibria is α or better.

Extension Theorem (Informal)

incomplete-info games

- i.e., uncertain payoffs

mixed Bayes-Nash equilibria

what we care about (e.g., for auctions)

Extension Theorem (Informal)

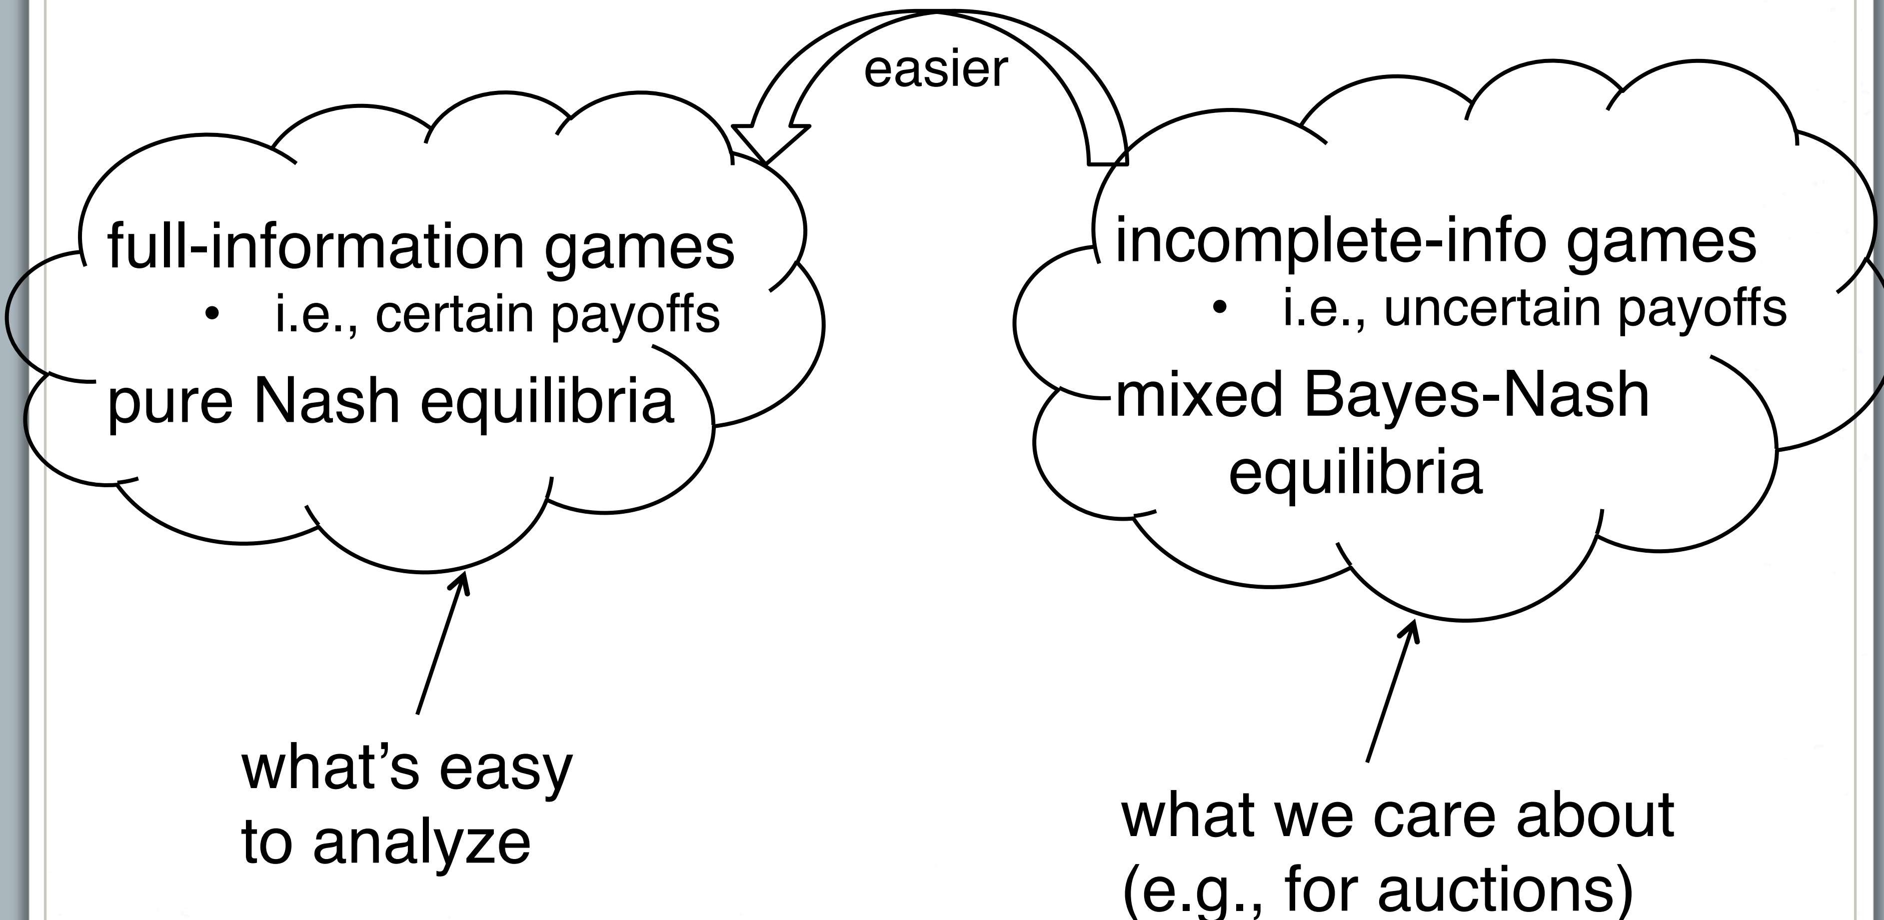

Extension Theorem (Informal)

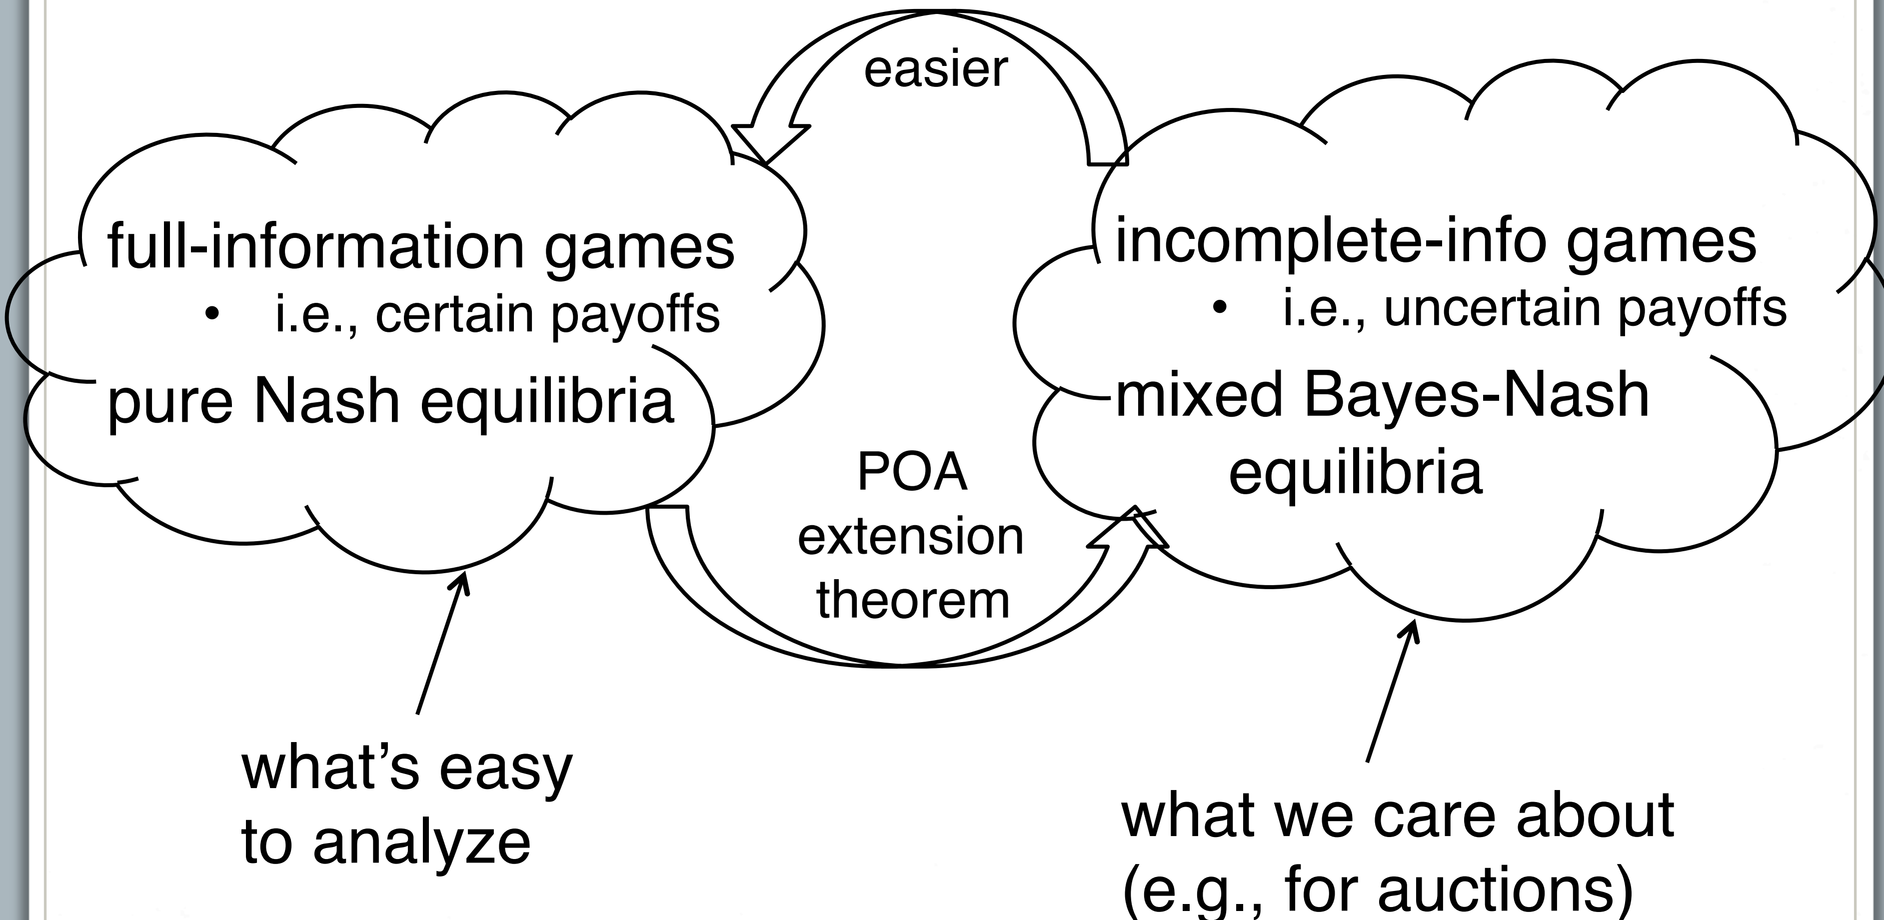

Smoothness Paradigm (Full Information)

1. Fix a game.

   (fixes optimal outcomes)

2. Choose baseline $\mathbf{s}^*$ = some optimal outcome.

   (in many games, only one option)

3. Fix outcome $\mathbf{s}$.

4. Prove $\Sigma_i C_i(s^*_i, \mathbf{s}_{-i}) \leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s})$.

5. Conclude that POA of no-regret sequences $\leq \lambda/(1-\mu)$.

Smoothness Paradigm (Full => Incomplete)

1. Fix a setting *and the private valuations.*
   (fixes optimal outcomes)

2. Choose baseline $\mathbf{s}^*$ = some optimal outcome.
   (in many games, only one option)

3. Fix outcome $\mathbf{s}$.

4. Prove $\Sigma_i C_i(s^*_i, \mathbf{s}_{-i}) \leq \lambda \bullet cost(\mathbf{s}^*) + \mu \bullet cost(\mathbf{s})$.

5. Conclude that POA of no-regret sequences $\leq \lambda/(1-\mu)$.

Smoothness Paradigm (Full => Incomplete)

1. Fix a setting *and the private valuations.*

   (fixes optimal outcomes)

2. Choose baseline $\mathbf{b}^*$ = some optimal outcome.

   (note the large number of possible options)

3. Fix outcome $\mathbf{s}$.

4. Prove $\Sigma_i C_i(s^*_i, \mathbf{s}_{-i}) \leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s})$.

5. Conclude that POA of no-regret sequences $\leq \lambda/(1-\mu)$.

Smoothness Paradigm (Full => Incomplete)

1. Fix a setting *and the private valuations.*

   (fixes optimal outcomes)

2. Choose baseline $\mathbf{b}^*$ = some optimal outcome.

   (note the large number of possible options)

3. Fix outcome **b**.

4. Prove $\Sigma_i C_i(s^*_i, \mathbf{s}_{-i}) \leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s})$.

5. Conclude that POA of no-regret sequences ≤ $\lambda/(1-\mu)$.

Smoothness Paradigm (Full => Incomplete)

1. Fix a setting *and the private valuations.*

   (fixes optimal outcomes)

2. Choose baseline $\mathbf{b}^*$ = some optimal outcome.

   (note the large number of possible options)

3. Fix outcome **b**.

4. Prove $\Sigma_i\, u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b})$.

   [Syrgkanis/ Tardos 13]

5. Conclude that POA of no-regret sequences ≤ λ/(1-μ).

Smoothness Paradigm (Incomplete Information)

1. Fix a setting *and the private valuations.*

   (fixes optimal outcomes)

2. Choose baseline $\mathbf{b}^*$ = some optimal outcome.

   (note the large number of possible options)

3. Fix outcome **b**.

4. Prove $\Sigma_i u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet$[OPT Welfare] – Revenue(**b**).

5. Conclude that POA of Bayes-Nash equilibria is ≥ λ.

[Syrgkanis/Tardos 13]

Smoothness Paradigm (Incomplete Information)

1. Fix a setting *and the private valuations.*

   (fixes optimal outcomes)

2. *Choose baseline* $\mathbf{b}^*$ *= some optimal outcome.*

   (note the large number of possible options)

3. Fix outcome **b**.

4. *Prove* $\Sigma_i u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \cdot$[OPT Welfare] – Revenue($\mathbf{b}$).

5. Conclude that POA of Bayes-Nash equilibria is ≥ λ.

first-price auctions: for suitable $\mathbf{b}^*$, $\lambda \geq ½$

First-Price Auctions

Claim: for suitable choice of $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i\, u_i(b^*_i,\mathbf{b}_{-i}) \;\geq\; ½•[\text{OPT Welfare}] – \text{Revenue}(\mathbf{b}).$$

Proof: Set $b^*_i = v_i/2$ for every i. (a la [Lucier/Paes Leme 11])

- since LHS ≥ 0, can assume $½•[\max_i v_i] > \max_i b_i$
- suppose bidder 1 has highest valuation. Then:

$$u_1(b^*_1,b_{-1}) = v_1 – (v_1/2) = v_1/2 \;\geq\; ½•[\text{OPT Welfare}]$$

Optimization: [Syrgkanis 12] 50% => 63% (different $\mathbf{b}^*$)

Smoothness Paradigm (Incomplete Information)

1. Fix a setting *and the private valuations.*
   (fixes optimal outcomes)

2. Choose baseline $\mathbf{b}^*$ = some optimal outcome.
   (note the large number of possible options)

3. Fix outcome **b**.

4. Prove $\Sigma_i\, u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet$[OPT Welfare] – Revenue(**b**).

5. *Conclude that POA of Bayes-Nash equilibria is ≥ λ.*

general extension theorem

Extension Theorem (PNE)

Assume: for suitable choice of $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i\, u_i(b^*_i,\mathbf{b}_{-i}) \;\geq\; \lambda\bullet[\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: POA of pure Nash equilibria is $\geq \lambda$.

Extension Theorem (PNE)

Assume: for suitable choice of $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i u_i(b^*_i,\mathbf{b}_{-i}) \geq \lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: POA of pure Nash equilibria is $\geq \lambda$.

Proof: Let $\mathbf{b}$ = a pure Nash equilibrium. Then:

$$\begin{aligned} \text{welfare}(\mathbf{b}) &= \text{Rev}(\mathbf{b}) + \Sigma_i u_i(\mathbf{b}) && \text{[defn of utility]} \\ &\geq \text{Rev}(\mathbf{b}) + \Sigma_i u_i(b^*_i,\mathbf{b}_{-i}) && [\mathbf{b} \text{ a Nash eq}] \\ &\geq \text{Rev}(\mathbf{b}) + [\lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b})] \\ &= \lambda \bullet [\text{OPT Welfare}] \end{aligned}$$

Extension Theorem (BNE)

Assume: for suitable choice of $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i\, u_i(b^*_i,\mathbf{b}_{-i}) \;\geq\; \lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: (≈[Lucier/Paes Leme 11]) for all (possibly correlated) valuation distributions, POA of Bayes-Nash eq is $\geq \lambda$.

Proof: Let $\mathbf{b}()$ = a Bayes-Nash equilibrium. Then:

$$
\begin{aligned}
E_v[\text{welfare}(\mathbf{b}(\mathbf{v}))] &= E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + \Sigma_i\, E_v[u_i(\mathbf{b}(\mathbf{v}))] && \text{[defn of utility]}\\
&\geq E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + \Sigma_i\, E_v[u_i(b^*_i(v_i),\mathbf{b}_{-i}(\mathbf{v}_{-i}))] && [\mathbf{b} \text{ a BNE}]\\
&\geq E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + [\lambda \bullet E_v[\text{OPT Welfare}] - E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))]]\\
&= \lambda \bullet E_v[\text{OPT Welfare}]
\end{aligned}
$$

First-Price Auctions

Summary: for all (possibly correlated) valuation distributions, every Bayes-Nash equilibrium of a first-price auction has welfare at least 50% (or even 63%) of the maximum possible.

- 63% is tight for correlated valuations [Syrgkanis 14]
- independent valuations = worst-case POA unknown
  - worst known example = 87% [Hartline/Hoy/Taggart 14]
- 63% extends to simultaneous single-item auctions (covered tomorrow)

Further Applications

- first-price sponsored search auctions [Caragiannis/Kaklamanis/Kanellopolous/Kyropoulou/Lucier/Paes Leme/Tardos 12]

- greedy pay-as-bid combinatorial auctions [Lucier/Borodin 10]

- pay-as-bid mechanisms based on LP rounding [Duetting/Kesselheim/Tardos 15]

Second-Price Rules

- simultaneous second-price auctions [Christodoulou/Kovacs/Schapira 08]
  - worst-case POA = 50%, and this is tight (even for PNE)
- truthful greedy combinatorial auctions [Borodin/Lucier 10]
  - worst-case POA close to greedy approximation ratio
- can be reinterpreted via modified smoothness condition [Roughgarden 12, Syrgkanis 12]
- "bluffing equilibria" => need a no overbidding condition for non-trivial POA bounds

Revenue Covering

- [Hartline/Hoy/Taggart 14] define "revenue covering"
- for every **b**, Rev(**b**) ≥ critical bids of winners in OPT
- implies smoothness condition
  - near-equivalent in some cases [Duetting/Kesselheim 15]
- application #1: POA bounds w.r.t. revenue objective
  - e.g., simultaneous first-price auctions with monopoly reserves
- application #2: [Hoy/Nekipelov/Syrgkanis 15] bound the "empirical POA" from data
  - do not need to explicitly estimate valuations!
  - can prove instance-by-instance bounds that beat the worst-case bound

Dynamic Auctions

[Lykouris/Syrgkanis/Tardos 15] first POA guarantees when bidder population changing (p fraction drops out each time step, replaced by new bidders).

- convergence to (Nash) equilibrium hopeless
- positive results for "adaptive learners" (assume agents use sufficiently good learning algorithm)
- need baseline near-optimal strategy profiles (one per time step) s.t. no player changes frequently
- novel use of differential privacy! (in the analysis)

Outline

1. Smooth Games, Extension Theorems, and Robust POA Bounds

2. Smooth Mechanisms and Bayes-Nash POA Bounds

3. *Reducing Complex Mechanisms to Simple Mechanisms Using Composition Theorems*

4. Complexity-Based POA Lower Bounds

Multi-Item Auctions

- suppose m different items
- for now: *unit-demand* valuations
- each bidder i has private valuation v_{ij} for each item j
- $v_i(S) := \max_{j \text{ in } S} v_{ij}$

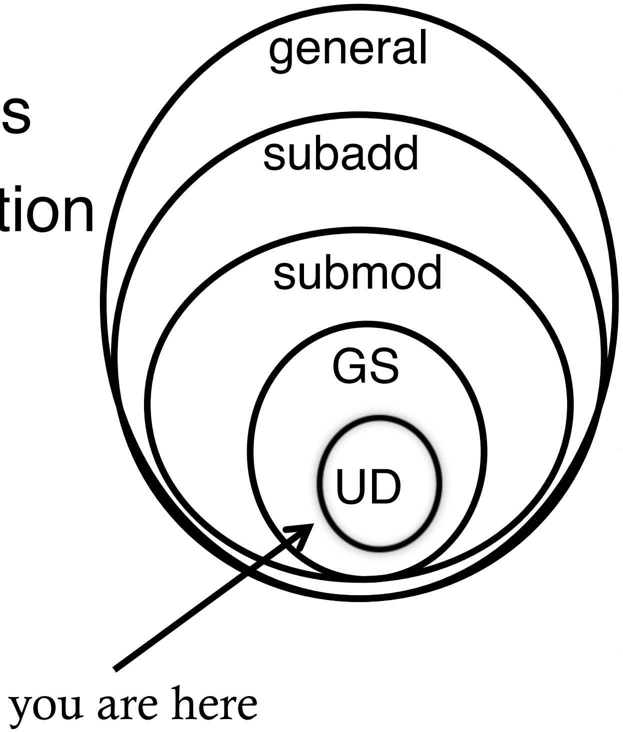

Simultaneous Composition

- suppose have mechanisms $M_1,\ldots,M_m$
- in their *simultaneous composition*:
  - new action space = product of the m action spaces
  - new allocation rule = union of the m allocation rules
  - new payment rule = sum of the m payment rules
- example: each M_j a single-item first-price auction

Question: as a unit-demand bidder, how should you bid?

(not so easy)

Composition Preserves Smoothness

Hypothesis: every single-item auction M_j is λ-smooth: for every **v**, there exists $\mathbf{b}^*$ such that, for every **b**,

$$\Sigma_i\, u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet [\text{OPT Welfare}(\mathbf{v})] - \text{Rev}(\mathbf{b}).$$

Theorem: [Syrgkanis/Tardos 13] if bidders are unit-demand, then composed mechanism is also λ-smooth.

- holds more generally from arbitrary smooth M_j's and "XOS" valuations (generalization of submodular)

Composition Preserves Smoothness

Hypothesis: every single-item auction M_j is λ-smooth: for every **v**, there exists **b*** such that, for every **b,**

$$\Sigma_i\, u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet [\text{OPT Welfare}(\mathbf{v})] - \text{Rev}(\mathbf{b}).$$

Theorem: [Syrgkanis/Tardos 13] if bidders are unit-demand, then composed mechanism is also λ-smooth.

Proof idea: Fix unit-demand valuations **v**, fixes OPT.

- baseline strategy for a bidder i that gets item j in OPT
  - bid 0 in mechanisms other M_j
  - in M_j, use assumed baseline strategy for M_j

Simultaneous First-Price Auctions (First Try)

Consequence: for all (possibly correlated) unit-demand valuation distributions, every Bayes-Nash equilibrium of simultaneous first-price auctions has welfare at least 50% (or even 63%) of the maximum possible.

- prove smoothness inequality for first-price auction
- use composition theorem to extend smoothness to simultaneous first-price auctions
- use extension theorem to conclude Bayes-Nash POA bound for simultaneous first-price auctions

Counterexample

Fact: [Feldman/Fu/Gravin/Lucier 13], following [Bhawalkar/Roughgarden 11] there are (highly correlated) valuation distributions over unit-demand valuations such that every Bayes-Nash equilibrium has expected welfare arbitrary smaller than the maximum possible.

- idea: plant a random matching plus some additional highly demanded items; by symmetry, a bidder can't detect the item "reserved" for it

Revised Statement

Consequence: for all *product* unit-demand valuation distributions, every Bayes-Nash equilibrium of simultneous first-price auction has welfare at least 50% (or even 63%) of the maximum possible.

- prove smoothness inequality for first-price auction
- use composition theorem to extend smoothness to simultaneous first-price auctions
- use *modified* extension theorem to conclude Bayes-Nash POA bound for simultaneous first-price auctions

Private Baseline Strategies

First-price auction: set $b^*_i = v_i/2$ for every i.

- independent of $\mathbf{v}_{-i}$ ("private" baseline strategies)

Simultaneous first-price auctions: b^*_i is "bid half your value only on the item j you get in OPT($\mathbf{v}$)".

- "public" baseline strategies
- not well defined unless $\mathbf{v}_{-i}$ known

Extension Theorem (BNE)

Assume: for suitable choice of $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i u_i(b^*_i, \mathbf{b}_{-i}) \geq \lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: (≈[Lucier/Paes Leme 11]) for all (possibly correlated) valuation distributions, POA of Bayes-Nash eq is ≥ λ.

Proof: Let **b()** = a Bayes-Nash equilibrium. Then:

$$E_v[\text{welfare}(\mathbf{b}(\mathbf{v}))] = E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + \Sigma_i E_v[u_i(\mathbf{b}(\mathbf{v}))] \quad \text{[defn of utility]}$$

$$\geq E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + \Sigma_i E_v[u_i(b^*_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \quad [\mathbf{b} \text{ a BNE}]$$

$$\geq E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))] + [\lambda \bullet E_v[\text{OPT Welfare}] - E_v[\text{Rev}(\mathbf{b}(\mathbf{v}))]]$$

$$= \lambda \bullet E_v[\text{OPT Welfare}]$$

deviation can depend on v_i but not $\mathbf{v}_{-i}$

Extension Theorem (BNE)

Assume: for suitable choice of *private* $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i\, u_i(b^*_i, \mathbf{b}_{-i}) \;\geq\; \lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: (≈[Lucier/Paes Leme 11]) for all (possibly correlated) valuation distributions, POA of Bayes-Nash eq is ≥ λ.

Proof: Let **b()** = a Bayes-Nash equilibrium. Then:

$$E_v[\text{welfare}(\mathbf{b(v)})] = E_v[\text{Rev}(\mathbf{b(v)})] + \Sigma_i\, E_v[u_i(\mathbf{b(v)})] \quad \text{[defn of utility]}$$

$$\geq E_v[\text{Rev}(\mathbf{b(v)})] + \Sigma_i\, E_v[u_i(b^*_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}))] \quad [\mathbf{b} \text{ a BNE}]$$

$$\geq E_v[\text{Rev}(\mathbf{b(v)})] + [\lambda \bullet E_v[\text{OPT Welfare}] - E_v[\text{Rev}(\mathbf{b(v)})]]$$

$$= \lambda \bullet E_v[\text{OPT Welfare}]$$

deviation can depend on v_i but not $\mathbf{v}_{-i}$

Modified Extension Theorem

Assume: for suitable choice of *public* $\mathbf{b}^*$, for every $\mathbf{b}$,

$$\Sigma_i\, u_i(b^*_i, \mathbf{b_{-i}}) \geq \lambda \bullet [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Theorem: [Syrgkanis/Tardos 13], following [Christodoulou/Kovacs/Schapira 08] for all *product* valuation distributions, POA of Bayes-Nash eq is $\geq \lambda$.

Proof idea: to transform public b^*_i to a deviation:

- sample $\mathbf{w_{-i}}$ from prior distribution
- play baseline strategy for valuation profile $(v_i, \mathbf{w_{-i}})$

Outline

1. Smooth Games, Extension Theorems, and Robust POA Bounds
2. Smooth Mechanisms and Bayes-Nash POA Bounds
3. Reducing Complex Mechanisms to Simple Mechanisms Using Composition Theorems
4. *Complexity-Based POA Lower Bounds*

Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14]

the worst-case POA of S1A's with subadditive bidder valuations is precisely 2.

monotone *subadditive* valuations:

- $v_i(A \cup B) \le v_i(A) + v_i(B)$ for all disjoint A,B

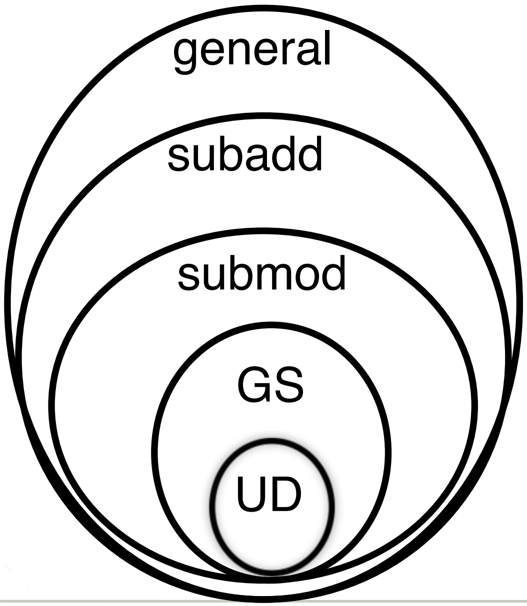

Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14]

the worst-case POA of S1A's with subadditive bidder valuations is precisely 2.

Explicit Lower Bound:

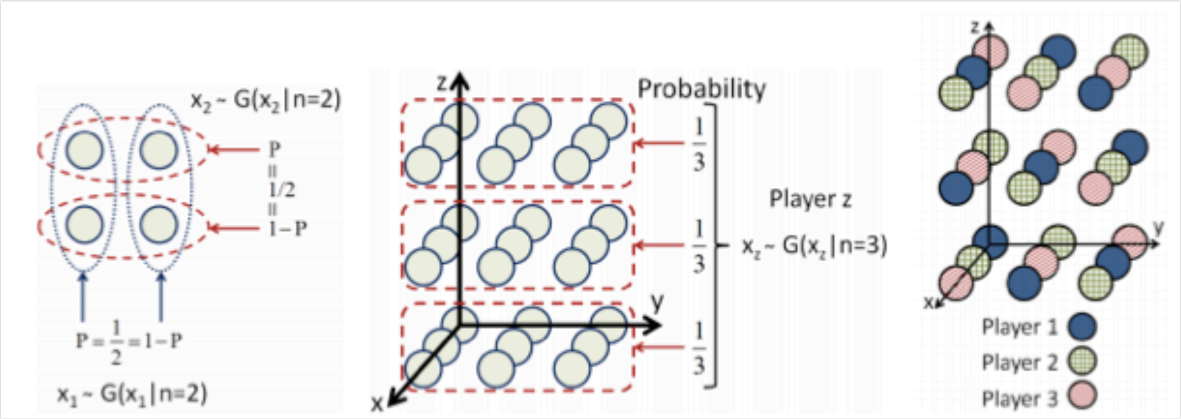

Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14]

the worst-case POA of S1A's with subadditive bidder valuations is precisely 2.

Question: Can we do better?

(without resorting to the VCG mechanism)

The Upshot

Meta-theorem: equilibria are generally bound by the same limitations as algorithms with polynomial computation or communication.

- lower bounds without explicit constructions!

Caveats: requires that equilibria are

- guaranteed to exist (e.g., mixed Nash equilibria)
- can be efficiently verified

Example consequence: no "simple" auction has POA < 2 for bidders with subadditive valuations.

From Protocol Lower Bounds to POA Lower Bounds

Theorem: [Roughgarden 14] Suppose:

- no nondeterministic subexponential-communication protocol approximates the welfare-maximization problem (with valuations V) to within factor of α.
  - i.e., impossible to decide $OPT \geq W^*$ vs. $OPT \leq W^* / \alpha$

*Then worst-case POA of **ε-approximate** mixed Nash equilibria of every "simple" mechanism is at least α.*

- simple = number of strategies sub-doubly-exponential in m
- ε can be as small as inverse polynomial in n and m

Point: : reduces lower bounds for equilibria to lower bounds for communication protocols.

Consequences

Corollary: (via [Nisan/Segal 06], [Dobsinski/Nisan/Schapira 05])

- With subadditive bidder valuations, no simple auction guarantees equilibrium welfare better than 50% OPT.
  - "simple": bid space dimension ≤ subexponential in # of goods
- With general valuations, no simple auction guarantees non-trivial equilibrium welfare.

Take-aways:

1. In these cases, S1A's optimal among simple auctions.
2. With complements, complex bid spaces (e.g., package bidding) necessary for welfare guarantees.

Why Approximate MNE?

Issue: in an S1A, number of strategies = $(V_{max} + 1)^m$

- valuations, bids assumed integral and poly-bounded

Consequence: can't efficiently guess/verify a MNE.

Theorem: [Lipton/Markakis/Mehta 03] a game with n players and N strategies per player has an ε-approximate mixed Nash equilibrium with support size polynomial in n, log N, and ϵ^{-1}.

- proof idea based on sampling from an exact MNE

Nondeterministic Protocols

- each of n players has a private valuation v_i
- a "referee" wants to convince the players that the value of some function $f(v_1,...,v_n)$ has the value z
- referees knows all v_i's and writes, in public view, an alleged proof P that $f(v_1,...,v_n) = z$
- protocol accepts if and only if every player i accepts the proof P (knowing only v_i)
- communication used = length (in bits) of proof P
- example: Non-Equality vs. Equality

From Protocol Lower Bounds to POA Lower Bounds

Theorem: [Roughgarden 14] Suppose:

- no nondeterministic subexponential-communication protocol approximates the welfare-maximization problem (with valuations V) to within factor of α.
  - i.e., impossible to decide OPT $\geq W^*$ vs. OPT $\leq W^*/\alpha$

*Then worst-case POA of **ε-approximate** mixed Nash equilibria of every "simple" mechanism is at least α.*

- simple = number of strategies sub-doubly-exponential in m
- ε can be as small as inverse polynomial in n and m

Point: : reduces lower bounds for equilibria to lower bounds for communication protocols.

Proof of Theorem

Suppose worst-case POA of ε-MNE is $\rho<\alpha$:

Input: game
G s.t. either
(i) $OPT \geq W^*$
or (ii) $OPT \leq W^*/\alpha$

Proof of Theorem

Suppose worst-case POA of ε-MNE is ρ<α:

Input: game G s.t. either (i) OPT ≥ W* or (ii) OPT ≤ W*/ α →

Protocol:
"proof" = ε -MNE x with small support (exists by LMM); players verify it privately

Proof of Theorem

Suppose worst-case POA of ε-MNE is $\rho<\alpha$:

Input: game G s.t. either (i) OPT ≥ W* or (ii) OPT ≤ W*/ α

→

Protocol: "proof" = ε -MNE x with small support (exists by LMM); players verify it privately

→

if E[wel(x)] > W*/ α then OPT > W*/ α so in case (i)

Proof of Theorem

Suppose worst-case POA of ε-MNE is $\rho < \alpha$:

Input: game G s.t. either (i) OPT ≥ W* or (ii) OPT ≤ W^*/α

→

Protocol: "proof" = ε-MNE x with small support (exists by LMM); players verify it privately

→ if $E[wel(x)] > W^*/\alpha$ then $OPT > W^*/\alpha$ so in case (i)

→ if $E[wel(x)] \leq W^*/\alpha$ then $OPT \leq (\rho/\alpha)W^* < W^*$ so in case (ii)

Key point: every ε-MNE is a short, efficiently verifiable certificate for membership in case (ii).

Exact vs. Approximate Equilibria

Claim: POA lower bounds for ε-MNE with small enough ε essentially as good as for exact MNE. Reasons:

1. All known upper bound techniques apply automatically to approximate equilibria.
   1. e.g., "smoothness proofs" [Roughgarden 09]
   2. so our lower bounds limit all known proof techniques
2. Lower bounds for approximate equilibria can sometimes be translated into bounds for exact equilibria.
3. If POA of exact equilibria << POA of approximate equilibria, the latter is likely more relevant (and robust).

More Applications

- optimality results for "simple" auctions with other valuation classes (general, XOS)
- analogous results for combinatorial auctions with succinct valuations (if coNP not in MA)
- impossibility results for low-dimensional price equilibria (assuming NP ≠ coNP) [Roughgarden/Talgam-Cohen 15]
- unlikely to reduce planted clique to ε-Nash hardness

Open Questions

1. Tight POA bounds for important auction formats
   1. e.g. first-price auctions with independent valuations

2. Best "simple" auction for submodular valuations?
   1. S1A's give 63% [Syrgkanis/Tardos 13], [Christodoulou et al 14]
   2. > 77% impossible [Dobzinski/Vondrak 13] + [R14]
   3. > 63% is possible with poly communication [Feige/Vondrak 06]

3. Design "natural" games with POA matching hardness lower bound for the underlying optimization problem.
   1. e.g., many auction and scheduling problems

FIN