Intrinsic Robustness of the Price of Anarchy

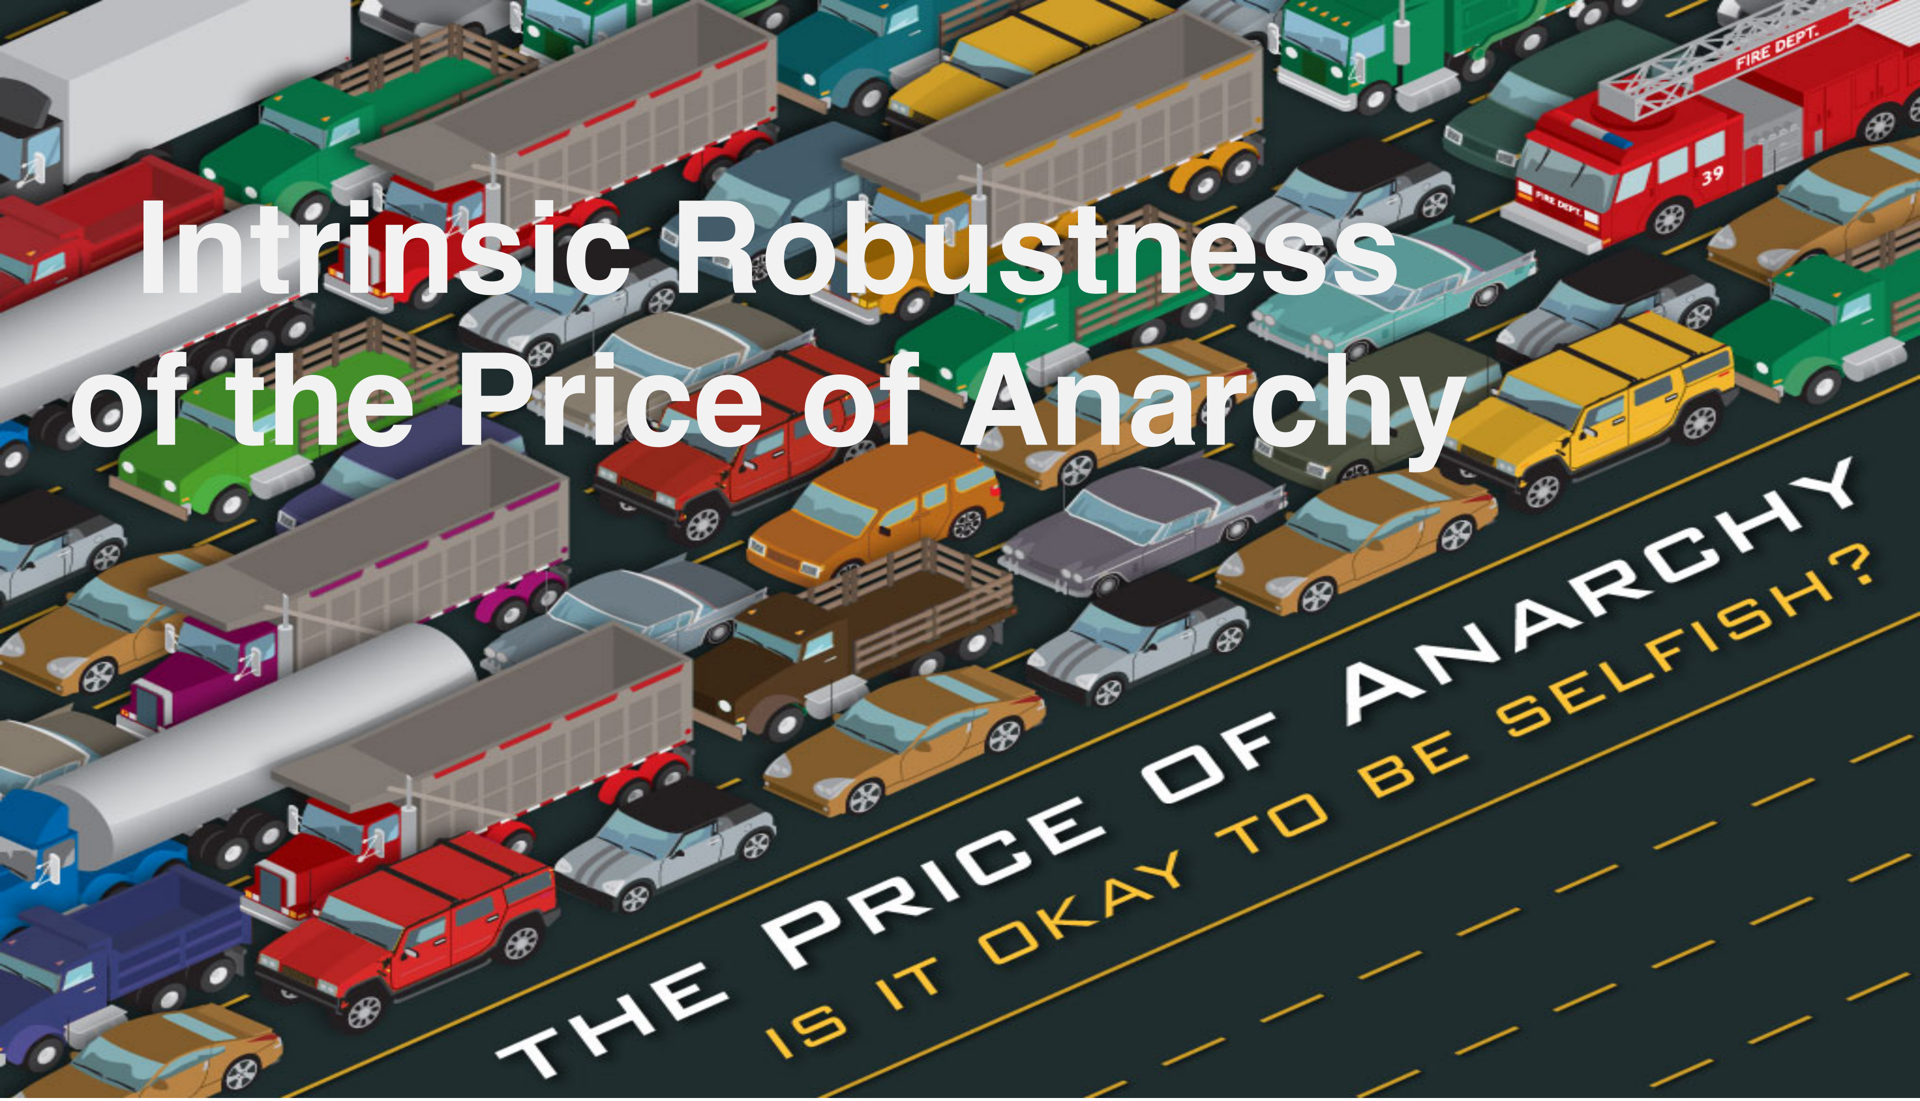

Tim Roughgarden (Stanford University)

2016 Kalai Prize Lecture

Kalai

The Price of Anarchy

Nash Equilibrium:

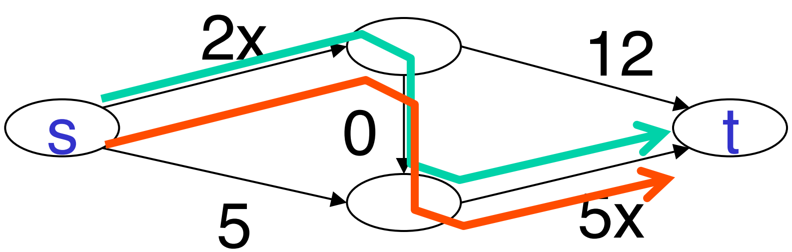

cost = 14+14 = 28

To Minimize Cost:

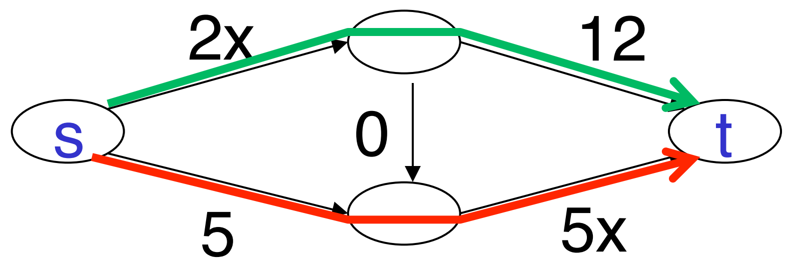

cost = 14+10 = 24

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

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

Price of Anarchy: Definition

Definition: [Koutsoupias/Papadimitriou 99]

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

Koutsoupias

Papadimitriou

A Representative Result

Theorem: [Roughgarden/Tardos 00] POA is at most 4/3 in every nonatomic selfish routing network with affine cost functions.

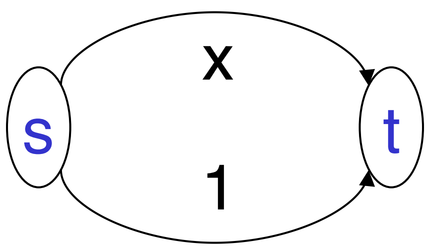

tight example (Pigou, 1920)

Tardos

The POA Goes Viral

Example domains: scheduling, routing, facility location, bandwidth allocation, network formation, network cascades, contention resolution, coordination games, firm competition, auctions, ...

The Price of Anarchy of Health Care

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

The price of anarchy in basketball

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

Do Players Reach an Equilibrium?

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: do our POA bounds really apply?

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]

Mirrokni

Vetta

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]

Christodoulou

Koutsoupias

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]
- coarse correlated equilibria (≈ no-regret sequences) [Blum/Even-Dar/Ligett 06], [Blum/Hajiaghayi/Ligett/Roth 08]


Blum


Hajiaghayi


Ligett


Roth

A Hierarchy of Equilibria

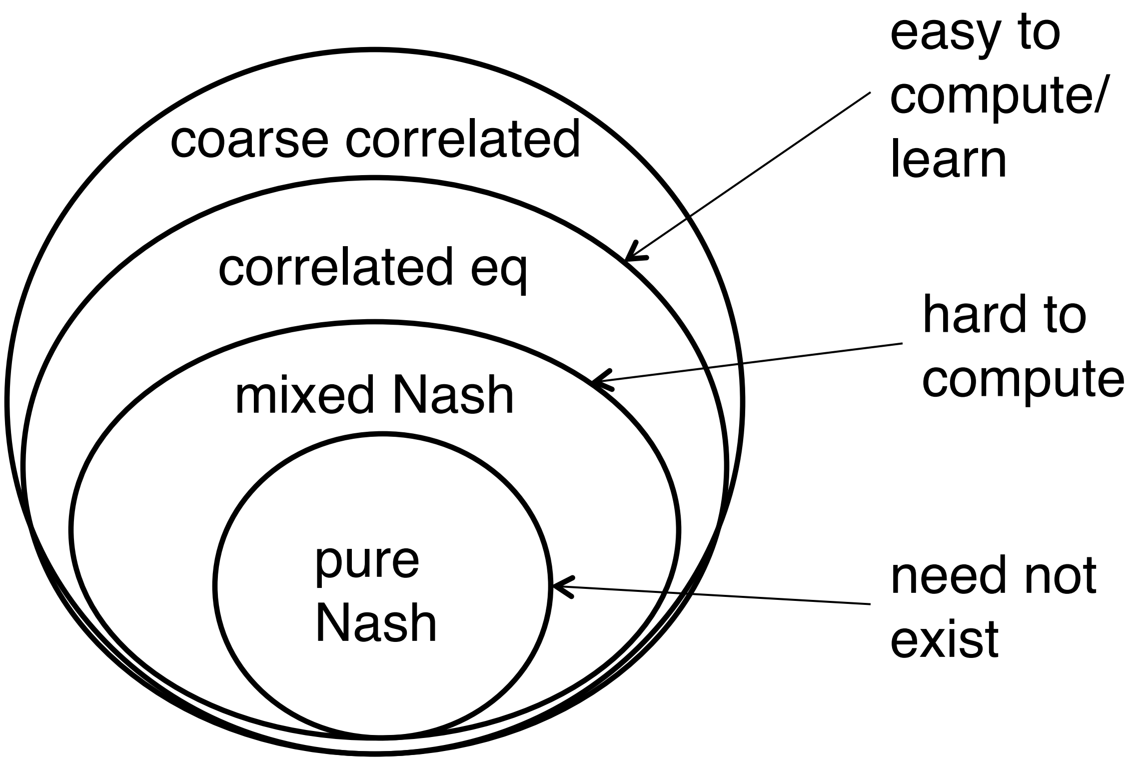

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

POA Bounds Without Convergence

Theorem: [Roughgarden 09] most known POA bounds hold *even if players do not reach a Nash equilibrium!*

Extension Theorems

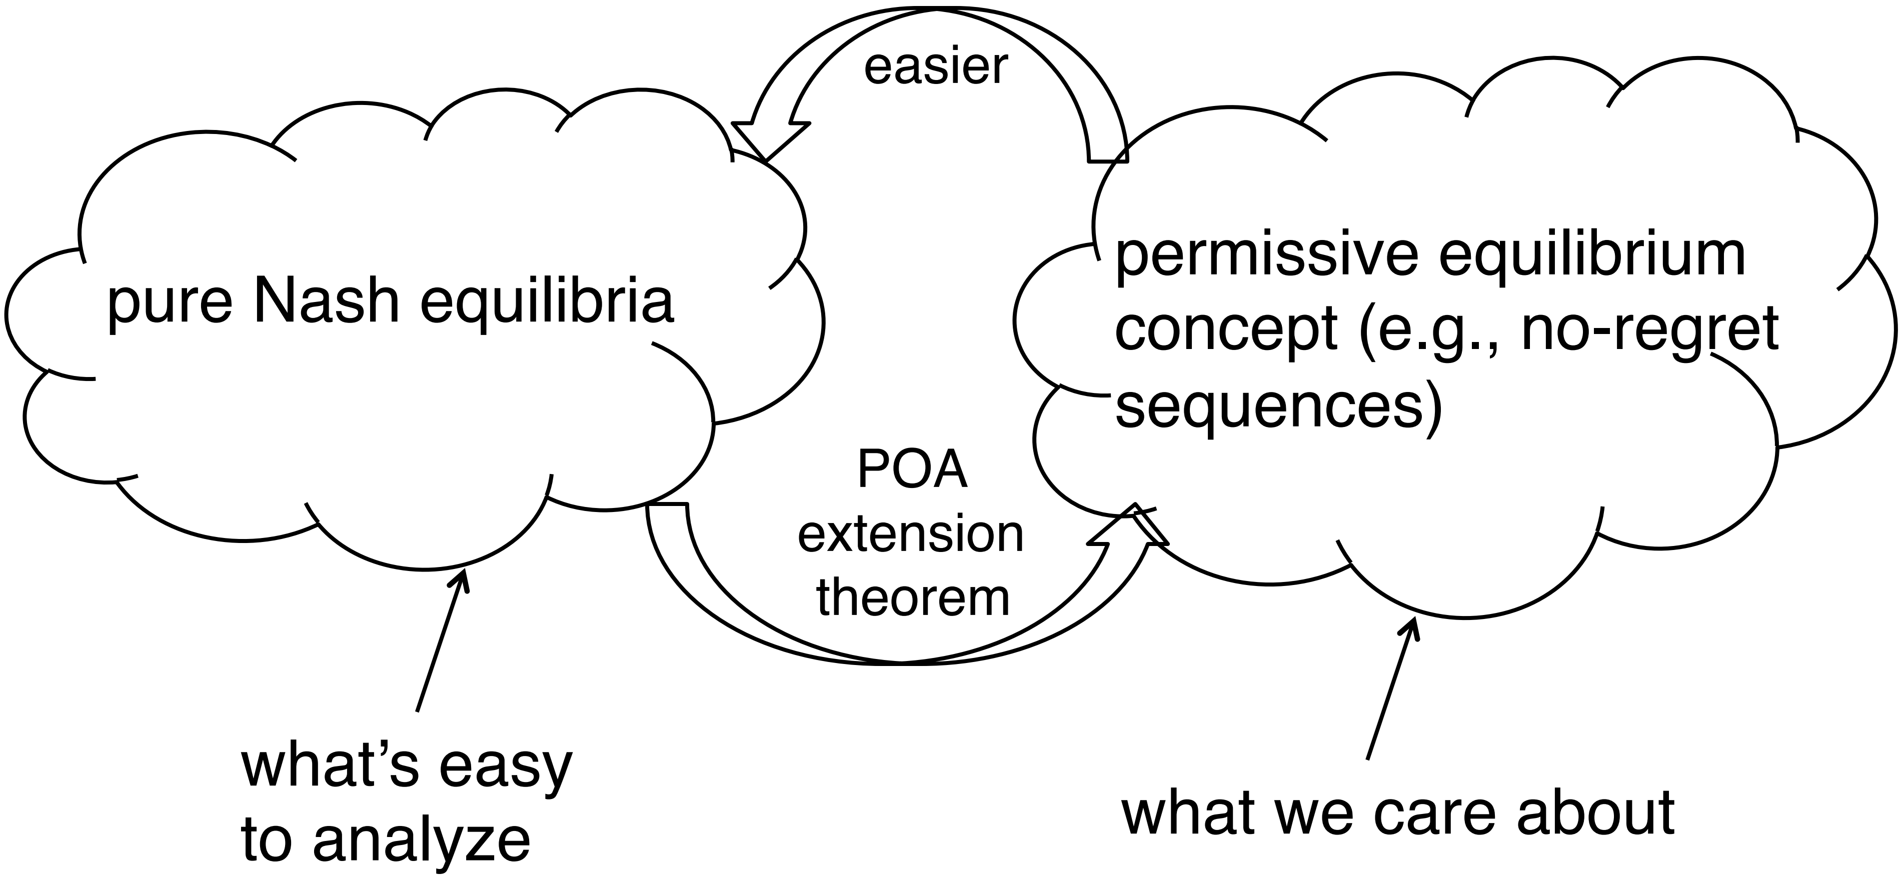

POA Bounds Without Convergence

Theorem: [Roughgarden 09] most known POA bounds hold *even if players do not reach a Nash equilibrium!*

Part I: [extension theorem] every POA bound proved for pure Nash equilibria *in a prescribed way* extends automatically, with no quantitative loss, to all coarse correlated equilibria.

- eludes non-existence/intractability critiques.

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

The Math

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

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

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

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

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]

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

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

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

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], ...

An Out-of-Equilibrium Bound

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

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

(the same bound we proved for pure Nash equilibria)

No-Regret Sequences

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

- for each player 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:

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

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

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

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

To finish proof: divide through by T.

Intrinsic Robustness

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

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

congestion games w/cost functions in C

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

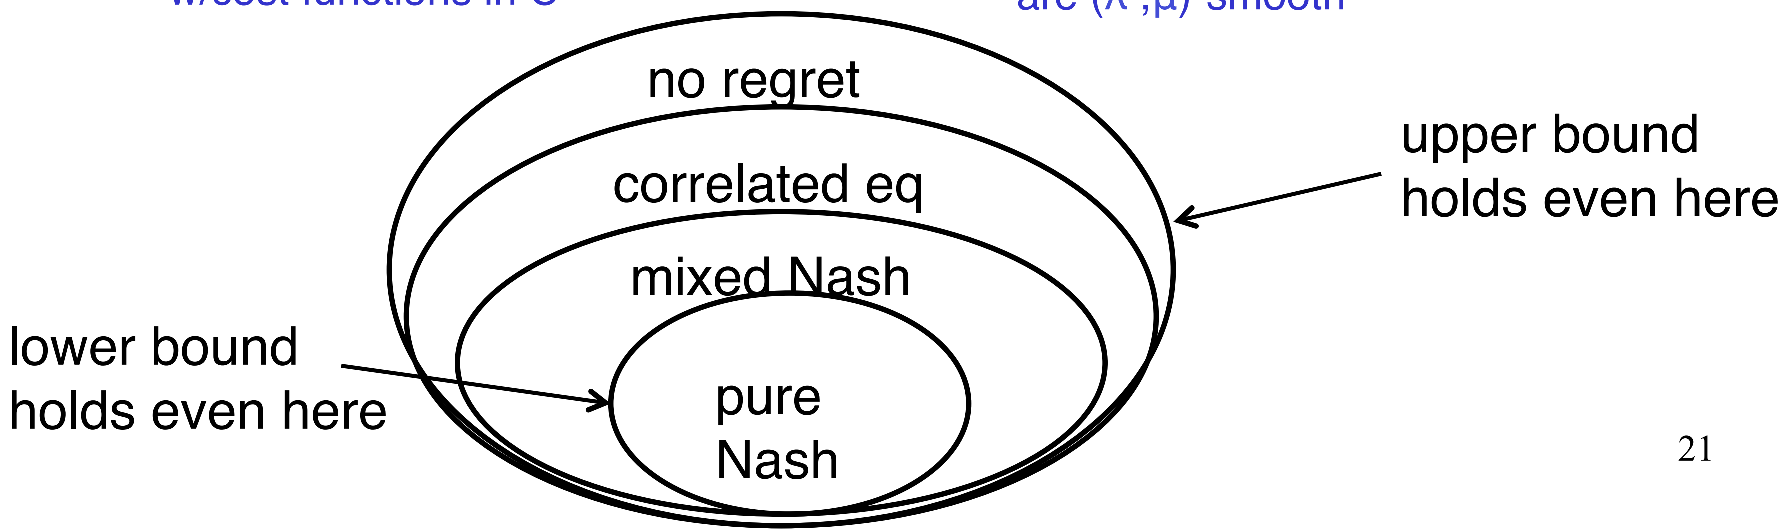

Incomplete Information

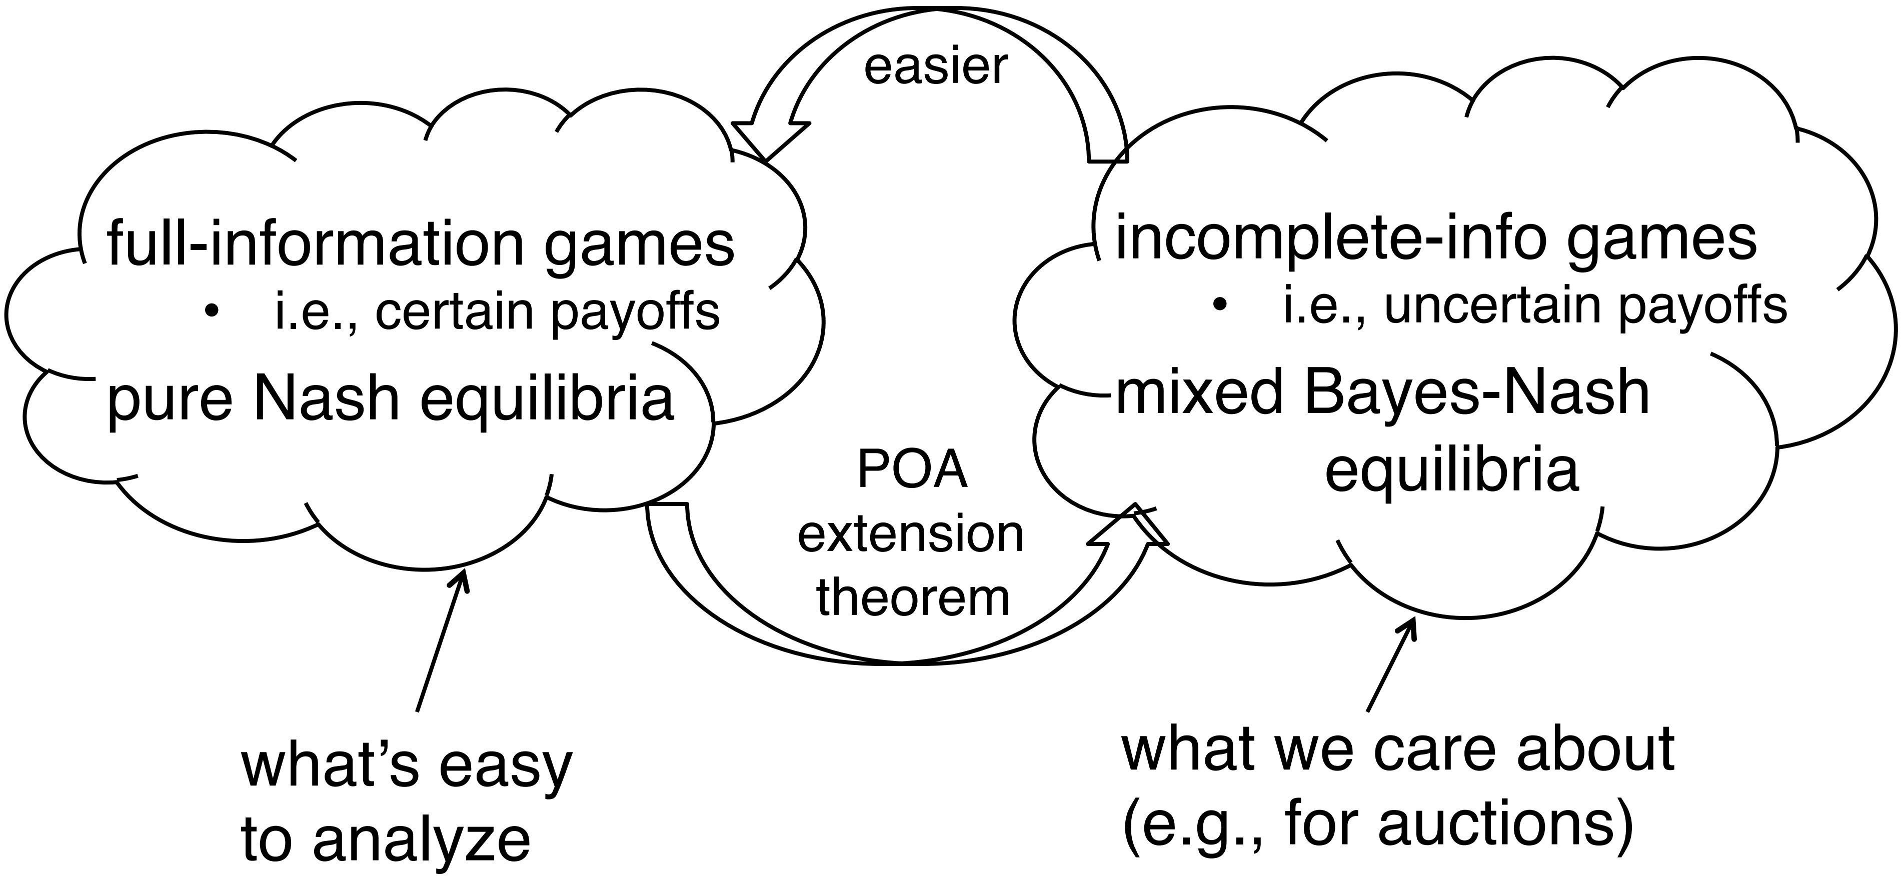

[Lucier/Paes Leme 11], [Roughgarden 12],
[Syrgkanis 12], [Syrgkanis/Tardos 13], ...

When Do Simple Mechanisms Suffice?

(P. Cramton, "The Efficiency of FCC Spectrum Auctions", 1998)

"The setting of spectrum auctions is too complex to guarantee full efficiency...Nonetheless, an examination of the bidding suggests that these problems, although present, probably did not lead to large inefficiencies."

Folklore belief: without strong complements, simple auctions work pretty well.

- loss in outcome quality appears small
- demand reduction exists, but not a dealbreaker

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, and
- for every (mixed) Bayes-Nash equilibrium,

expected welfare within 63% of max possible.

Modular proof: first-price auction is smooth + composition theorem + extension theorem

Conclusions

- price of anarchy: informative in many different application domains
  - try it on your favorite model!
- scales to large, complex games
  - where equilibrium characterization is hopeless
- user-friendly toolbox for proving POA bounds
  - extension theorems, composition theorems, etc.

THANKS!