Extension Theorems for the Price of Anarchy

Tim Roughgarden (Stanford University)

PRICE OF ANARCHY (1G)

The Price of Anarchy

Network with 2 players:

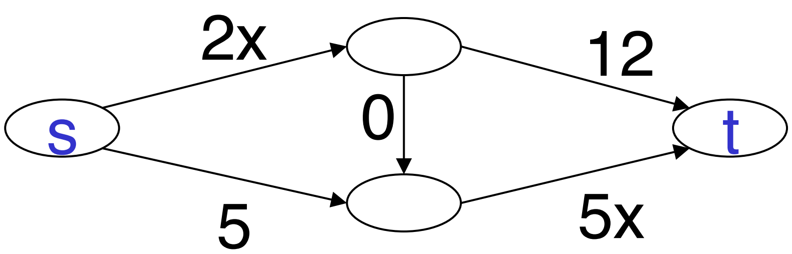

The Price of Anarchy

Nash Equilibrium:

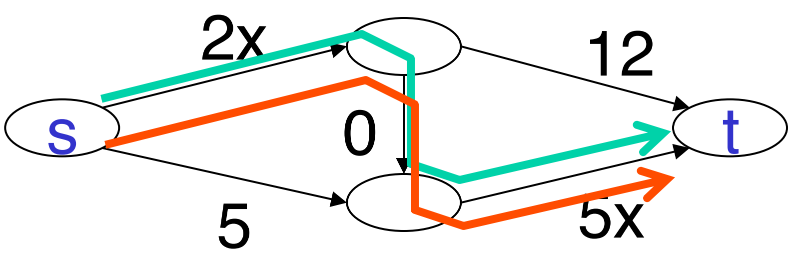

cost = 14+14 = 28

The Price of Anarchy

Nash Equilibrium:

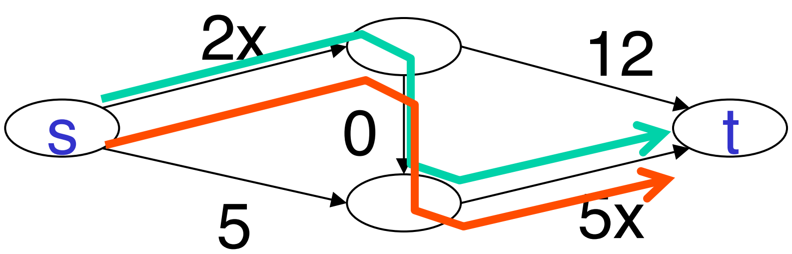

cost = 14+14 = 28

To Minimize Cost:

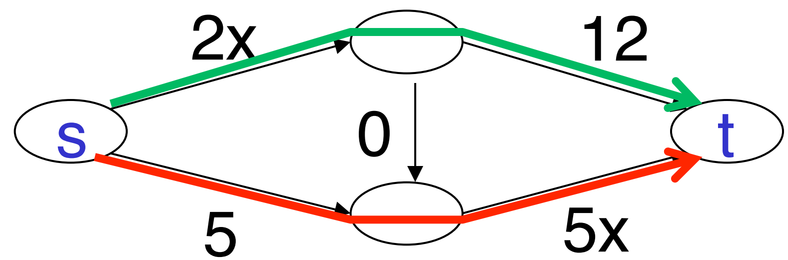

cost = 14+10 = 24

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

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

Price of Anarchy: Definition

Definition: [Koutsoupias/Papadimitriou STACS 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

Well-studied goal: when is the POA small?

– benefit of centralized control is small
– can suggest engineering rules of thumb: [Roughgarden STOC 02]: *10% extra network capacity guarantees POA for routing ≤ small constant*

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

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.

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

PRICE OF ANARCHY (2G)

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"?

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] most known POA bounds hold *even if system is not at Nash equilibrium!*

- e.g., if game is played repeatedly, no-regret conditions or a few myopic best responses are enough

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 aka “price of total anarchy” aka “no-regret players”
  - [Blum/Even-Dar/Ligett 06], [Blum/Hajiaghayi/Ligett/Roth 08]

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] most known POA bounds hold *even if system is not at Nash equilibrium!*

A Hierarchy of Equilibria

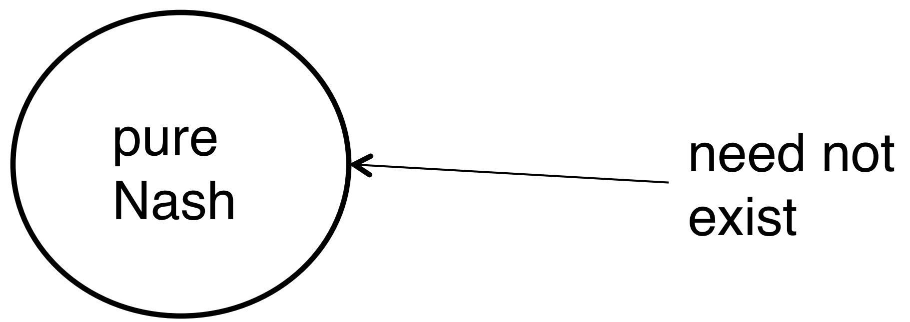

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

A Hierarchy of Equilibria

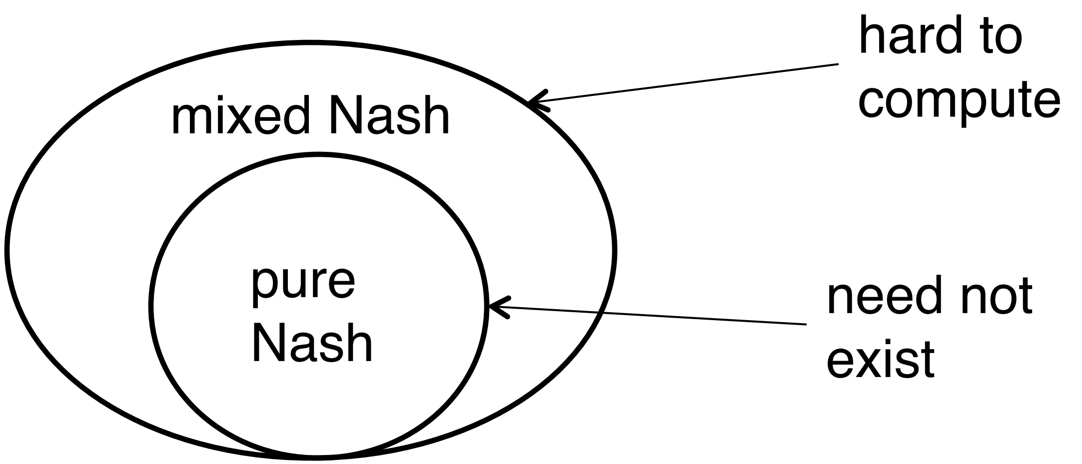

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

A Hierarchy of Equilibria

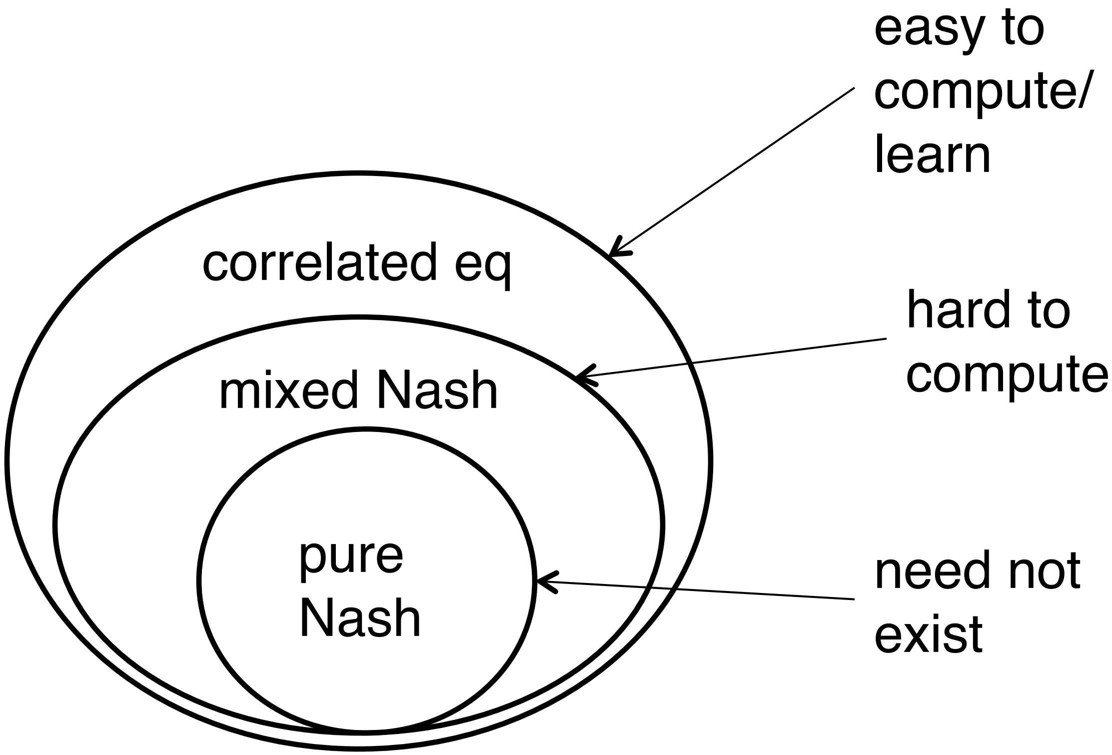

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

A Hierarchy of Equilibria

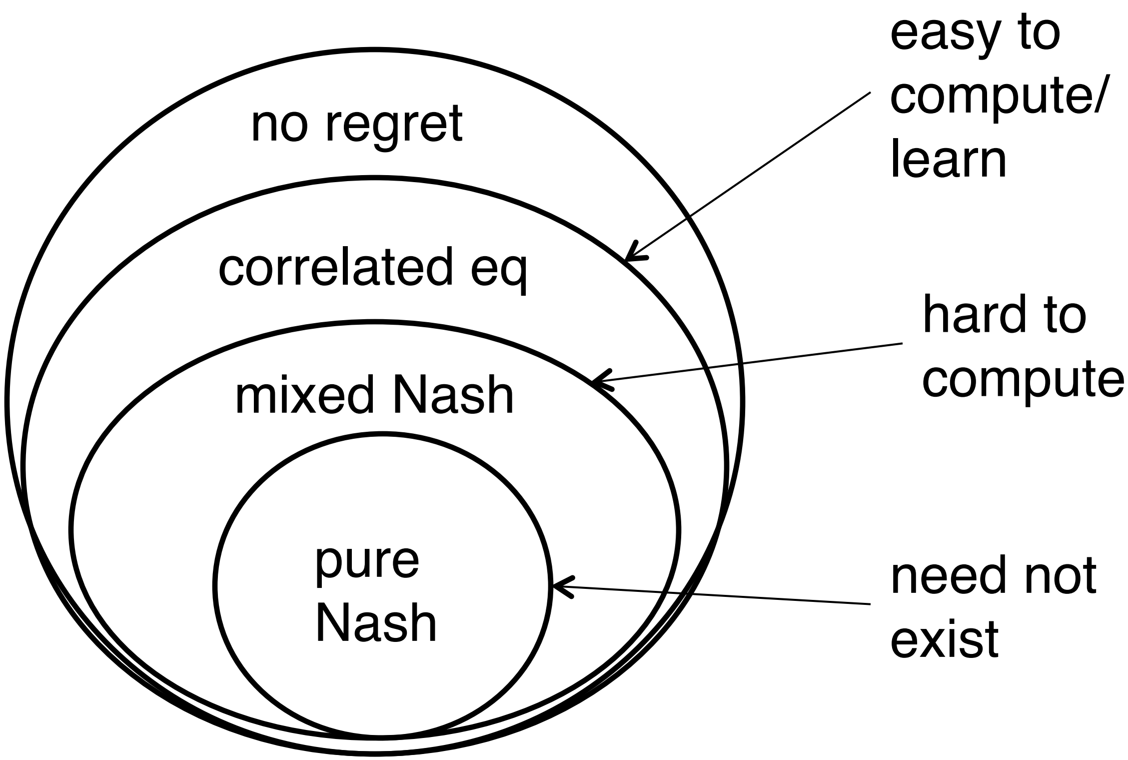

Recall: POA determined by *worst* equilibrium (only increases with the equilibrium set).

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] most known POA bounds hold *even if system is not at Nash equilibrium!*

Extension Theorems

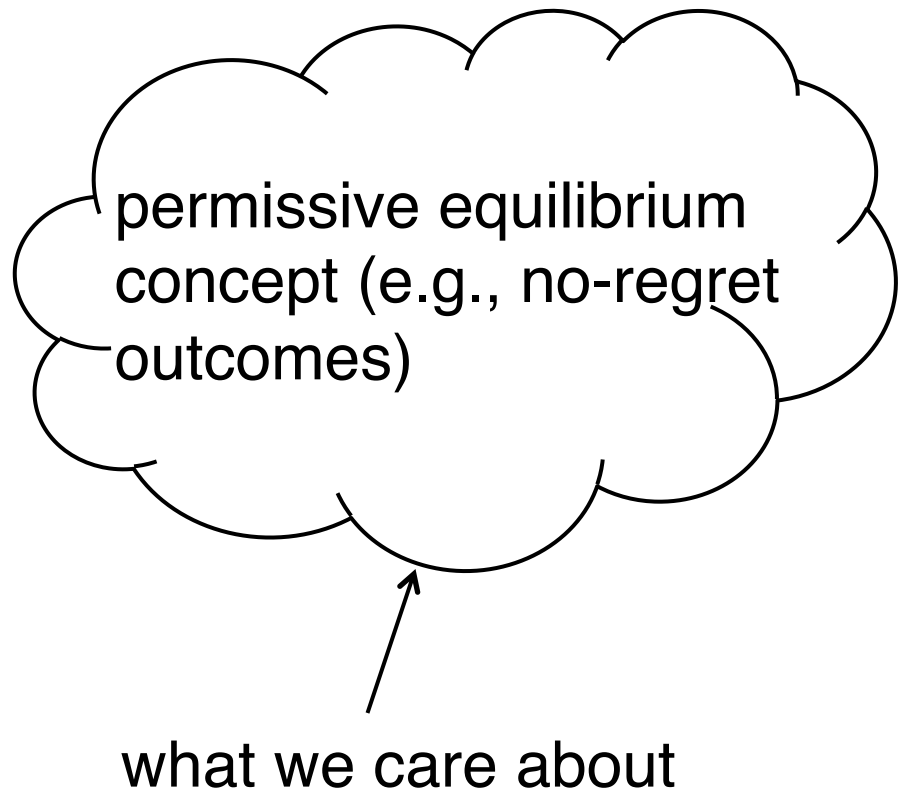

Extension Theorems

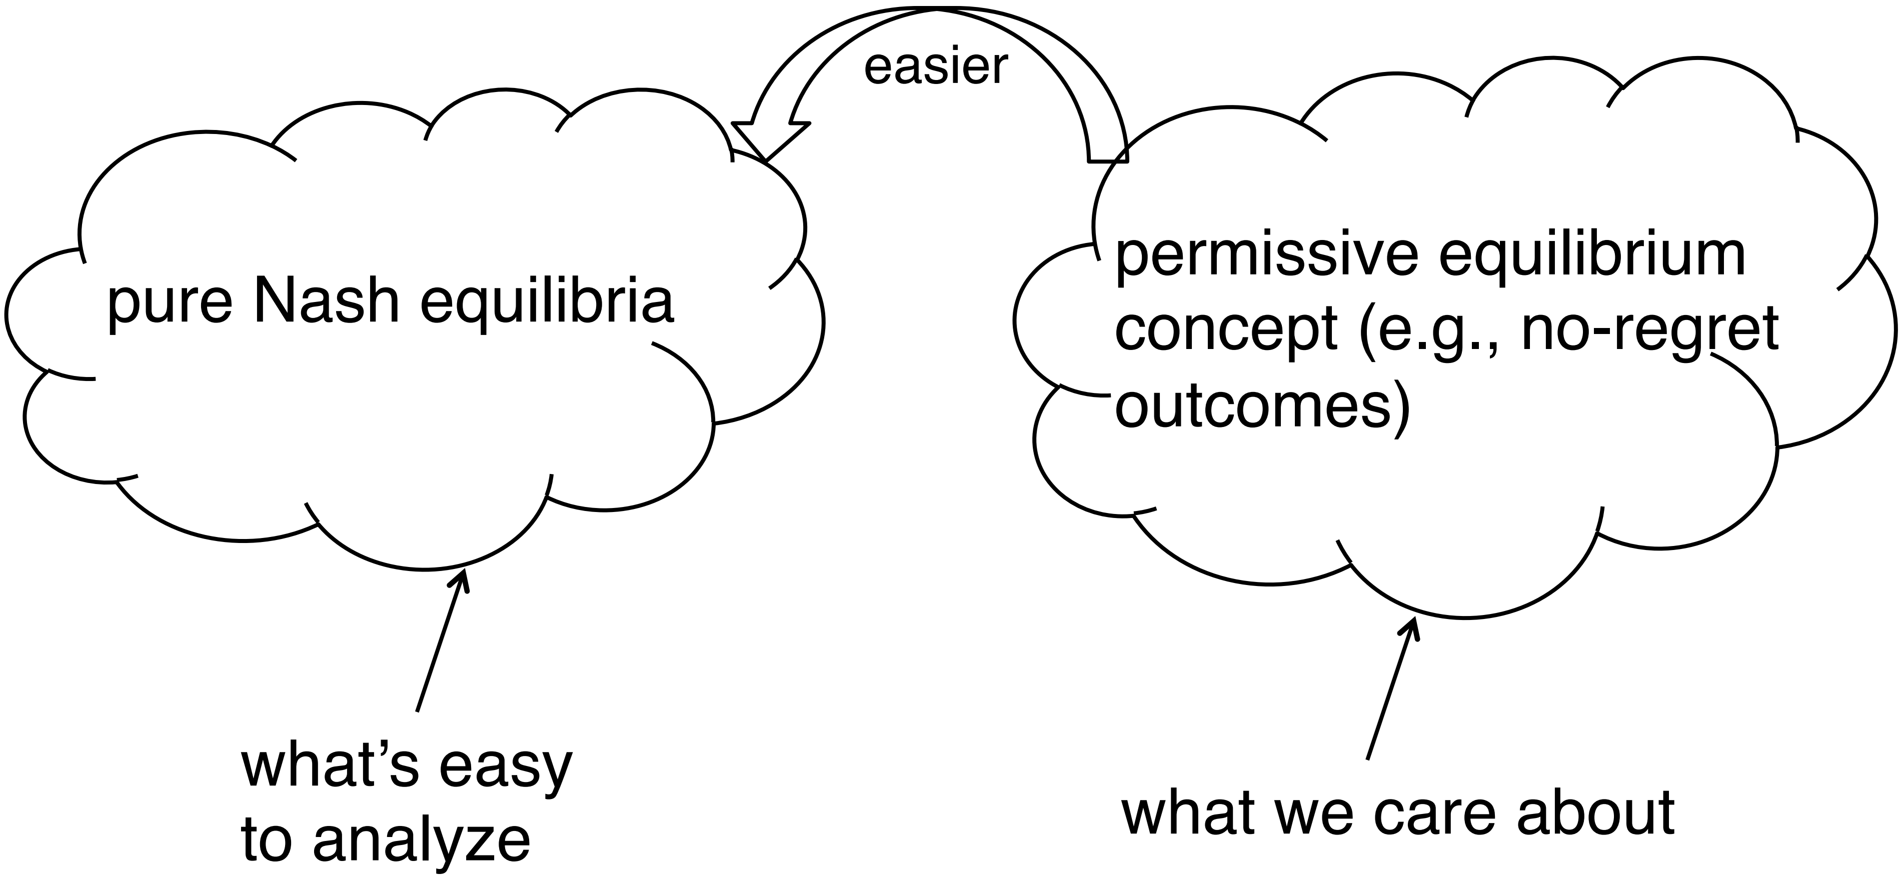

Extension Theorems

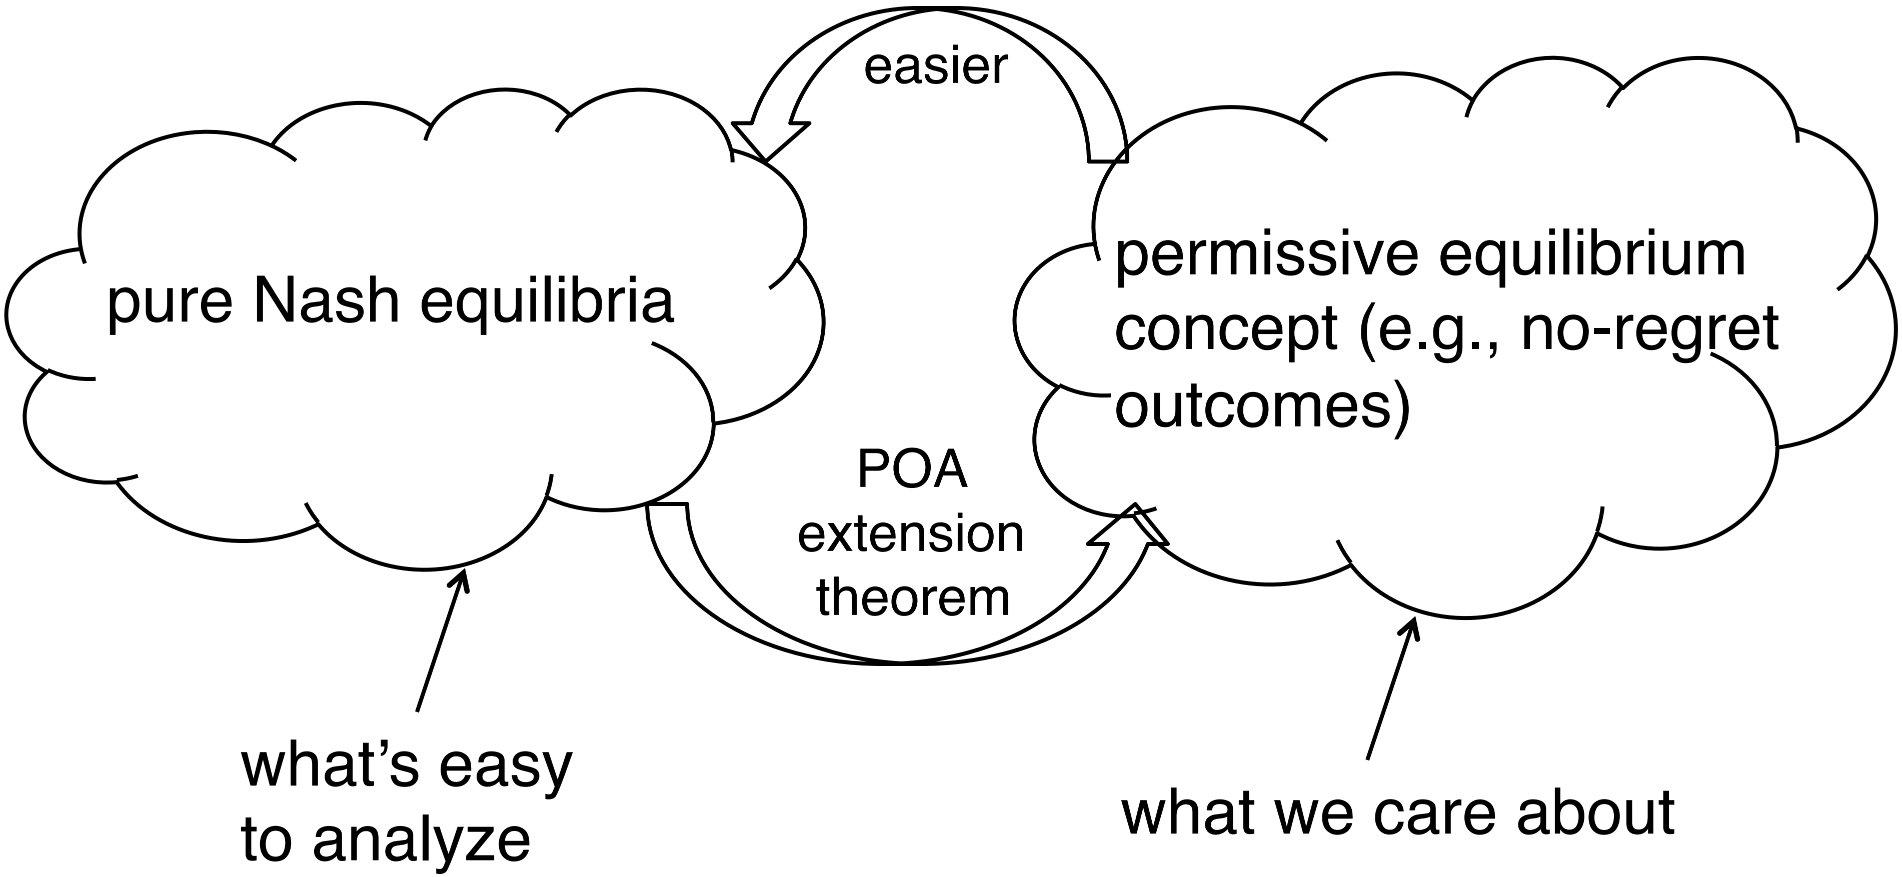

Extension Theorems

Worries about proving robust bounds:

- Approximation guarantee could get worse
- Seems like a lot of work!

"Extension Theorem": automatically extends a POA bound for pure Nash equilibria to more general equilibria, with no approximation loss.

Problem: too good to be true?

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] most known POA bounds hold *even if players do not reach a Nash equilibrium!*

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] 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 no-regret outcomes.

- eludes non-existence/intractability critiques.

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] 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 no-regret outcomes.

- eludes non-existence/intractability critiques.

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

How To Bound the POA?

Goal: prove that every pure Nash equilibrium has cost close to the minimum possible.

Observation: proof must apply NE hypothesis (once per player, with a candidate deviation).

"Smoothness proof": makes only this minimal use of equilibrium hypothesis.

- rest of proof should combine these n inequalities with game structure to yield a POA bound
- candidate deviations = independent of the Nash eq

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)

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

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

Smooth Games

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

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

Smooth Games

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

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

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

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

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

Note: only needed (*) to hold in special case where **s** = a Nash eq and **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]

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]

Some Smoothness Bounds

Claim: (5/3,1/3)-smoothness in atomic, affine case

- [Christodoulou/Koutsoupias 05]: for all integers y,z:

  $y(z+1) \leq (5/3)y^2 + (1/3)z^2$

Some Smoothness Bounds

Claim: (5/3,1/3)-smoothness in atomic, affine case

- [Christodoulou/Koutsoupias 05]: 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 ≥ 0

Some Smoothness Bounds

Claim: (5/3,1/3)-smoothness in atomic, affine case

- [Christodoulou/Koutsoupias 05]: 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 ≥ 0

- so: $\Sigma_e [a_e(x_e+1) + b_e)x_e^*] \leq (5/3) \Sigma_e [(a_e x_e^* + b_e)x_e^*]$

  $+ (1/3) \Sigma_e [(a_e x_e + b_e)x_e]$

Some Smoothness Bounds

Claim: (5/3,1/3)-smoothness in atomic, affine case

- [Christodoulou/Koutsoupias 05]: 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 STOC 09]
in a (λ,μ)-smooth game, average cost of every no-regret sequence is at most

$[\lambda/(1-\mu)]$ 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 (for suff large T).

- [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:

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

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

Smooth => No-Regret Bound

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

Assuming (λ,μ)-smooth:

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

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

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

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]

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

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

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

To finish proof: divide through by T.

The Limits of Smoothness

Theorem: [Nadav/Roughgarden WINE 10] Consider a (λ,μ)-smooth game for optimal choices of λ,μ. Then there is an "aggregate" coarse correlated equilibrium with cost = [λ/(1-μ)] x OPT.

- optimal smoothness bound governed by worst distribution with non-positive average (rather than per-player) regret with respect to optimal $\mathbf{s}^*$

Proof: convex duality.

Intrinsic Robustness

Theorem: [Roughgarden STOC 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

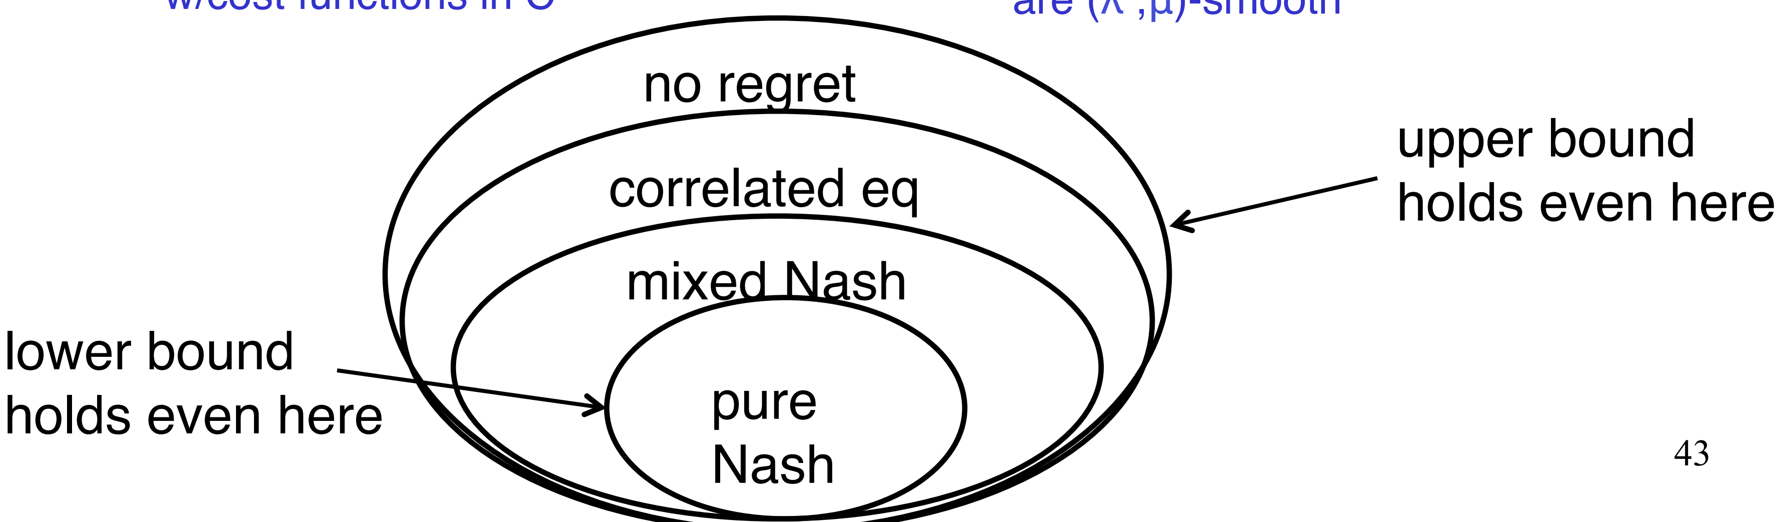

POA with Incomplete Information: The Best-Case Scenario

Observation: the more general the equilibrium concept, the worse the POA.

- full-information Nash equilibria = special case of incomplete-info Bayes-Nash equilibria (fixed type vector)

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.

Extension Theorem (Informal)

Consider a game of incomplete information *with stochastically independent types.*

- fixing the (random) types induces a full-info game

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

Conclusion: the POA of (mixed) Bayes-Nash equilibria is at least α.

- no matter what the common prior distribution is

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)

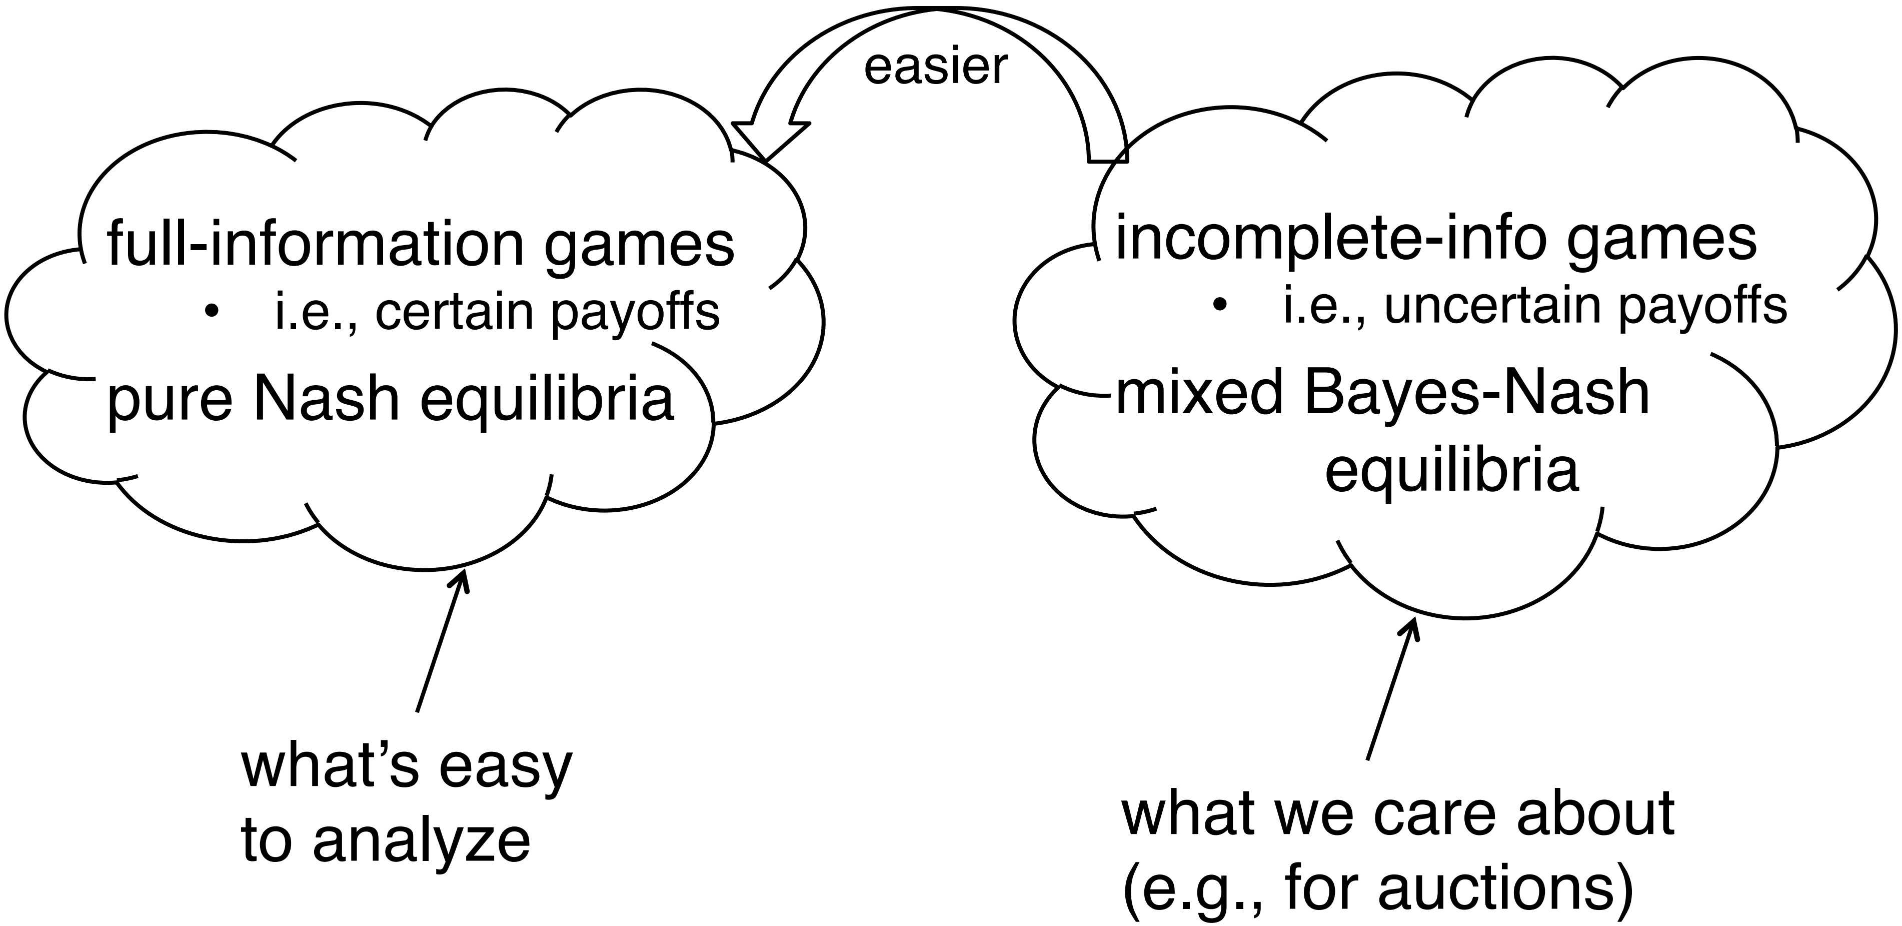

Extension Theorem (Informal)

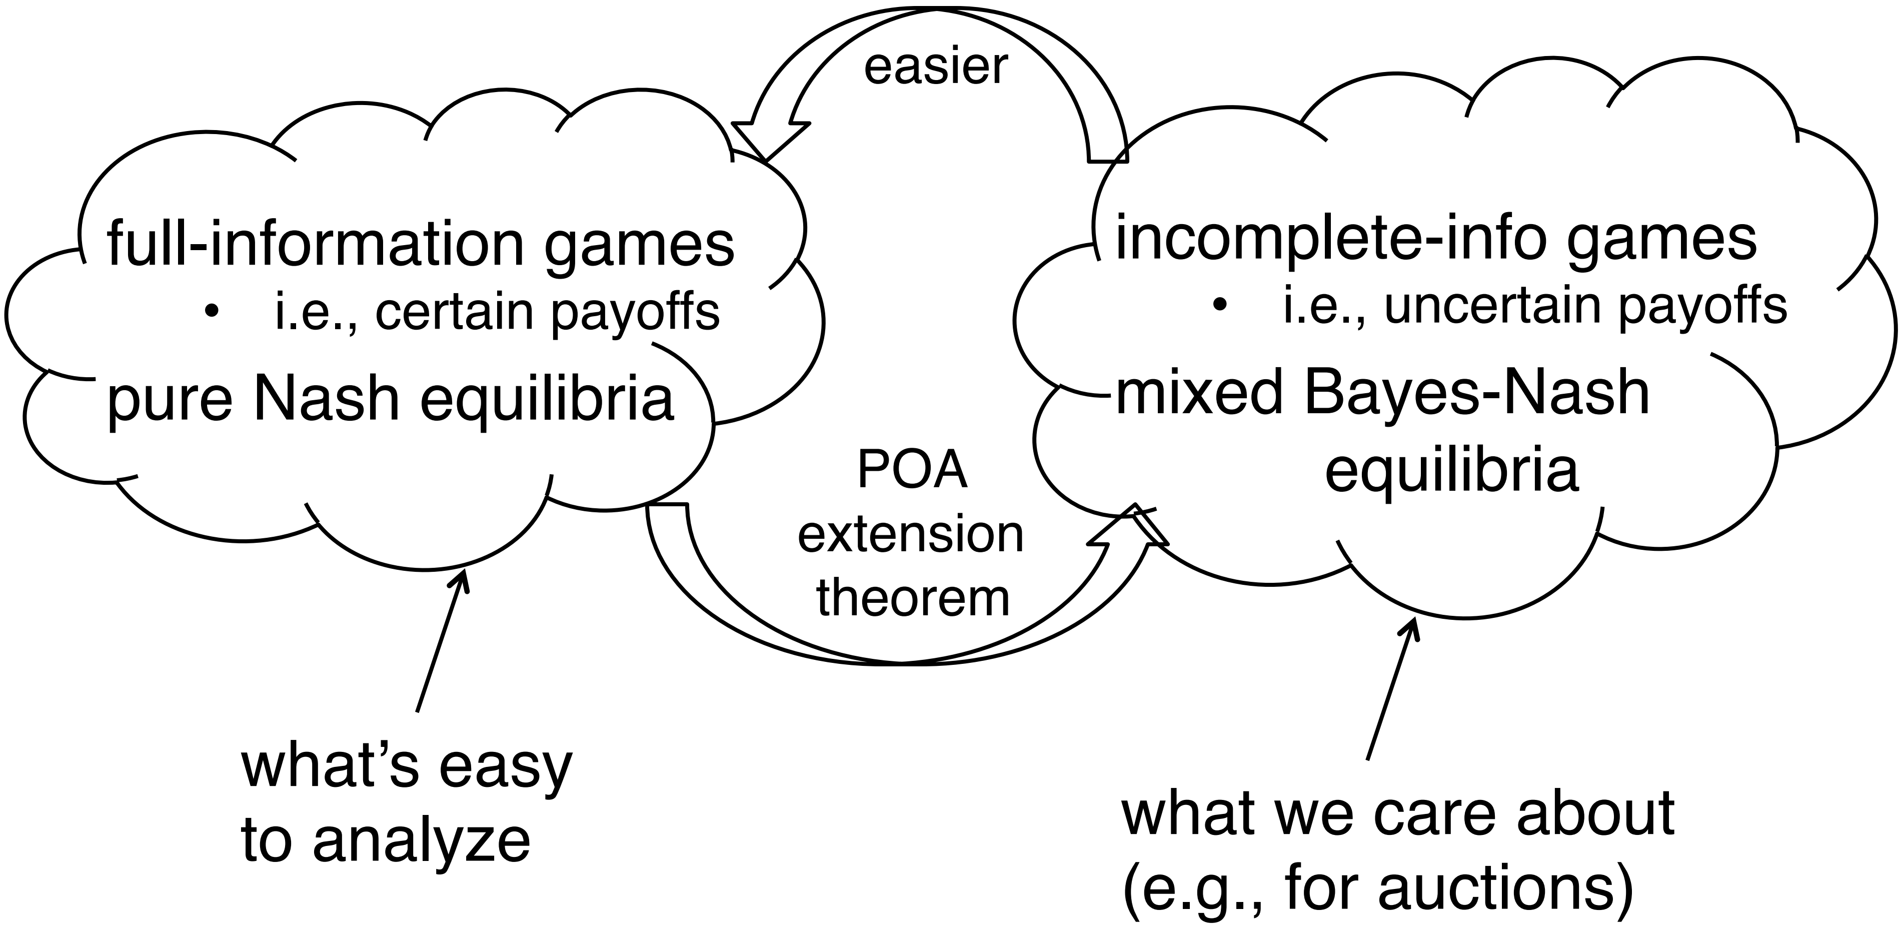

Extension Theorem (Formal)

Theorem: [Roughgarden EC 12, Syrgkanis 12] for every (λ,μ)-smooth game of incomplete information, and every prior product distribution over types, the POA of (mixed) Bayes-Nash equilibria is at most $\lambda/(1-\mu)$.

Proof idea: in 3 stages.

1. Use Bayes-Nash equilibrium hypothesis.
2. Use independence hypothesis.
3. Use smoothness hypothesis.

Applications

- simultaneous single-item auctions [Christodoulou/Kovacs/Schapira 08], [Bhawalkar/Roughgarden 11]
- greedy combinatorial auctions [Lucier/Borodin 10]
- sponsored search auctions [Caragiannis/Kaklamanis/Kanellopolous/Kyropoulou/Lucier/Paes Leme/Tardos 12]
- routing games with incomplete info (new)
  - full-information POA bounds carry over to uncertain source-sink pairs and/or player weights

Necessity of Independent Types

[Bhawalkar/Roughgarden SODA 11] extension theorem false without independence

[Caragiannis/Kaklamanis/Kanellopolous/Kyropoulou/Lucier/Paes Leme/Tardos 12]: give additional conditions for extension theorem to work with correlated player types.

- application: sponsored search auctions

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 models
- extensions: local smoothness for correlated equilibria; also Bayes-Nash equilibria with independent types