Lecture 3

Large Routing Games: Wardrop or Poisson?

Roberto Cominetti
Universidad Adolfo Ibáñez

Based on joint work with:

Marco Scarsini (LUISS)
Marc Schröder (RWTH Aachen)
Nicolás Stier-Moses (Facebook)

Journées SMAI MODE 2020

1. Convergence of Large Games
   - Splittable routing games $\longrightarrow$ Wardop equilibrium
   - Weighted atomic games $\longrightarrow$ Wardop equilibrium
   - Games with random players $\longrightarrow$ Poisson equilibrium
   - Convergence of PoA for sequences of ARGs

2. Price-of-Anarchy for Atomic Routing Games
   - Smoothnes framework
   - PoA for Bernoulli ARGs

You are planning your commute route for tomorrow, not sure about exact departure time nor who will be on the road...

Games with *"many small players"* are frequently modeled as nonatomic games with a continuum of players.

In which sense is a continuous model close to the discrete system ?

The answer depends on what we mean by *"small players"*...

- player i has a small load $w_i \approx 0$ to be transported with certainty
- player i has a unit load but is present with small probability $p_i \approx 0$

Each interpretation yields a different continuous limit.

Network Routing Games

We are given a graph (V, E) with

- a set of *edges* $e \in E$ with continuous non-decreasing costs $c_e : \mathbb{R}_+ \to \mathbb{R}_+$
- a set of *OD pairs* $\kappa \in \mathscr{K}$ with corresponding routes $r \in \mathscr{R}_\kappa \subseteq 2^E$
- a set of *demands* $d_\kappa \geq 0$ for each $\kappa \in \mathscr{K}$

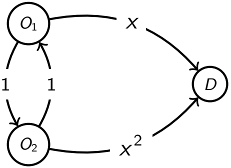

Demands can be...

- non-atomic: continuous, infinitesimal players $\to$ Wardrop
- atomic
  - splittable: continuous, few players $\to$ fluids, sand, telecom
  - unsplittable: discrete, few players $\to$ vessels, airplanes
  - random: unpredictable $\to$ packets or vehicles on a network

Non-Atomic Routing Games — Wardrop equilibrium

Let $\mathscr{F}$ be the set of splittings (y, x) of the demands d_κ into *route-flows* $y_r \geq 0$, together with their induced *edge-loads* x_e :

$$d_\kappa = \textstyle\sum_{r\in\mathscr{R}_\kappa} y_r \quad (\forall\kappa \in \mathscr{K}),$$
$$x_e = \textstyle\sum_{r\ni e} y_r \qquad (\forall e \in E).$$

A Wardrop equilibrium *is a pair* $(\hat{y}, \hat{x}) \in \mathscr{F}$ *that uses only shortest routes:*

$$(\forall\kappa \in \mathscr{K})(\forall r, r' \in \mathscr{R}_\kappa) \quad \hat{y}_r > 0 \Rightarrow \sum_{e\in r} c_e(\hat{x}_e) \leq \sum_{e\in r'} c_e(\hat{x}_e).$$

Characterized as the optimal solutions of the convex program

$$\min_{(y,x)\in\mathscr{F}} \sum_{e\in E} \int_0^{x_e} c_e(z)\, dz.$$

Atomic Splittable Routing Games

Atomic Splittable Routing Games

Atomic splittable routing games are similar in that demands are continuous and can be split over different routes, except that now there are finitely many players $i \in N = \{1, \ldots, n\}$, each one controlling a non-negligible amount of traffic $d_i > 0$, on a given OD pair $\kappa_i \in \mathscr{K}$.

- Player i splits traffic over routes $y_{ir} \geq 0$ with $d_i = \sum_{r \in \mathscr{R}_{\kappa_i}} y_{ir}$
- This induces edge loads $x_{ie} = \sum_{r \ni e} y_{ir}$
- Player i's cost is $C_i(x) = \sum_{e \in E} x_{ie}\, c_e(x_{ie} + x_{\text{-}ie})$ with $x_{\text{-}ie} = \sum_{j \neq i} x_{je}$.

Definition

A *splittable Nash equilibrium* is a family of feasible flows $(y_i, x_i)_{i \in N}$ such that

$$C_i(x_i, x_{\text{-}i}) \leq C_i(x_i', x_{\text{-}i}) \qquad \forall i \in N,\ \forall x_i' \text{ feasible flow}$$

Atomic Splittable Routing Games

Debreu-Glicksberg-Fan's extension of Nash's Theorem yields

Theorem (Rosen, 1965)

1. $c_e(\cdot)$ *non-decreasing and convex* $\Rightarrow \quad \exists$ *splittable Nash equilibria.*
2. *...under suitable monotonicity conditions it is unique.*

As the number of players increases and their individual demands become smaller, one expects convergence towards a Wardrop equilibrium.

Proved by (Haurie-Marcotte, 1985; Milchtaich, 2000; Jacquot-Wan, 2018) tipically assuming that each player controls the same amount of traffic. The following more general statement seems to be new... though expected!

Convergence to Wardrop — Vanishing Traffic Demands

Let (y^n, x^n) be equilibria for a sequence of atomic splittable routing games with smooth costs $c_e(\cdot) \in C^1$ and

$$\begin{cases} a) & |N^n| \to \infty \\ b) & \max_{i \in N^n} d_i^n \to 0 \\ c) & d_\kappa^n \triangleq \sum_{i:\kappa_i^n = \kappa} d_i^n \to d_\kappa \quad \text{for all } \kappa \in \mathscr{K} \end{cases}$$

Theorem

1. *The aggregate route flows* $y_r^n = \sum_{i \in N} y_{ir}^n$ *and edge loads* $x_e^n = \sum_{i \in N} x_{ie}^n$ *are bounded and each limit point* $(\bar{y}, \bar{x})$ *is a Wardrop equilibrium for demands* d_κ*.*
2. *If the* c_e*'s are strictly increasing then* $\bar{x}$ *is unique and* $x^n \to \bar{x}$*.*

Weighted Atomic Routing Games

Weighted Atomic Routing Games— $\mathscr{G}(w)$

A *weighted routing game* has a finite set of players $i \in N$ with OD pairs $\kappa_i \in \mathscr{K}$, and unsplittable weights $w_i > 0$ that must be routed over a single path $r_i \in \mathscr{R}_{\kappa_i}$ chosen at random using a mixed strategy $\pi_i \in \Delta(\mathscr{R}_{\kappa_i})$.

- $Y_r = \sum_{i \in N} w_i \, \mathbb{1}_{\{r_i = r\}}$ are the random route-flows
- $X_e = \sum_{i \in N} w_i \, \mathbb{1}_{\{e \in r_i\}}$ are the corresponding edge-loads

Definition

A mixed strategy profile $\pi = (\pi_i)_{i \in N}$ is a Nash equilibrium *iff for each player $i \in N$ and routes $r, r' \in \mathscr{R}_{\kappa_i}$ with $\pi_i(r) > 0$ we have*

$$\mathbb{E}\left[\sum_{e \in r} c_e(X_e) | r_i = r\right] \leq \mathbb{E}\left[\sum_{e \in r'} c_e(X_e) | r_i = r'\right]$$

- Weighted ARGs with identical weights $w_i \equiv w$ are potential games and admit pure equilibria (Rosenthal'73). The potential for a profile $\mathbf{r} = (r_i)_{i \in N}$ is

$$\Phi(\mathbf{r}) = \sum_{e \in E} \sum_{k=1}^{n_e(\mathbf{r})} c_e(k\,w)\,w \qquad ; \qquad n_e(\mathbf{r}) \triangleq |\{i \in N : e \in r_i\}|.$$

- For heterogeneous weights we only have existence of mixed equilibria.

Example: Routing n players over 2 identical parallel links.

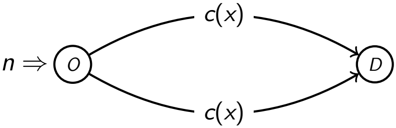

Symmetric mixed equilibrium: each player randomizes $(\frac{1}{2}, \frac{1}{2})$.

If players' weights are $w_i \equiv d/n$ then we have random edge-loads

$$X_e \sim \tfrac{d}{n}\,\mathrm{Binomial}(n, \tfrac{1}{2})$$

which converge almost surely to the Wardrop equilibrium $(\frac{d}{2}, \frac{d}{2})$.

What happens for non-symmetric equilibria? What if weights are not homogeneous? And with different costs? And more complex topologies?

Wardrop Convergence for Vanishing Weights

Let π^n be a sequence of mixed equilibria for weighted ARGs $\mathscr{G}(w^n)$ with

$$\begin{cases} a) & |N^n| \to \infty \\ b) & \max_{i\in N^n} w_i^n \to 0 \\ c) & d_\kappa^n \triangleq \sum_{i:\kappa_i^n=\kappa} w_i^n \to d_\kappa \quad \text{for all } \kappa \in \mathscr{K} \end{cases}$$

Theorem

1. *The expected flows $(y^n, x^n) = (\mathbb{E}Y^n, \mathbb{E}X^n)$ are bounded and each cluster point $(\hat{y}, \hat{x})$ is a Wardrop equilibrium with demands d_κ and costs $c_e(\cdot)$.*
2. *Along any convergent subsequence, the random route-flows and edge-loads (Y^n, X^n) converge in L^2 to the (constant) Wardrop equilibrium $(\hat{y}, \hat{x})$.*
3. *If the costs $c_e(\cdot)$ are strictly increasing, then $\hat{x}$ is unique and $X^n \xrightarrow{L^2} \hat{x}$.*
4. *If $c_e \in C^2$ with $c_e'(\cdot) > 0$, then there is a constant κ such that*

$$\|X^n - \hat{x}\|_{L^2} \leq \kappa\left(\sqrt{\max_{i\in N} w_i^n} + \sqrt{\|d^n - d\|_1}\right).$$

Nice and simple... but reality looks more like this

Copenhagen – *Source: DTU Transport (www.transport.dtu.dk)*

Figure 2: Example of real time illustration of congestion (Source: Vejdirektoratet, www.trafikken.dk)

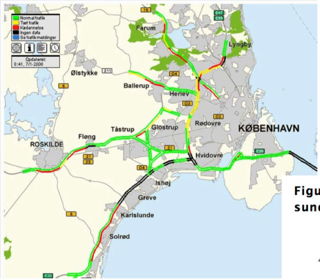

Figure 7: Observations of travel time by time of day. Frederikssundsvej, inward direction

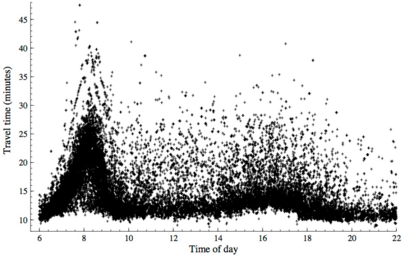

Traffic count data – Dublin 2017-2018

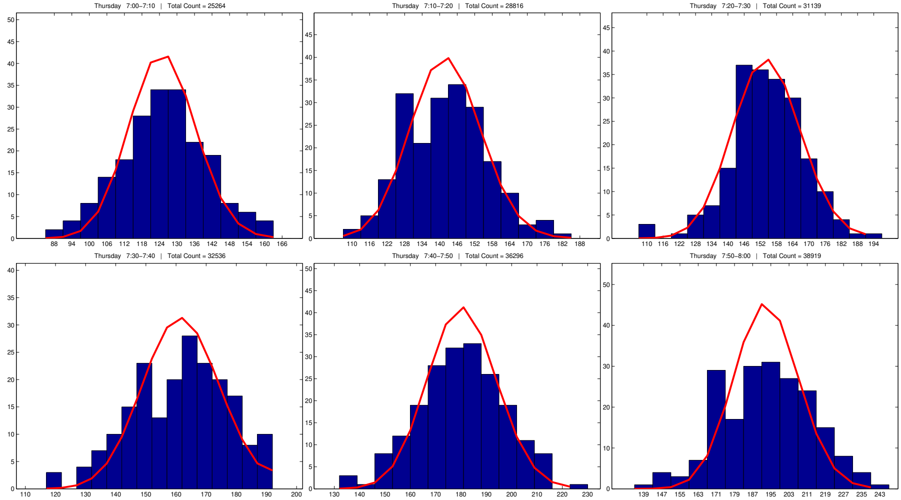

Data source: Transport Infrastructure Ireland

https://www.tii.ie/roads-tolling/operations-and-maintenance/traffic-count-data/

Bernoulli Routing Games

Bernoulli Atomic Routing Games — $\mathscr{G}(p)$

A *Bernoulli routing game* has a finite set of players $i \in N$ with OD pairs $\kappa_i \in \mathscr{K}$, unit weights $w_i = 1$, and a probability of being active $p_i = \mathbb{P}(U_i = 1)$. Each player $i \in N$ selects a route $r_i \in \mathscr{R}_{\kappa_i}$ using a mixed strategy $\pi_i \in \Delta(\mathscr{R}_{\kappa_i})$.

- $Y_r = \sum_{i \in N} U_i \, \mathbb{1}_{\{r_i = r\}}$ are the random route-flows
- $X_e = \sum_{i \in N} U_i \, \mathbb{1}_{\{e \in r_i\}}$ are the corresponding edge-loads

Definition

A strategy profile $\pi = (\pi_i)_{i \in N}$ is a Bayes-Nash equilibrium *if for each player $i \in N$ and routes $r, r' \in \mathscr{R}_{\kappa_i}$ with $\pi_i(r) > 0$ we have*

$$\mathbb{E}\left[\sum_{e \in r} c_e(X_e) | U_i = 1, r_i = r\right] \leq \mathbb{E}\left[\sum_{e \in r'} c_e(X_e) | U_i = 1, r_i = r'\right].$$

REMARK. Costs need only be defined over the integers $c_e : \mathbb{N} \to \mathbb{R}_+$.

Bernoulli ARGs are Potential Games

Proposition

Every Bernoulli ARG is a potential game with potential

$$\Phi(\mathbf{r}) \triangleq \mathbb{E}\left[\sum_{e\in E}\sum_{k=1}^{N_e(\mathbf{r})} c_e(k)\right] \quad ; \quad N_e(\mathbf{r}) \triangleq \sum_{i:e\in r_i} U_i$$

Corollary

Every Bernoulli ARG has Nash equilibria in pure strategies.

Example: Routing n random players over 2 identical parallel links.

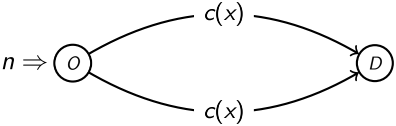

Symmetric mixed equilibrium: each player randomizes $(\frac{1}{2}, \frac{1}{2})$.

If each player is present with probability $p_i = d/n$, the random edge-loads are

$$X_e \sim \mathrm{Binomial}(n, \tfrac{d}{2n})$$

which for large n converges to a $\mathrm{Poisson}(\frac{d}{2})$.

What happens for other non-symmetric equilibria? What if players are not homogeneous? And with different costs? And more complex topologies?

Toolkit — Sums of Bernoullis ≈ Poisson

The total variation distance between two integer-valued random variables U, V is

$$d_{\mathrm{TV}}(U,V) = \frac{1}{2}\sum_{k\in\mathbb{N}} \big|\mathbb{P}(U=k) - \mathbb{P}(V=k)\big|.$$

Theorem (Barbour & Hall 1984, Borizov & Ruzankin 2002)

Let $S = X_1 + \ldots + X_n$ be a sum of independent Bernoullis with $\mathbb{P}(X_i = 1) \leq p$, and $X \sim \mathrm{Poisson}(x)$ with the same expectation $\mathbb{E}[X] = x = \mathbb{E}[S]$. Then

$$d_{TV}(S,X) \leq p.$$

Moreover, if $h : \mathbb{N} \to \mathbb{R}$ is such that $\mathbb{E}|\Delta^2 h(X)| \leq \nu$, then

$$|\mathbb{E}h(S) - \mathbb{E}h(X)| \leq \frac{x\nu}{2}\frac{p\, e^p}{(1-p)^2}.$$

Remark: $\Delta^2 h(x) \triangleq h(x+2) - 2h(x+1) + h(x)$.

Poisson Convergence for Vanishing Probabilities

Standing Assumption: $\mathbb{E}[X^2 c_e(1+X)] < \infty$ for all $e \in E$ and $X \sim \mathrm{Poisson}(x)$.

This is a mild condition. It holds for costs with polynomial or exponential growth. It fails for fast growing costs such as factorials $k!$ or bi-exponentials $\exp(\exp(k))$.

We introduce the expected cost functions $\tilde{c}_e : \mathbb{R}_+ \to \mathbb{R}_+$ defined by

$$\boxed{\tilde{c}_e(x) \triangleq \mathbb{E}[c_e(1+X)] = \sum_{k=0}^{\infty} c_e(1+k) e^{-x} \frac{x^k}{k!}.}$$

Poisson Convergence for Vanishing Probabilities

Let π^n be a sequence of Bayes-Nash equilibria for Bernoulli ARGs $\mathscr{G}(p^n)$ with

$$\begin{cases} a) & |N^n| \to \infty, \\ b) & \max_{i\in N^n} p_i^n \to 0, \\ c) & d_\kappa^n \triangleq \sum_{i:\kappa_i^n=\kappa} p_i^n \to d_\kappa \quad \text{for all } \kappa \in \mathscr{K}. \end{cases}$$

Theorem

1. *The expected flows $(y^n, x^n) = (\mathbb{E}Y^n, \mathbb{E}X^n)$ are bounded and each cluster point $(\tilde{y}, \tilde{x})$ is a Wardrop equilibrium with demands d_κ and costs $\tilde{c}_e(\cdot)$.*
2. *Along any convergent subsequence we have*
   - *the edge-loads X_e^n converge in total variation to $X_e \sim$ Poisson$(\tilde{x}_e)$,*
   - *the route-flows Y_r^n converge in total variation to $Y_r \sim$ Poisson$(\tilde{y}_r)$,*
   - *the Poisson limits Y_r are independent.*

Poisson convergence for vanishing probabilities

Corollary

If the costs $c_e(k)$ are non-decreasing and non-constant, then the $\tilde{c}_e(\cdot)$'s are strictly increasing, the edge-loads $\tilde{x}_e$ are the same in all Wardrop equilibria, and for every sequence π^n of Bayes-Nash equilibria we have

$$X_e^n \xrightarrow{\mathrm{TV}} X_e \sim Poisson(\tilde{x}_e).$$

Theorem

If $c_e(2) > c_e(1)$ for all $e \in E$ then there is a constant κ such that

$$d_{TV}(X_e^n, X_e) \leq \kappa\left(\sqrt{\max_{i \in N} p_i^n} + \sqrt{\|d^n - d\|_1}\right).$$

Convergence of PoA along sequences of ARGs

For an atomic routing game $\mathscr{G}$ we denote

$$C(\pi) = \mathbb{E}_\pi \left[\textstyle\sum_{e \in E} X_e \, c_e(X_e) \right] \qquad \text{(expected social cost)}$$

$$C_{opt}(\mathscr{G}) = \min_\pi C(\pi) \qquad \text{(minimum social cost)}$$

$$\mathrm{PoA}(\mathscr{G}) = \max_{\pi \in \mathscr{E}(\mathscr{G})} C(\pi) / C_{opt}(\mathscr{G}) \qquad \text{(price-of-anarchy)}$$

Theorem

Under the same conditions of the convergence theorems for weighted and Bernoulli ARGs, we have

$$\mathrm{PoA}(\mathscr{G}(w^n)) \longrightarrow \mathrm{PoA}(\mathit{Wardrop})$$

$$\mathrm{PoA}(\mathscr{G}(p^n)) \longrightarrow \mathrm{PoA}(\mathit{Poisson})$$

Summary and Comments

1. Both $w_i^n \to 0$ and $p_i^n \to 0$ lead to different non-atomic limit games:
   - For vanishing weights, the random edge-loads X_e^n converge in L^2 to the constants edge-loads $\hat{x}_e$.
   - For vanishing probabilities, X_e^n remain random in the limit and converge in total variation to $X_e \sim \text{Poisson}(\tilde{x}_e)$.
2. The Poisson limit is consistent with empirical data on traffic counts. Also $p_i^n \to 0$ is quite natural... congestion depends on players that are present on a small window around your departure time.
3. The Poisson limit is a special case of Myerson's Poisson games: the normalized loads $\sigma(r|t) = y_r/d_\kappa$ for $r \in \mathscr{R}_\kappa$ yield an equilibrium in the Poisson game (Myerson, Int J Game Theory 1998).
4. Poisson games were defined without reference to a limit process, so our convergence result as well as the connection with Wardrop seem new.

Price-of-Anarchy for Atomic Routing Games

PoA for Atomic Routing – Smoothness Framework

Consider an unweighted ARG and let the social cost be

$$C(r) = \sum_{i \in N} C_i(r) = \sum_{e \in E} n_e \, c_e(n_e).$$

Definition

The game is called (λ, μ)-smooth with $\lambda > 0$ and $0 < \mu < 1$ if for all strategy profiles $r = (r_i)_{i \in N}$ and $r' = (r_i')_{i \in N}$ we have

$$\textstyle\sum_{i \in N} C_i(r_i', r_{-i}) \leq \lambda \, C(r') + \mu \, C(r).$$

Lemma (Dumrauf-Gairing, 2006; Harks-Vegh, 2007; Roughgarden, 2015)

For (λ, μ)-smooth atomic routing games we have $\mathrm{PoA} \leq \frac{\lambda}{1-\mu}$.

PoA for Atomic Routing – Smoothness Framework

Consider an unweighted ARG and let the social cost be

$$C(r) = \sum_{i\in N} C_i(r) = \sum_{e\in E} n_e\, c_e(n_e).$$

Definition

The game is called (λ,μ)-smooth with $\lambda > 0$ and $0 < \mu < 1$ if for all strategy profiles $r = (r_i)_{i\in N}$ and $r' = (r_i{}')_{i\in N}$ we have

$$\textstyle\sum_{i\in N} C_i(r_i', r_{-i}) \leq \lambda\, C(r') + \mu\, C(r).$$

Lemma (Dumrauf-Gairing, 2006; Harks-Vegh, 2007; Roughgarden, 2015)

For (λ,μ)-smooth atomic routing games we have $\mathrm{PoA} \leq \frac{\lambda}{1-\mu}$.

Proof. For every Nash equilibrium r we have

$$C(r) = \textstyle\sum_{i\in N} C_i(r_i, r_{-i}) \leq \sum_{i\in N} C_i(r_i', r_{-i}) \leq \lambda\, C(r') + \mu\, C(r).$$

hence $C(r) \leq \frac{\lambda}{1-\mu} C(r')$ and we conclude by minimizing over r'. □

PoA for Atomic Routing – Smoothness Framework

Translated in terms of arc costs (λ, μ)-smoothness becomes

$$k\, c_e(m+1) \leq \lambda\, k\, c_e(k) + \mu\, m\, c_e(m) \qquad \forall k, m \in \mathbb{N}$$

Theorem (Christodoulou-Koutsoupias, 2005; Awerbuch-Azar-Epstein, 2005**)**

Atomic routing games with affine costs are $(\frac{5}{3}, \frac{1}{3})$-smooth and $\mathrm{PoA} \leq \frac{5}{2}$.

Almost twice larger than the bound $\mathrm{PoA} \leq \frac{4}{3}$ for non-atomic routing.

PoA for Atomic Routing – Smoothness Framework

Translated in terms of arc costs (λ, μ)-smoothness becomes

$$k\, c_e(m+1) \leq \lambda\, k\, c_e(k) + \mu\, m\, c_e(m) \qquad \forall k, m \in \mathbb{N}$$

Theorem (Christodoulou-Koutsoupias, 2005; Awerbuch-Azar-Epstein, 2005**)**

Atomic routing games with affine costs are $(\frac{5}{3}, \frac{1}{3})$*-smooth and* $\mathrm{PoA} \leq \frac{5}{2}$.

Almost twice larger than the bound $\mathrm{PoA} \leq \frac{4}{3}$ for non-atomic routing.

What happens in Bernoulli ARGs as we move from the deterministic case $p_i \equiv 1$ to the Wardrop limit when $p_i \downarrow 0$?

PoA for Bernoulli ARGs

Consider a Bernoulli ARG with $p_i = \mathbb{P}(W_i = 1)$ and social cost

$$C(\pi) = \sum_{i \in N} C_i(\pi) = \mathbb{E}\left[\sum_{e \in E} X_e\, c_e(X_e)\right].$$

Theorem

Let $\mathscr{G}^p$ denote the set of Bernoulli ARGs with $p_i \leq p$ for all players. The largest values of $\mathrm{PoA}(\mathscr{G}(p))$ occur for homogeneous players with $p_i \equiv p$.

Proposition

A Bernoulli ARG with homogeneous players $p_i \equiv p$ is equivalent to a deterministic unweighted ARG for the auxiliary costs

$$c_e^p(k) = \mathbb{E}[c_e(1+B)] \text{ with } B \sim \mathrm{Binomial}(k-1, p)$$

$\Rightarrow$ any (λ, μ)-bound for this equivalent deterministic game remains valid for non-homogeneous players with $p_i \leq p$.

PoA for Bernoulli ARGs — Homogeneous players

From now on we focus on the homogeneous case and study PoA and PoS as a function of p when we move from the deterministic case $p = 1$ to the limit $p \downarrow 0$.

$$\begin{aligned} \mathrm{PoA}(p) &= \sup_{\mathscr{G}^p} \max_{\pi \in \mathscr{E}(\mathscr{G}^p)} C(\pi)/C_{opt}(\mathscr{G}^p) \quad \text{(Price-of-Anarchy)} \\ \mathrm{PoS}(p) &= \sup_{\mathscr{G}^p} \min_{\pi \in \mathscr{E}(\mathscr{G}^p)} C(\pi)/C_{opt}(\mathscr{G}^p) \quad \text{(Price-of-Stability)} \end{aligned}$$

We search for the best (λ, μ)-smoothness parameters

$$k\, c_e^p(m+1) \leq \lambda\, k\, c_e^p(k) + \mu\, m\, c_e^p(m) \qquad \forall k, m \in \mathbb{N}.$$

For the special case of affine costs we get $\mathrm{PoA} \leq \frac{5}{2}$.
But we expect sharper bounds with $\mathrm{PoA} \approx \frac{4}{3}$ for small p.

PoA for Bernoulli ARGs – Affine Costs

For affine costs $c_e(x) = \alpha_e x + \beta_e$ with $\alpha_e, \beta_e \geq 0$ we have

$$\begin{aligned} c_e^p(k) &= \mathbb{E}[c_e(1 + B(k-1, p))] \\ &= \alpha_e(1 + (k-1)p) + \beta_e \end{aligned}$$

and (λ, μ)-smoothness reduces to

$$k\,(1+p\,m) \leq \lambda\,k\,(1-p+p\,k) + \mu\,m\,(1-p+p\,m) \qquad \forall k, m \in \mathbb{N}. \tag{1}$$

The best combination of λ and μ for fixed p requires to solve

$$\mathrm{PoA} \leq B(p) \triangleq \min_{\lambda>0,\mu\in(0,1)} \left\{ \frac{\lambda}{1-\mu} : \text{ subject to (1)} \right\}$$

which reduces to a 1D problem noting that the smallest λ compatible with (1) is

$$\lambda = \sup_{(k,m)\in\mathscr{P}} \frac{k(1+pm)-\mu\,m(1-p+pm)}{k(1-p+pk)}$$

PoA for Bernoulli ARGs – Affine Costs

The previous reduction leads to the equivalent minimization problem

$$B(p) = \inf_{\mu \in (0,1)} \varphi_p\left(\tfrac{\mu}{1-\mu}\right) = \inf_{y>0} \varphi_p(y)$$

where $\varphi_p(\cdot)$ is the convex envelop function

$$\varphi_p(y) = \sup_{(k,m)\in \mathscr{P}} \tfrac{1+pm}{1-p+pk} + \tfrac{k(1+pm)-m(1-p+pm)}{k(1-p+pk)}\, y.$$

For each p the unique optimum y can be found explicitly, and then we recover the optimal combination (λ, μ).

This reveals some unexpected phase transitions !

PoA for Bernoulli ARGs – Affine Costs

Theorem (C-Scarsini-Schröder-Stier, 2019)

Let $\bar{p}_0 = \frac{1}{4}$ and $\bar{p}_1 \sim 0.3774$ the unique real root of $8p^3 + 4p^2 = 1$. For Bernoulli ARGs with affine costs and $p_i \leq p$ we have

$$\mathrm{PoA} \leq B(p) = \begin{cases} \frac{4}{3} & \text{if } 0 < p \leq \bar{p}_0 \\ \frac{1+p+\sqrt{p(2+p)}}{1-p+\sqrt{p(2+p)}} & \text{if } \bar{p}_0 \leq p \leq \bar{p}_1 \\ 1 + p + \frac{p^2}{1+p} & \text{if } \bar{p}_1 \leq p \leq 1 \end{cases}$$

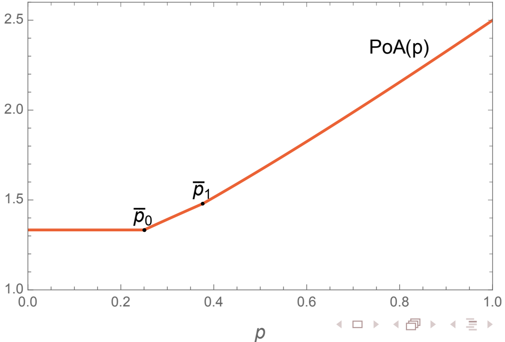

Tight lower bounds for large p

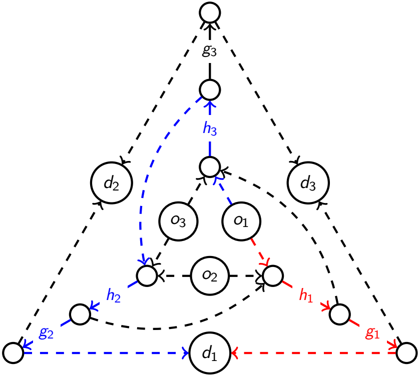

$$c_e(x) = \begin{cases} x & \text{if } e = h_i \\ p\,x & \text{if } e = g_i \\ 0 & \text{if dashed} \end{cases}$$

$$\Rightarrow \mathrm{PoA}(\mathscr{G}^p) = 1 + p + \frac{p^2}{1+p}.$$

Tight lower bounds for small p

$$c_e(x) = \begin{cases} \frac{1}{1+2kp}\, x & \text{if } e = \bar{e} \\ x & \text{if } e = e_i \\ 0 & \text{if dashed} \end{cases}$$

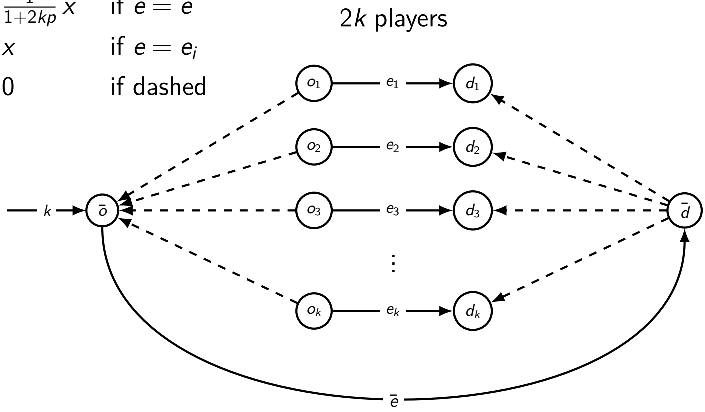

$$\Rightarrow \mathrm{PoA}(\mathscr{G}^p) = \mathrm{PoS}(\mathscr{G}^p) \geq \frac{4kp+2-2p}{3kp+2-p} \to \frac{4}{3} \text{ as } k \to \infty$$

Tight lower bounds for intermediate p

$$c_e(x) = \begin{cases} \alpha\, x & \text{if } e = h_i \\ x & \text{if } e = g_i \\ 0 & \text{if dashed} \end{cases}$$

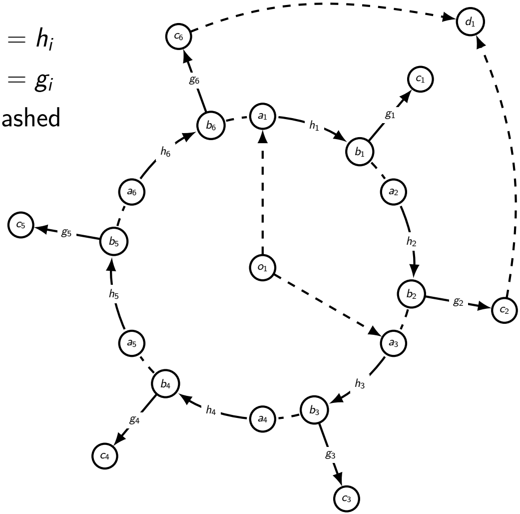

Bounds on the Price-of-Anarchy are Tight

$$\mathrm{PoA}(p) = B(p) = \begin{cases} \frac{4}{3} & \text{if } 0 < p \leq \bar{p}_0 \\ \frac{1+p+\sqrt{p(2+p)}}{1-p+\sqrt{p(2+p)}} & \text{if } \bar{p}_0 \leq p \leq \bar{p}_1 \\ 1+p+\frac{p^2}{1+p} & \text{if } \bar{p}_1 \leq p \leq 1 \end{cases}$$

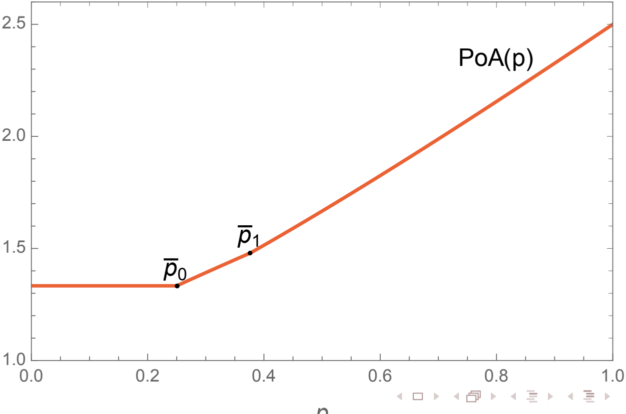

Price-of-Anarchy vs Price-of-Stability

Combined with (Kleer-Schäfer, 2018) we also get tight bounds for PoS

$$\mathrm{PoS}(p) = \begin{cases} \frac{4}{3} & \text{if } 0 < p \leq \bar{p}_0 \\ 1 + \sqrt{p/(2+p)} & \text{if } p \geq \bar{p}_0 \end{cases}$$

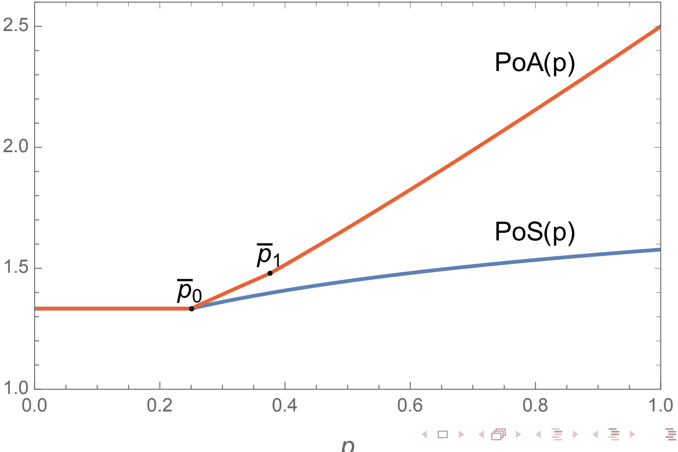

Conclusion

1. Convergence of ARGs towards non-atomic games:
    - vanishing weights $\longrightarrow$ Wardrop
    - vanishing probabilities $\longrightarrow$ Poisson/Wardrop
2. Convergence of PoA/PoS, plus sharp bounds for affine costs
3. Some open questions
    - Mixed limits: weights & probabilities
    - Tight bounds for polynomial costs

Questions ?