London Business School

Spatial Matching under Multihoming

Alireza Amanihamedani (LBS)

Joint work with
Ali Aouad (LBS)
Daniel Freund (MIT)

Online matching: myopic vs. forward-looking

Greedy can be sub-optimal to minimize cost

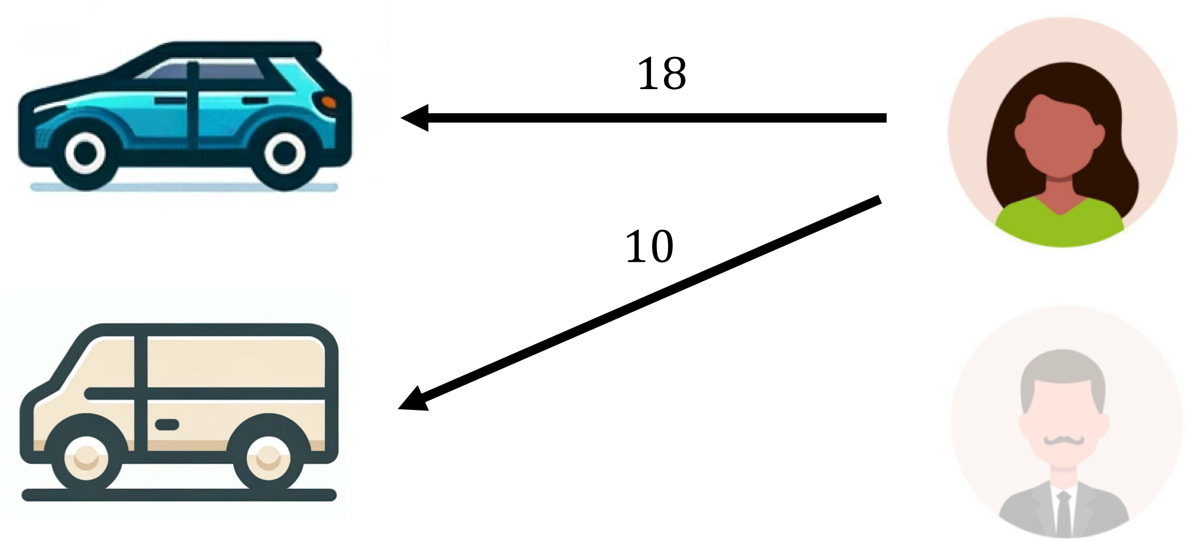

Online matching: myopic vs. forward-looking

Greedy can be sub-optimal to minimize cost

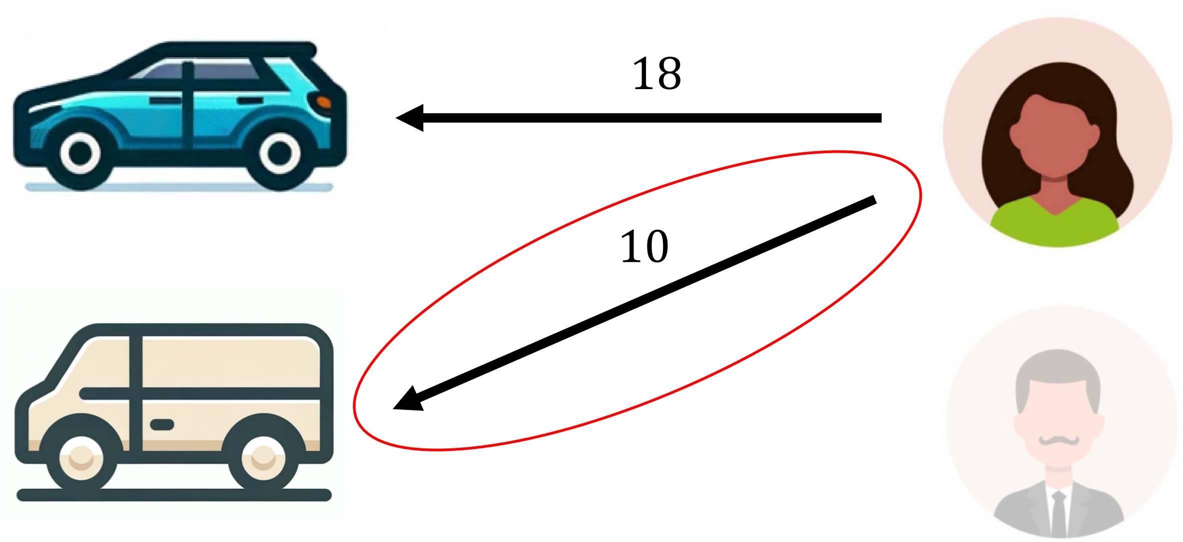

Online matching: myopic vs. forward-looking

Greedy can be sub-optimal to minimize cost

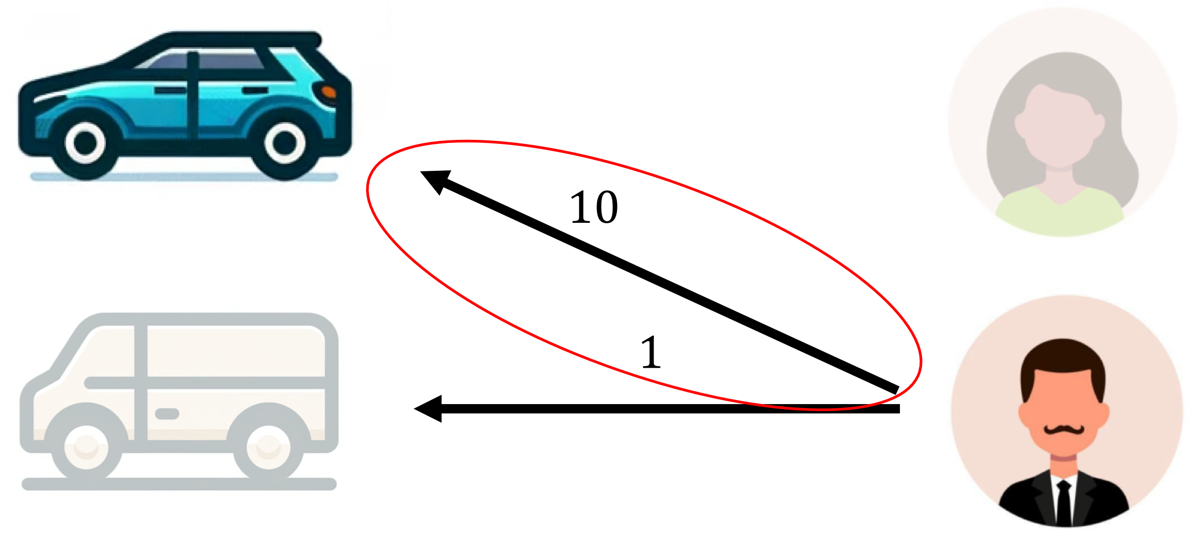

Online matching: myopic vs. forward-looking

Greedy can be sub-optimal to minimize cost

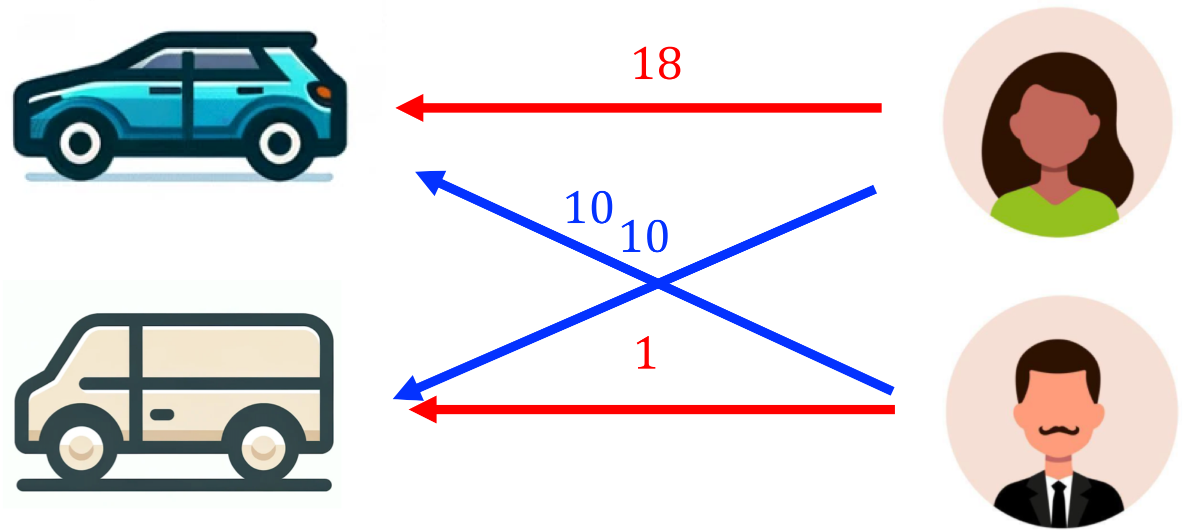

$$Greedy\ cost = 10 + 10 = \ 20$$

Online matching: myopic vs. forward-looking

Greedy can be sub-optimal to minimize cost

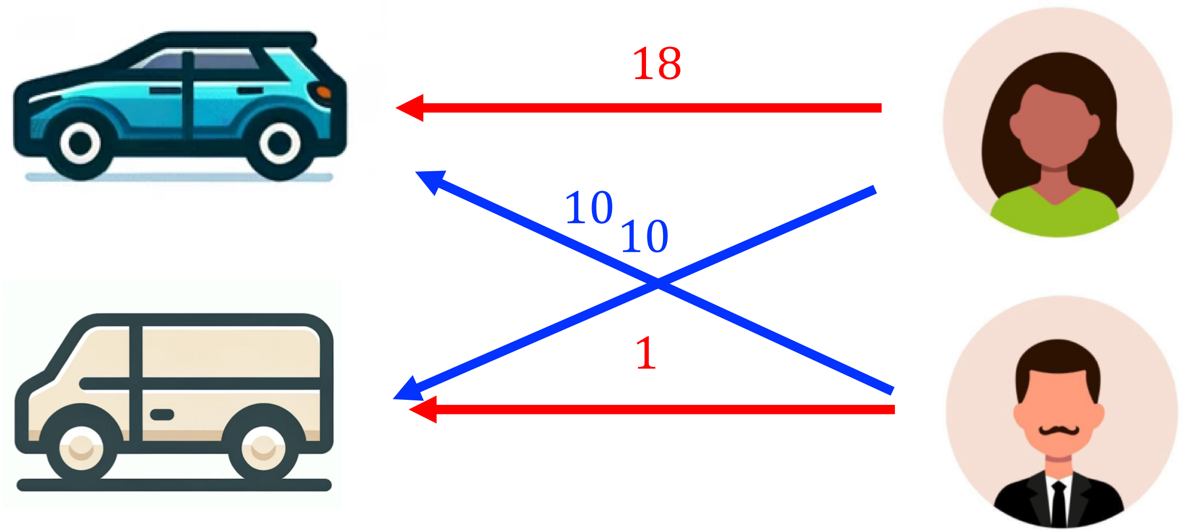

$Greedy\ cost = 10 + 10 = 20$
$\text{“}Forward\ looking\text{”}\ cost = 18 + 1 = 19$

Effect of competition

Competition decreases the value of being forward looking

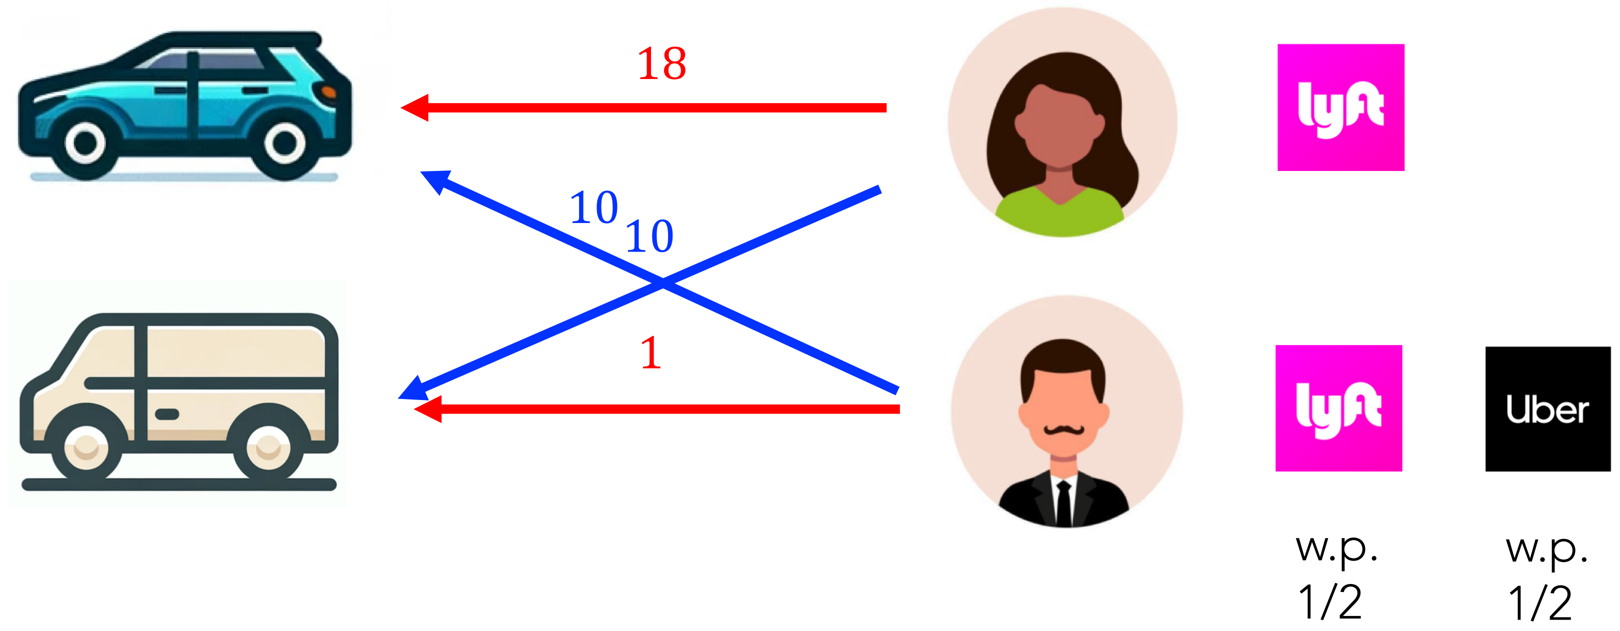

Effect of competition

Competition decreases the value of being forward looking

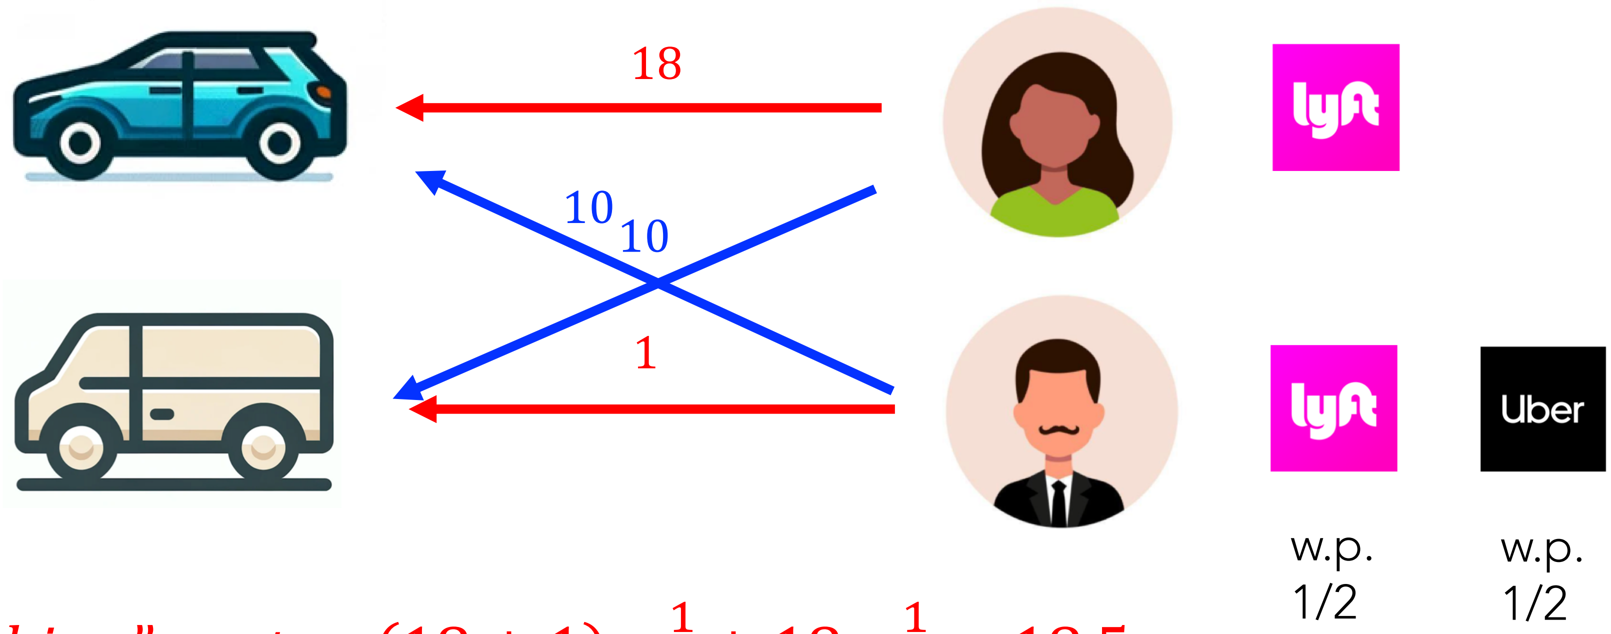

$$\text{“Forward looking” cost} = (18 + 1) * \frac{1}{2} + 18 * \frac{1}{2} = 18.5$$

Effect of competition

Being greedy can be better since the resources may be "stolen"

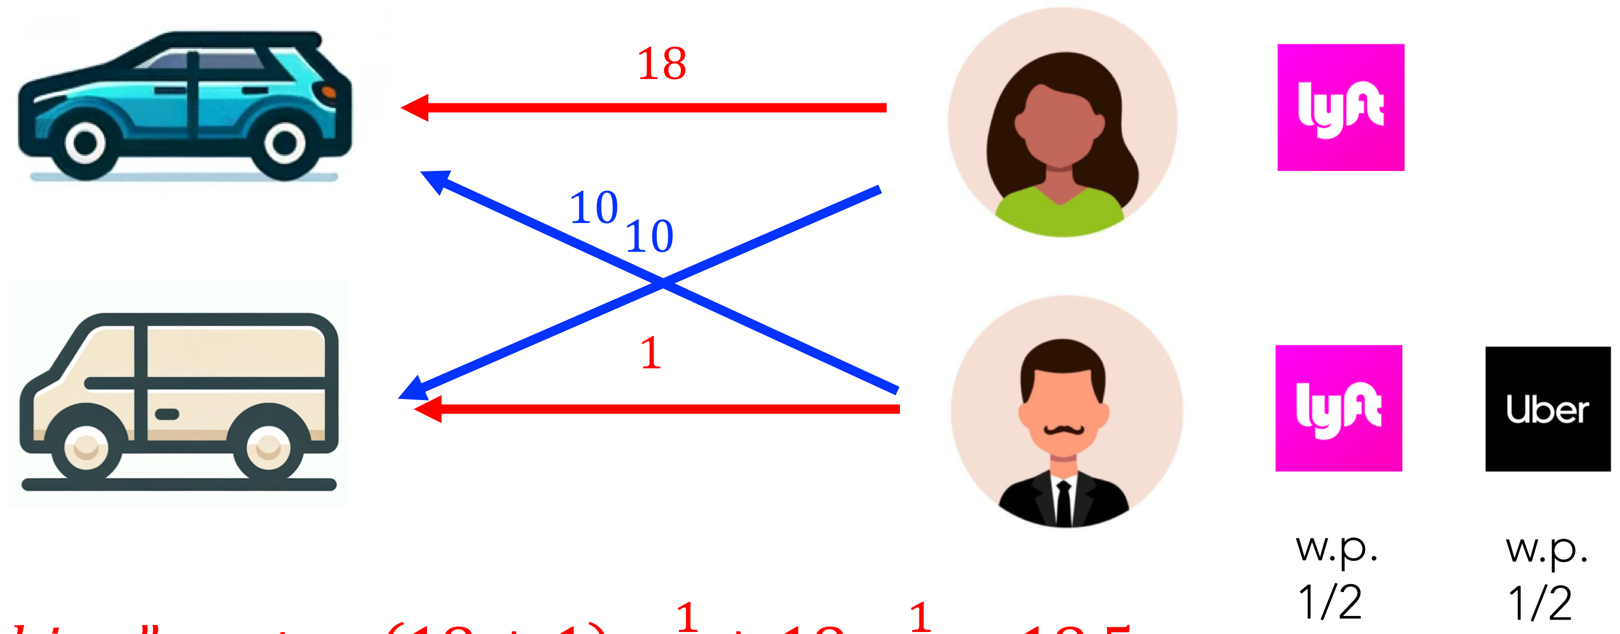

$$\text{"Forward looking" cost} = (18+1)*\frac{1}{2} + 18*\frac{1}{2} = 18.5$$

$$\text{Greedy cost} = (10+10)*\frac{1}{2} + 10*\frac{1}{2} = 15$$

Spatial matching & pooling

- Focus on spatial matching services (e.g., ride-hailing, delivery)

  more available suppliers ⇒ better dispatches

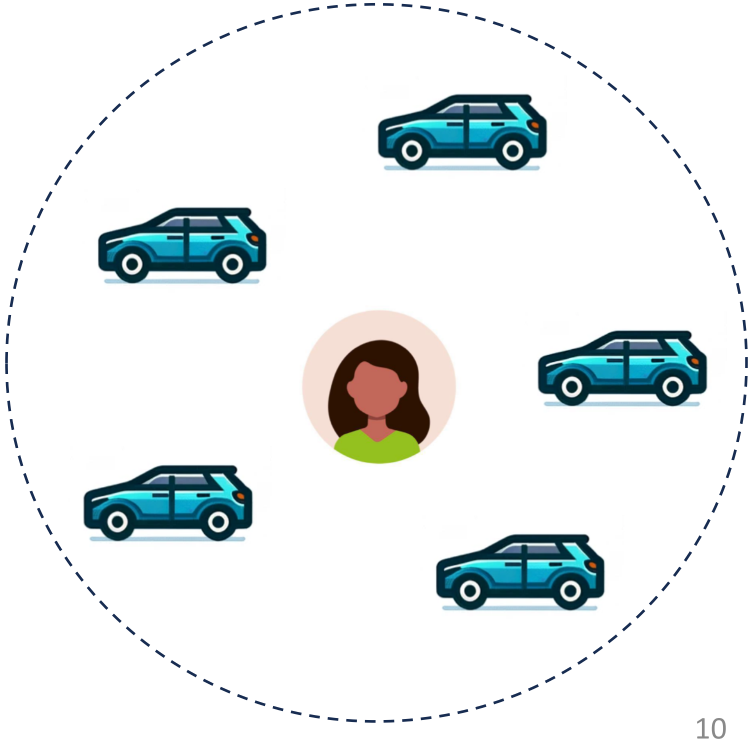

Spatial matching & pooling

- Focus on spatial matching services (e.g., ride-hailing, delivery)

  more available suppliers ⇒ better dispatches

- Forward-looking policies: keep a large buffer of idle supply (“spatial pooling”), Castillo et al. [‘17]

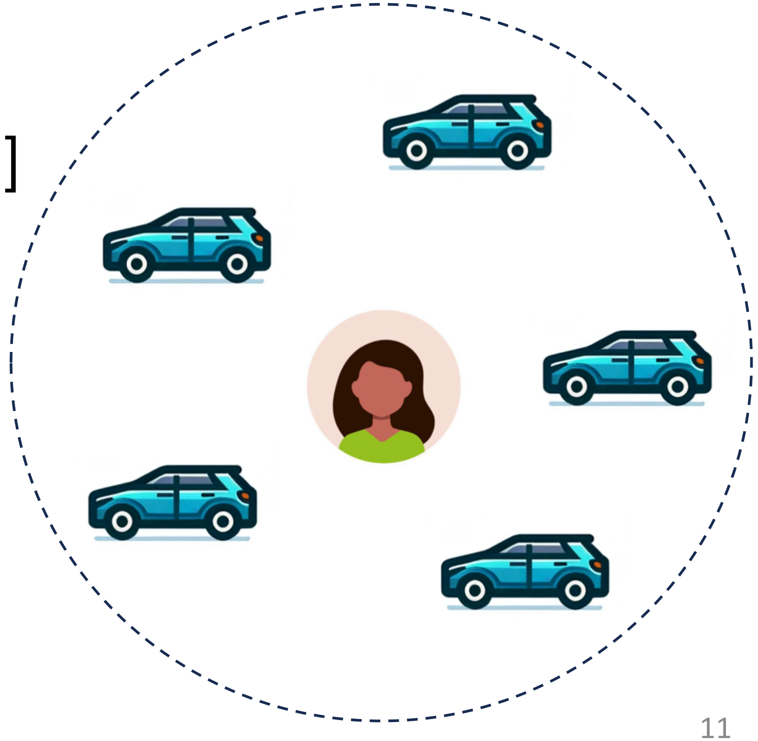

Spatial matching & pooling

- Focus on spatial matching services (e.g., ride-hailing, delivery)

  more available suppliers ⇒ better dispatches

- Forward-looking policies: keep a large buffer of idle supply (“spatial pooling”), Castillo et al. [‘17]

- Multihoming and competing platforms?

  tension between

  spatial pooling
  race to the bottom

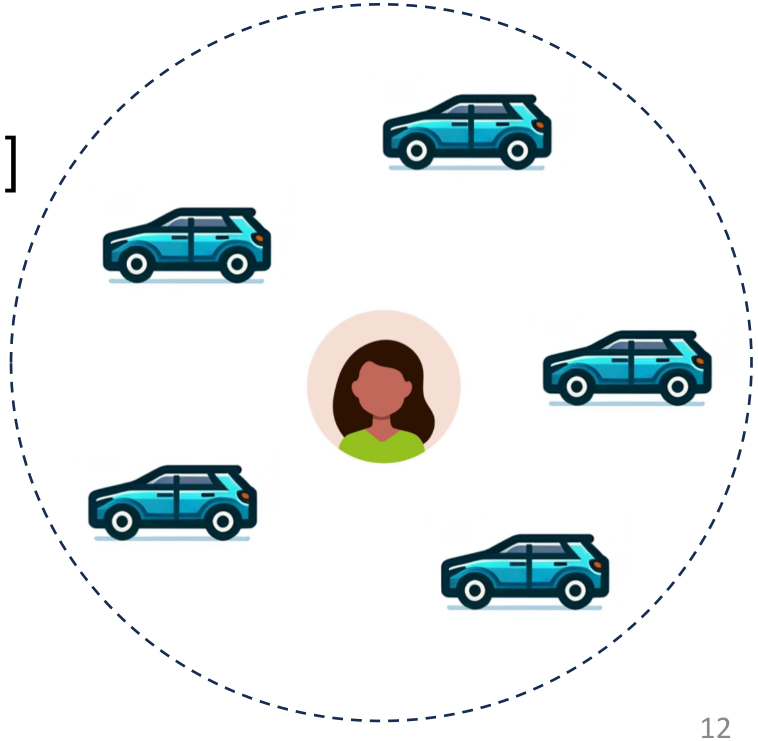

Spatial matching & pooling

- Focus on spatial matching services (e.g., ride-hailing, delivery)

  more available suppliers ⇒ better dispatches

- Forward-looking policies: keep a large buffer of idle supply ("spatial pooling"), Castillo et al. ['17]

- Multihoming and competing platforms?

Does multihoming in a duopoly lead to inefficient matching policies in equilibrium?

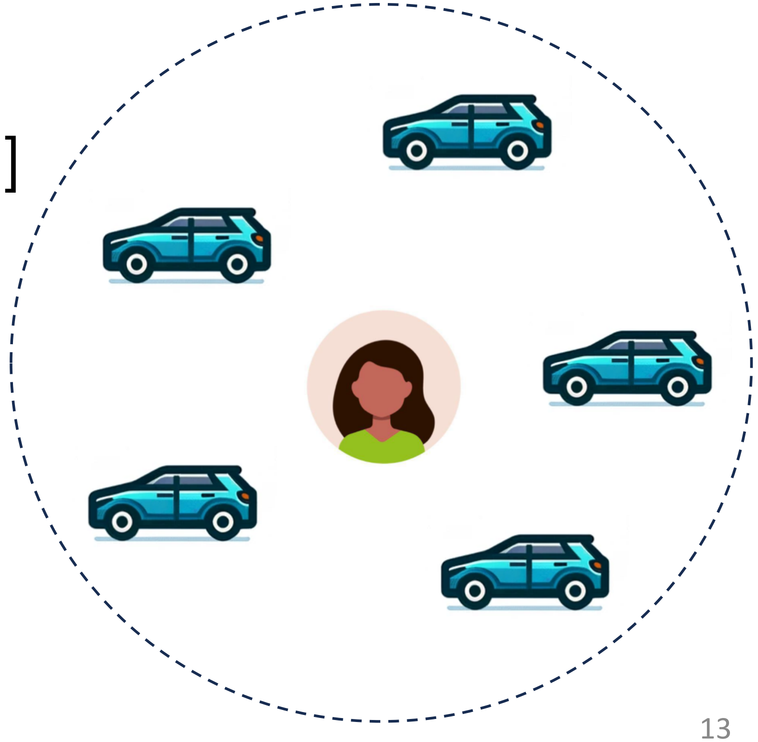

Outline

- Modeling approach:
  - Duopoly with dispatch policies in a stochastic system
  - Trade-offs between different cost types
- Equilibrium analysis: scale-efficient and scale-inefficient regimes
- Characterization of market efficiency and insights

Related literature

- Staffing & capacity planning:
  - Halfin & Whitt [’81], Ward [‘12], Atar [‘12]
  - Spatial capacity planning, Besbes et al. [‘21]
- Matching and pricing in two-sided platforms:
  - Wild Goose Chase, Castillo et al. [‘17]
  - Optimal control: Banerjee et al. [‘16], Feng et al. [‘20], Freund & van Ryzin [‘21], Kanoria [‘22], Akbarpour et al. [‘21], Aouad & Saritac [‘22]
- Competition in gig economy platforms:
  - Competition via pricing, Ahmadinejad et al. [‘19]
  - Labor participation and workers’ earnings, Lian et al. [‘21]

A queuing duopoly model

- Suppliers: Poisson process with rate Λ, shared between platforms, located randomly in a 1D ball

- Customers: two disjoint Poisson processes with rates $\lambda_1 \cdot \Lambda$ and $\lambda_2 \cdot \Lambda$

$$\lambda_1 + \lambda_2 > 1$$

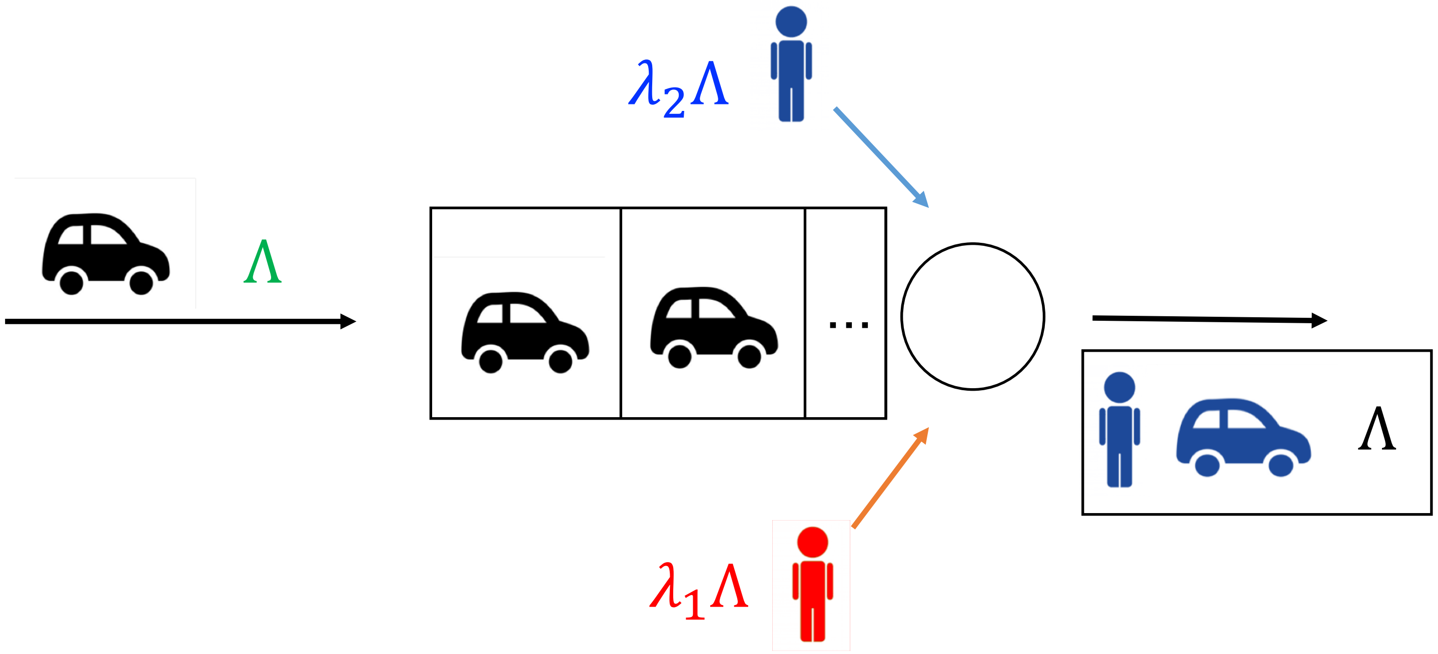

A queuing duopoly model

- Suppliers: Poisson process with rate Λ, shared between platforms, located randomly in a 1D ball

- Customers: two disjoint Poisson processes with rates $\lambda_1 \cdot \Lambda$ and $\lambda_2 \cdot \Lambda$

- Example:
  - UberX vs. Lyft
  - UberX vs. DoorDash

A queuing duopoly model

- Suppliers: Poisson process with rate Λ, shared between platforms, located randomly in a 1D ball

- Customers: two disjoint Poisson processes with rates $\lambda_1 \cdot \Lambda$ and $\lambda_2 \cdot \Lambda$

- Example:
  - UberX vs. Lyft
  - UberX vs. DoorDash

- Large-market limit: $\Lambda \to \infty$

Cost-minimization game

1. Dispatch cost: $c_D \times \mathbb{E}\left[\text{dispatch distance}\right] \times (\text{rate of fulfilled demand})$

2. Idle cost: $c_I \times \mathbb{E}\left[\text{number of idle suppliers}\right] \times (\text{market share})$

3. Unfulfillment cost: $(\text{rate of rider requests that are not served})$

$c_D \in \mathbb{N}, c_I \geq 0$

Admission control policies

- Threshold on critical #idle suppliers

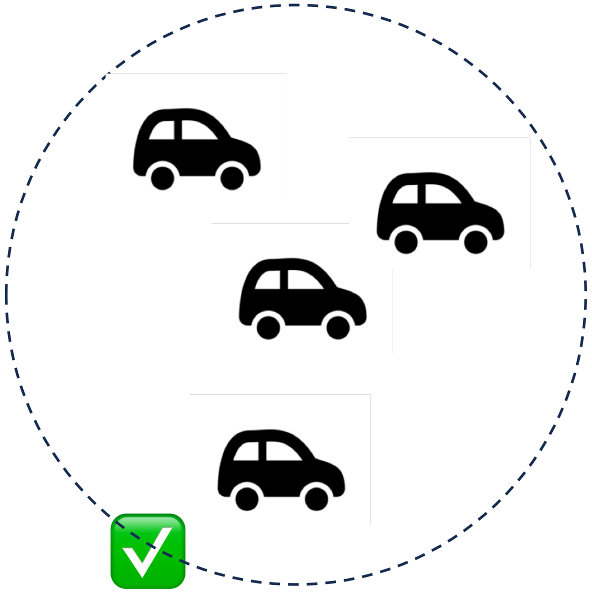

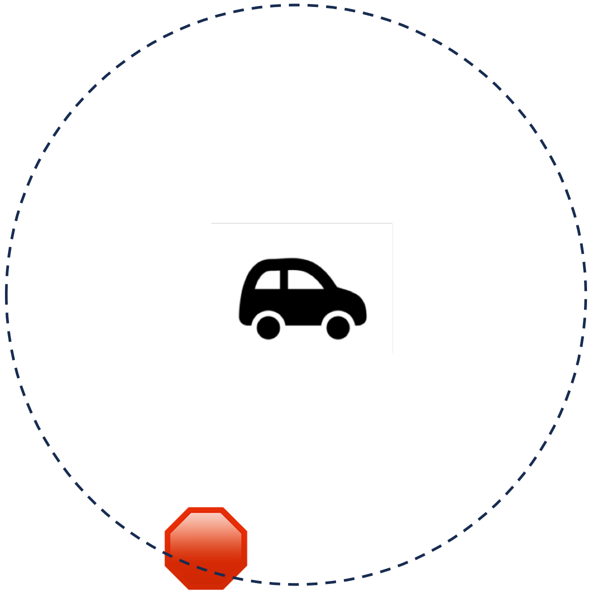

Admission control policies

- Threshold on critical #idle suppliers: (1, $n + 1$)

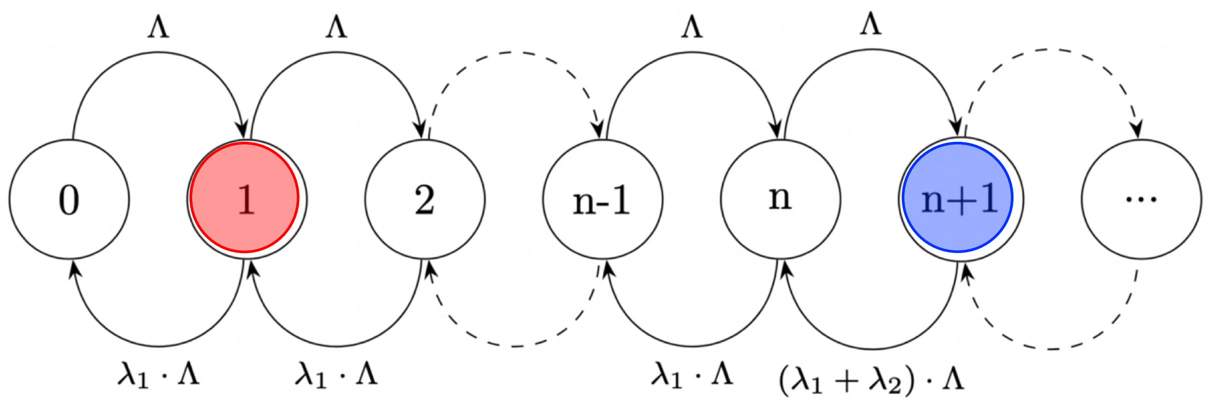

Cost-minimization game

$$C_j(n_j, n_{-j}) \triangleq \underbrace{\Pr[N(n_1,n_2) < n_j]}_{\text{UC}} + \underbrace{\mathbb{E}\left[\frac{c_D}{N(n_1,n_2)+1} \cdot \mathbb{I}[N(n_1,n_2) \geq n_j]\right]}_{\text{DC}}$$

$$+ \underbrace{c_I \cdot \frac{1}{\Lambda} \cdot \Pr[N(n_1,n_2) \geq n_j] \cdot \mathbb{E}[N(n_1,n_2)]}_{\text{IC}}.$$

Monopolist setting

Proposition [Folklore]:
The monopolist chooses an optimal threshold of $\Theta(\sqrt{\Lambda})$ (economies of scale)

Simple idea:

Unfulfillment cost constant (Little's law)
$n^*(\Lambda) = \Theta(\sqrt{\Lambda})$ to balance the idle cost ($\approx n/\Lambda$) and dispatch cost ($\approx 1/n$)

Equilibrium notion

- Equilibrium notion:

Fixing the opponent threshold, no platform reduces cost by deviating

ε-equilibrium:

Fixing the opponent threshold, no platform reduces cost too much by deviating

Equilibrium notion

- Equilibrium notion:

$$C_j^I(n_j, n_{-j}) \leq C_j^I(m, n_{-j}) \;\; \forall m \in \mathbb{N}$$

ε-equilibrium:

$$C_j^I(n_j, n_{-j}) \leq C_j^I(m, n_{-j}) + \varepsilon \;\; \forall m \in \mathbb{N}$$

Equilibrium analysis: definitions

With instance $I = (c_D, c_I, \lambda_1, \lambda_2, 1)$,

- (n_1, n_2) is a scale-inefficient equilibrium if an equilibrium for any large enough Λ
- I is a scale-inefficient instance if all equilibria are scale-inefficient equilibria
- $\left(n_1(\Lambda), n_2(\Lambda)\right)$ is a family of scale-efficient ε-equilibria for any large enough Λ and $\max\{n_1(\Lambda), n_2(\Lambda)\} = n^*(\Lambda)$

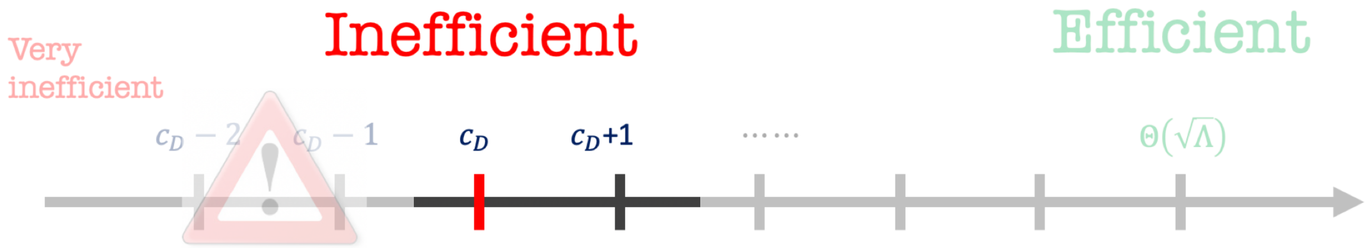

Equilibrium analysis: definitions

With instance $I = (c_D, c_I, \lambda_1, \lambda_2, 1)$,

- (n_1, n_2) is a scale-inefficient equilibrium if an equilibrium for any large enough Λ
- I is a scale-inefficient instance if all equilibria are scale-inefficient equilibria
- $\left(n_1(\Lambda), n_2(\Lambda)\right)$ is a family of scale-efficient ε-equilibria for any large enough Λ and $\max\{n_1(\Lambda), n_2(\Lambda)\} = n^*(\Lambda)$

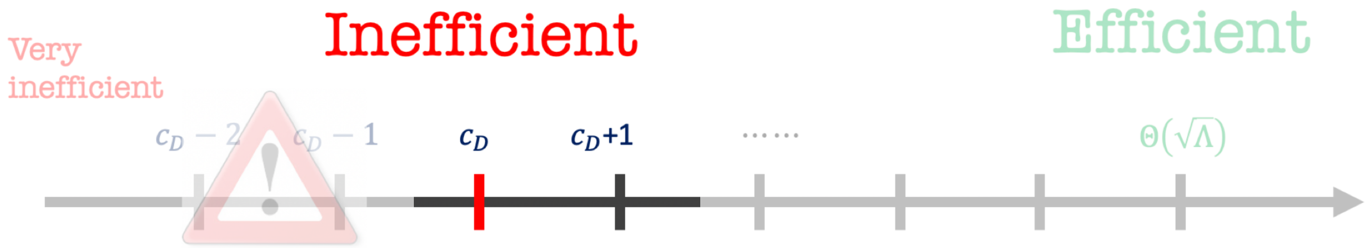

Equilibrium analysis: definitions

With instance $I = (c_D, c_I, \lambda_1, \lambda_2, 1)$,

- (n_1, n_2) is a scale-inefficient equilibrium if an equilibrium for any large enough Λ
- I is a scale-inefficient instance if all equilibria are scale-inefficient equilibria
- $(n_1(\Lambda), n_2(\Lambda))$ is a family of scale-efficient ε-equilibria for any large enough Λ and $\max\{n_1(\Lambda), n_2(\Lambda)\} = n^*(\Lambda)$

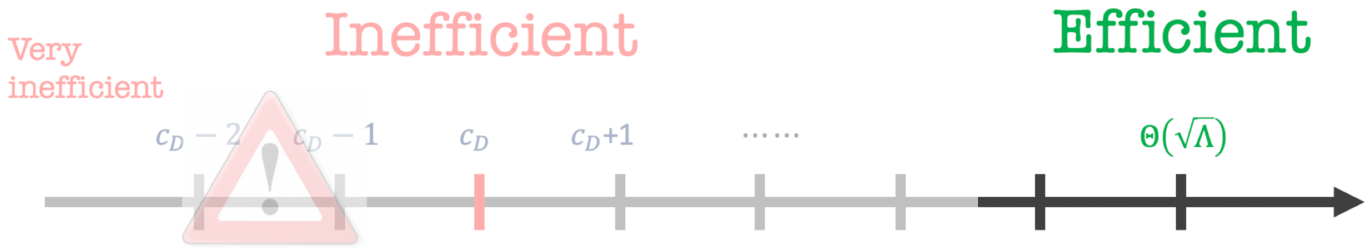

Equilibrium analysis

Informal Theorem 1 [A., Aouad, Freund, ’23]:
Any instance can be classified into two mutually disjunctive outcomes:

1. Scale-inefficient instance and equilibrium (c_D, c_D) with no efficiency of scale (Λ)

2. Scale-efficient ε-equilibria of the form $(c_D, n^*(\Lambda))$,
   where one platform generates efficiencies of scale

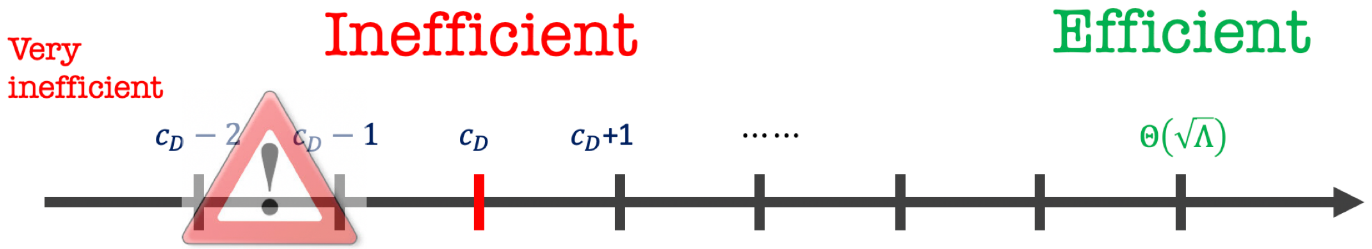

Equilibrium classification

Equilibrium classifier:

For $\lambda_1 \leq \lambda_2$, *let* $g \triangleq \lambda_2 - \lambda_1 \cdot \left(\sum_{i=1}^{+\infty} \frac{c_D}{c_D + i} \cdot \left(\frac{1}{\lambda_1 + \lambda_2} \right)^i \right)$.

Equilibrium classification

Theorem 1 [A., Aouad, Freund, '23]:

1. If $g > 0$

   I is scale-inefficient with equilibrium (c_D, c_D)

   no scale-efficient ε-equilibria with $n_1(\Lambda), n_2(\Lambda) \geq c_D$

2. If $g < 0$

   (c_D, c_D) is not a scale-inefficient equilibrium

   Scale-efficient ε-equilibria $(c_D, n^*(\Lambda))$

   🌟 If $\lambda_1 < 1/(c_D + 1)$, scale-efficient equilibria $(c_D, n^*(\Lambda) \pm 1)$

3. If $g = 0$

   (c_D, c_D) is a scale-inefficient equilibrium

   Scale-efficient ε-equilibria $(c_D, n^*(\Lambda))$

Equilibrium classification

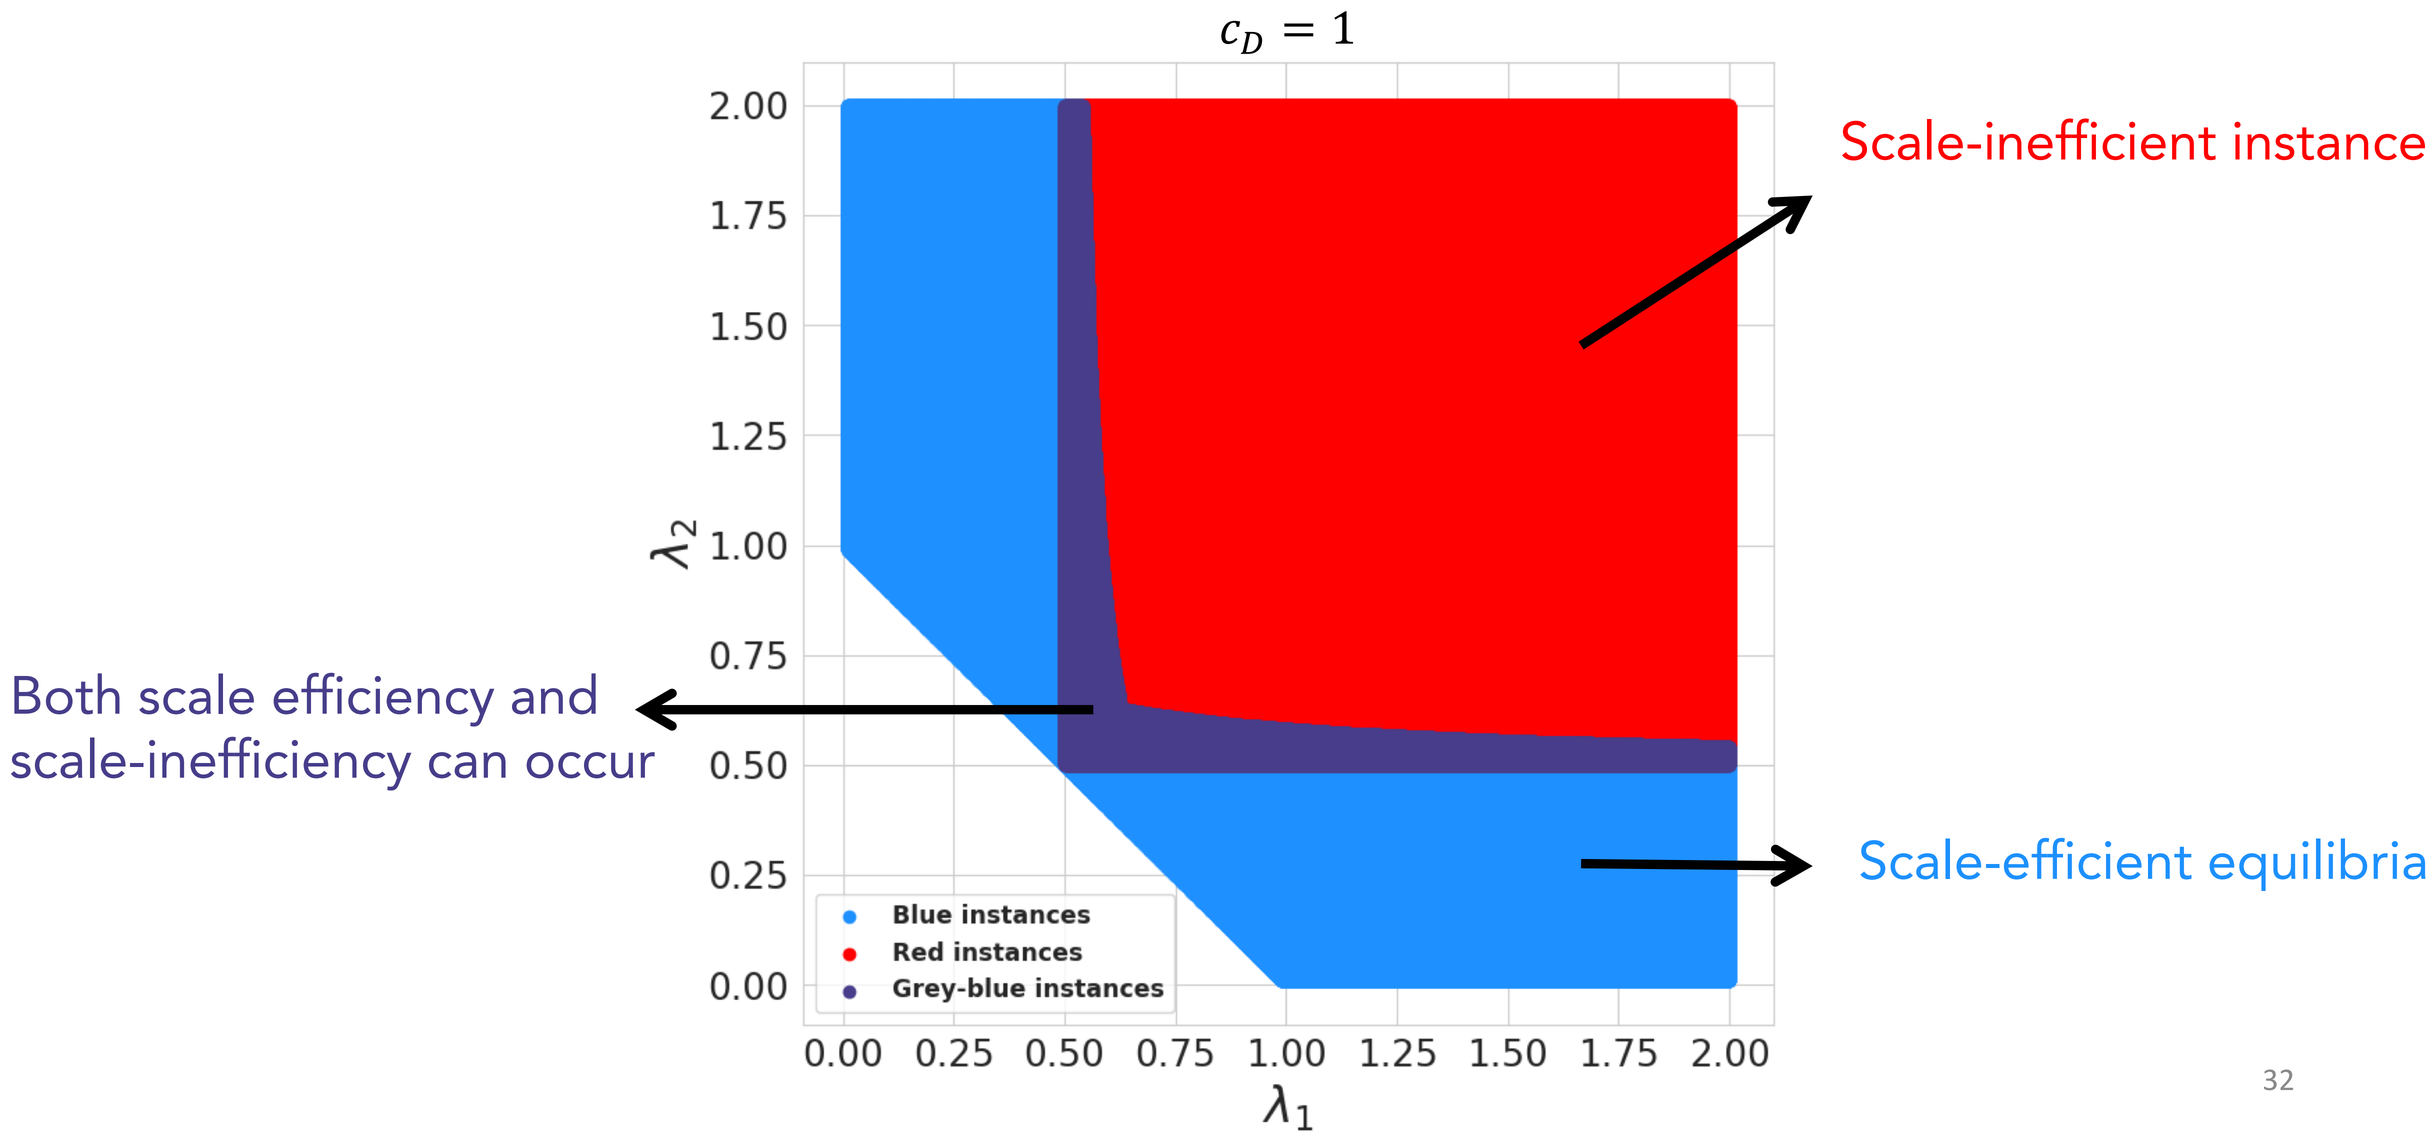

Equilibrium classification: main lemma

Lemma:
For a fixed $p \geq c_D$, the best response to p satisfies exactly one of these:
(i) is in $[c_D, p]$ for every large enough Λ
(ii) is in $[n^*(\Lambda) - 1, n^*(\Lambda) + 1]$ for every large enough Λ

$\Rightarrow$ best response is either a (smaller) constant or close to monopolist optimum

Equilibrium classification: proof challenges

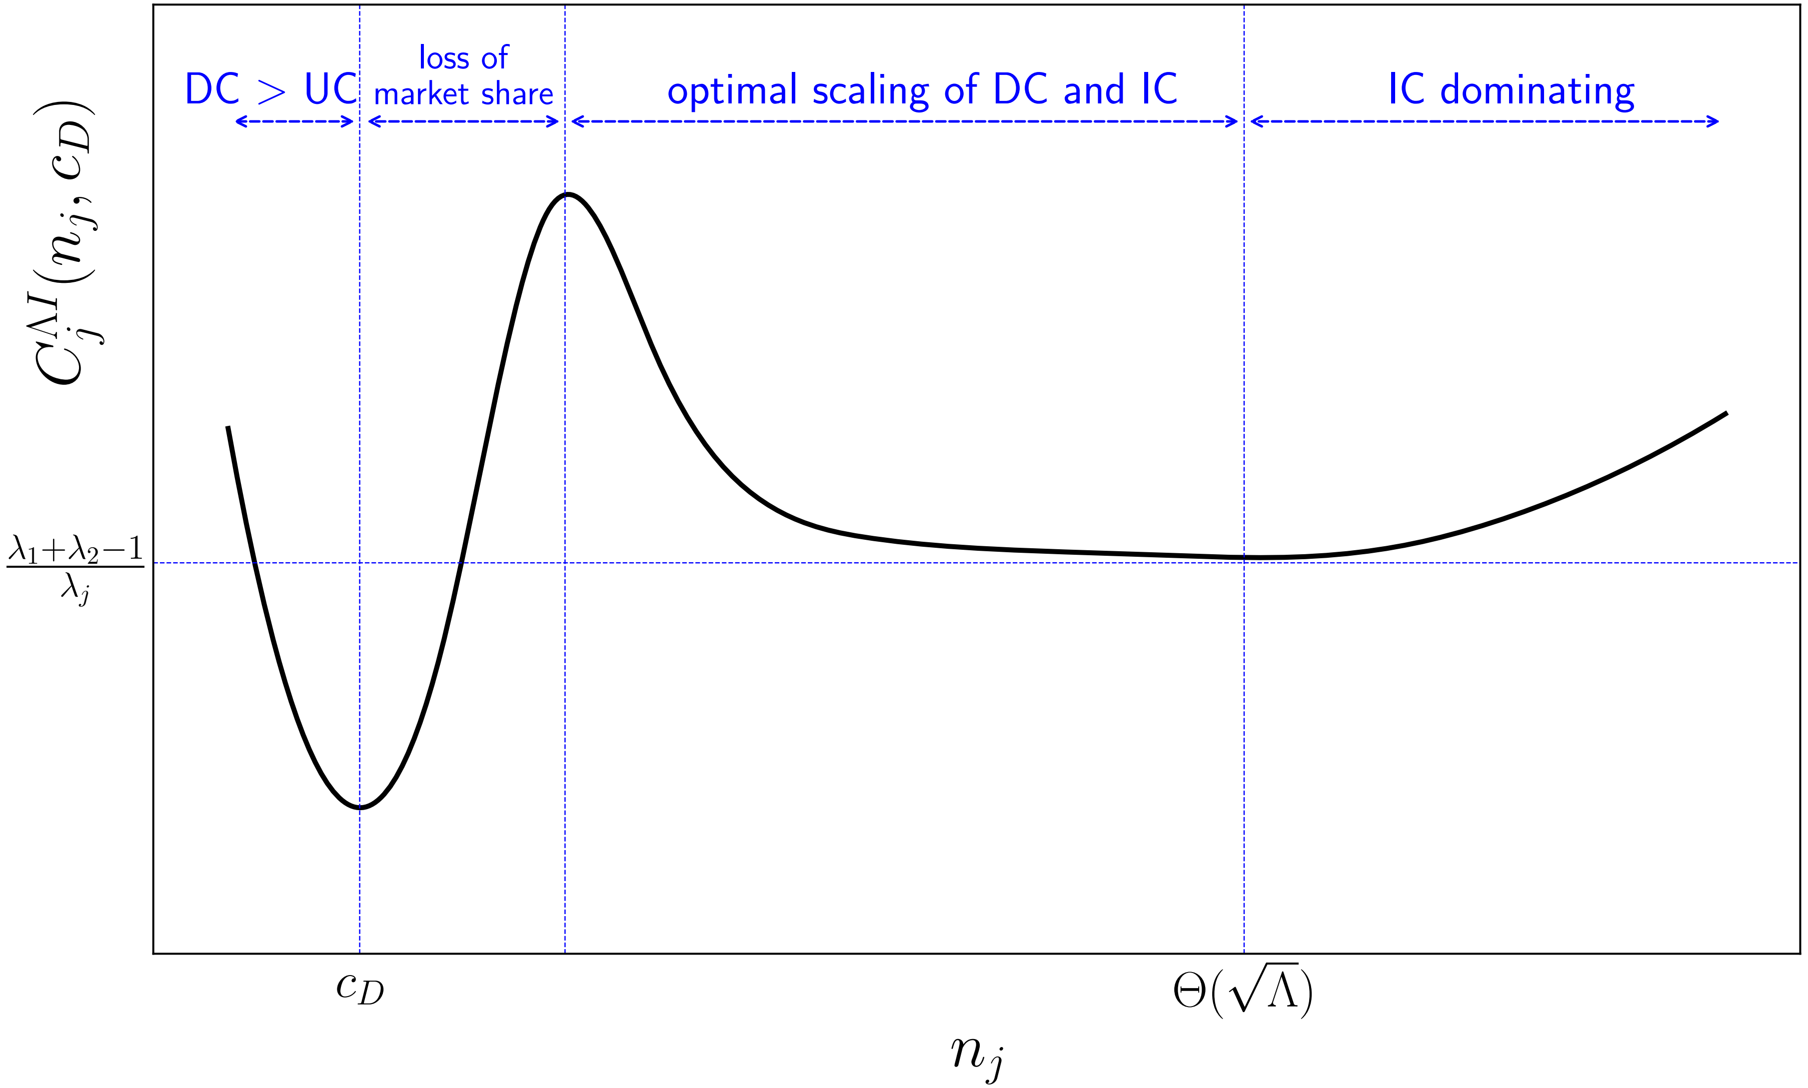

What is the resulting efficiency loss?

Price of Anarchy PoA (or Price of Stability PoS):

Ratio of worst (best) equilibrium cost to monopolist optimal cost

$$R(n_1, n_2) = \frac{C_1(n_1, n_2) + C_2(n_1, n_2)}{(\lambda_1 + \lambda_2) C_M(n^*)}$$

$$\text{PoA} = \limsup_{\Lambda \to +\infty} \sup_{\text{equilibrium } (n_1, n_2)} R(n_1, n_2)$$

$$\text{PoS}_\varepsilon = \limsup_{\Lambda \to +\infty} \inf_{\varepsilon\text{-equilibrium } (n_1, n_2)} R(n_1, n_2)$$

What is the resulting efficiency loss?

Price of Anarchy PoA (or Price of Stability PoS):

Ratio of worst (best) equilibrium cost to monopolist optimal cost

Theorem 2 [A., Aouad, Freund, '23]:

1. If $g > 0$, $1 < PoS \leq PoA \leq 2$.
2. If $g \leq 0$ and $\lambda_1 < (c_D + 1)^{-1}$, $PoA = PoS = 1$.
3. If $g \leq 0$ and $\lambda_1 \geq (c_D + 1)^{-1}$, $PoS_\varepsilon = 1$.

What is the resulting efficiency loss?

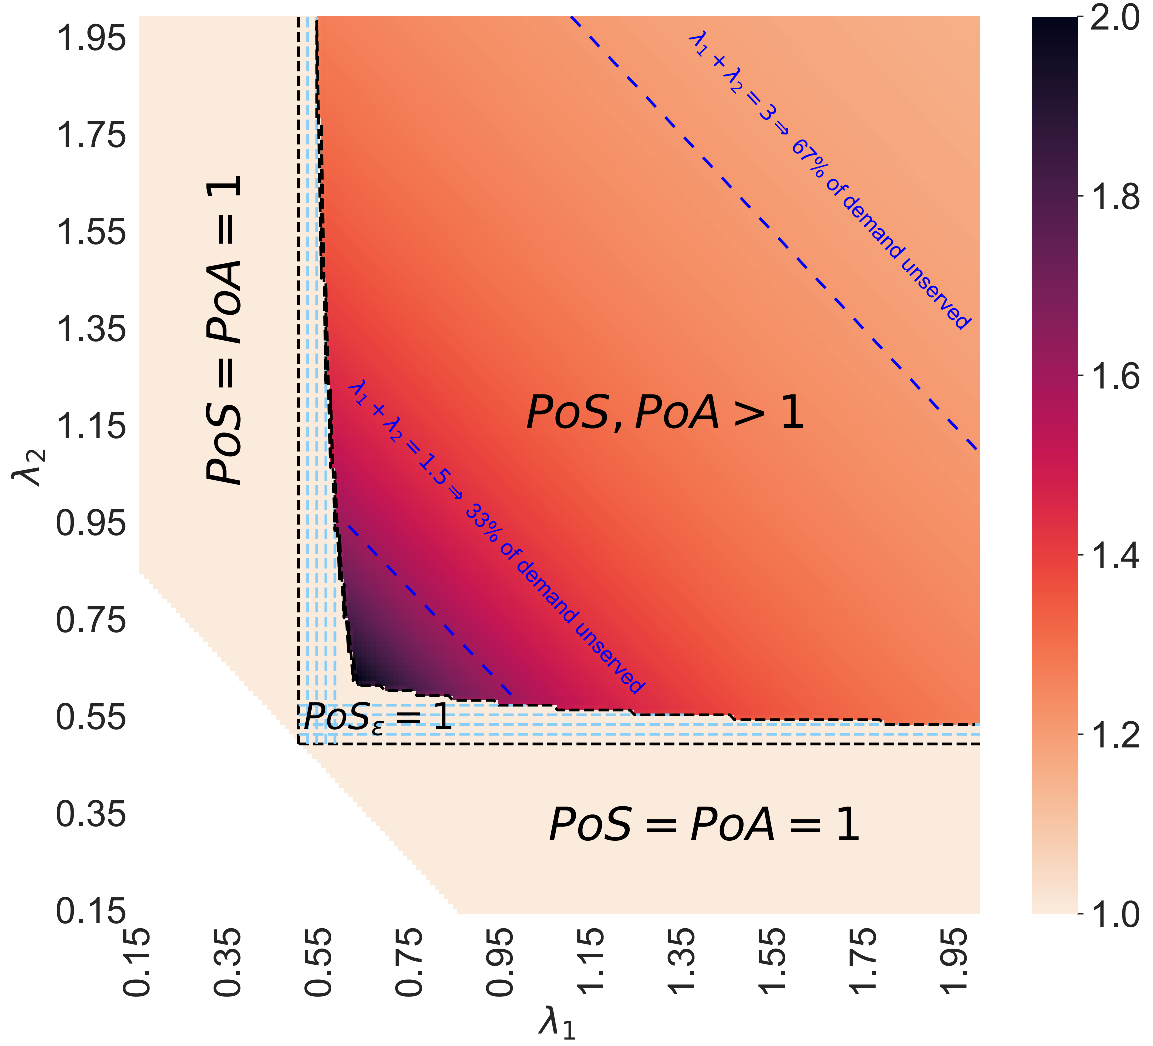

Efficiency loss: proof challenges

Lemma:
There exists a constant B such that for every large enough Λ, every equilibrium (n_1, n_2) satisfies exactly one of these:

(i) $n_1 = n_2 \in [c_D, B]$

(ii) $\min\{n_1, n_2\} \in [c_D, B]$ and $\max\{n_1, n_2\} \in [n^*(\Lambda) - 1, n^*(\Lambda) + 1]$

Ruling out equilibria of $\Theta(\Lambda)$ is very challenging since it requires an analysis of second-order cost terms

Extension to distance thresholds

- Threshold on pickup distances: (τ_1, τ_2)

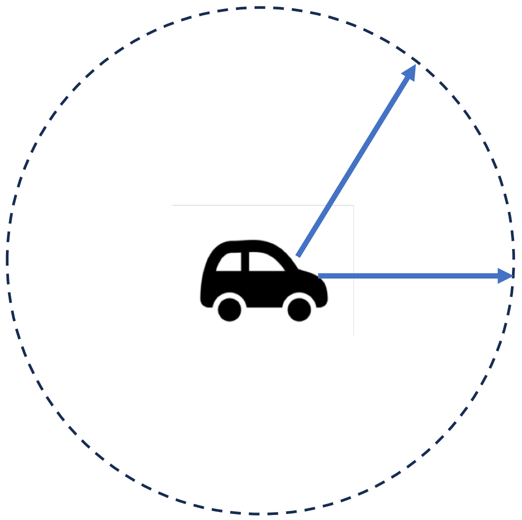

Maximum dispatch radius τ

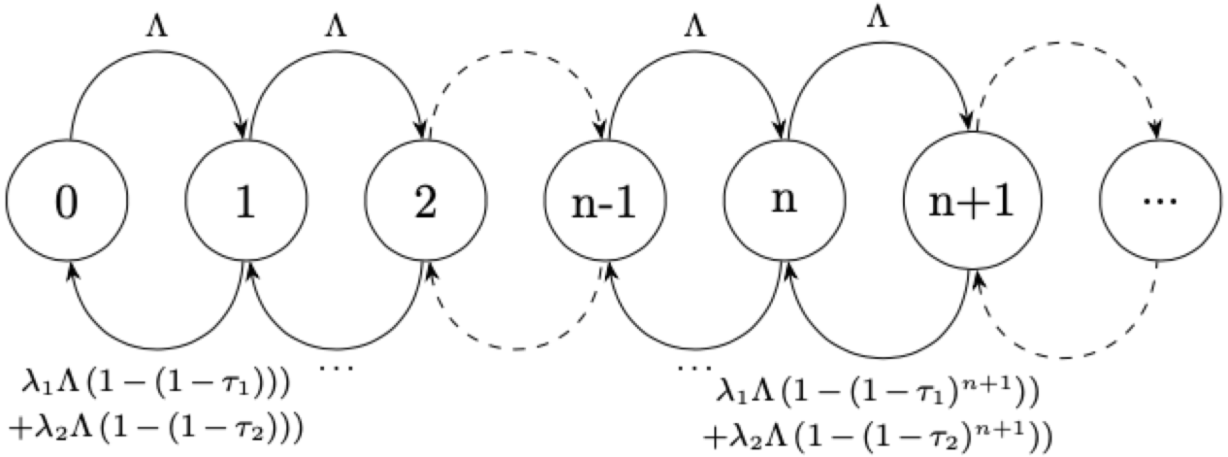

Extension to distance thresholds

Informal Theorem 3 [A., Aouad, Freund '23]:
With $I = (c_D, c_I, \lambda_1, \lambda_2, 1)$, at least one of the following holds for sufficiently large Λ:

1. Instance I is scale-inefficient.

2. There exists a family of scale-efficient ε-equilibria

Same structure but weaker result & harder to analyze

Extension to distance thresholds

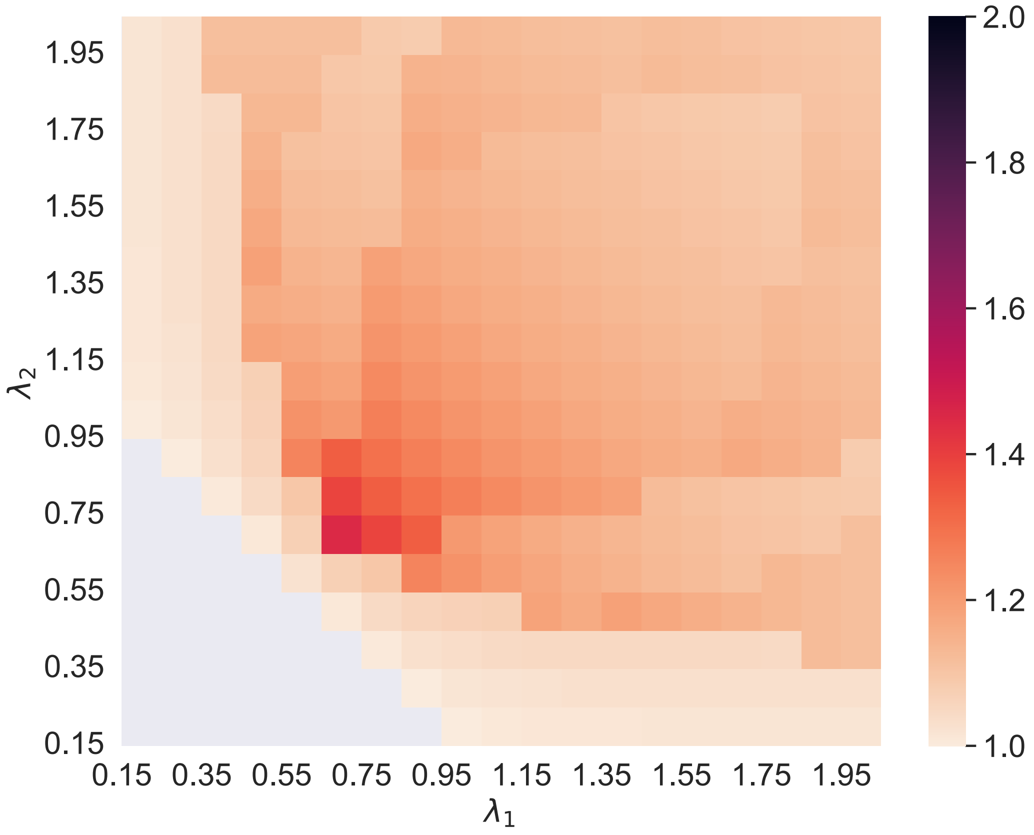

Distance-threshold PoA and PoS plots with calibrated parameters

Stylized calibration, NYC 2021-2023
$\Lambda = 8000,\ c_D = 4.29,\ c_I = 3.896$

Extension to 2-dimensional dispatch cost

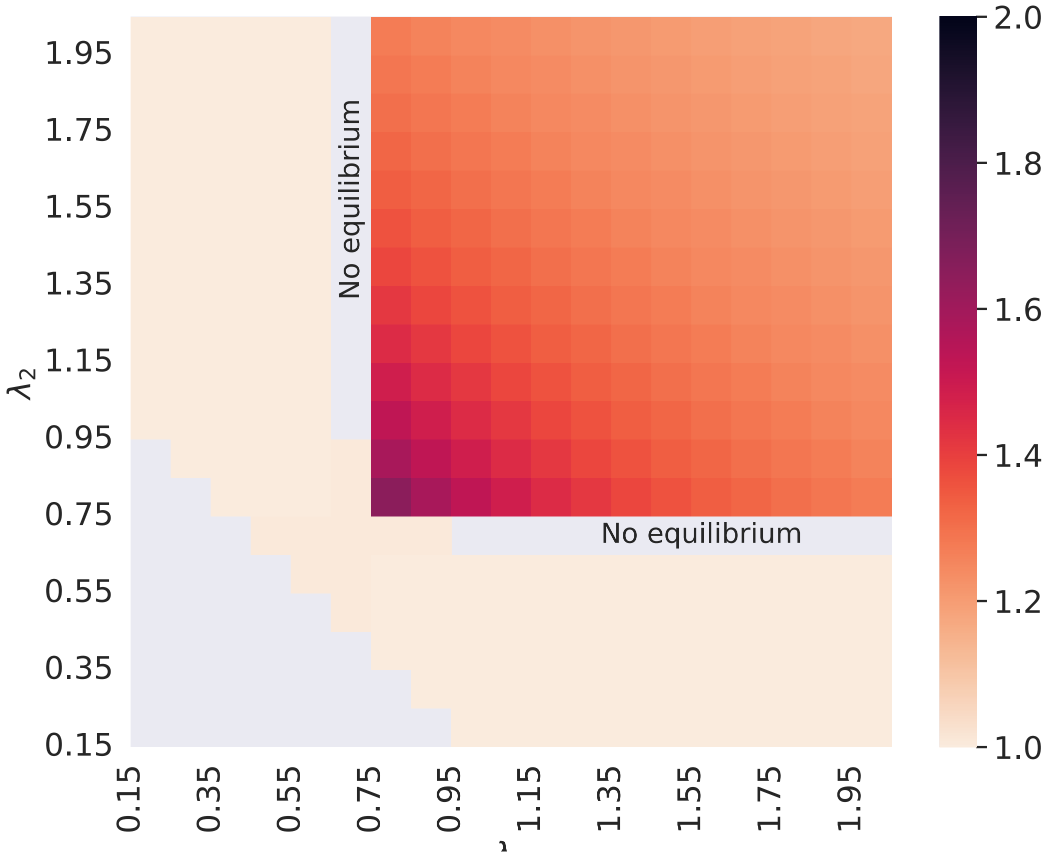

PoA and PoS with marginal dispatch cost $\approx 1/\sqrt{n}$

($\Lambda = 10^8$, $c_D = 1$, $c_I = 1$)

Fragmentation vs. matching competition?

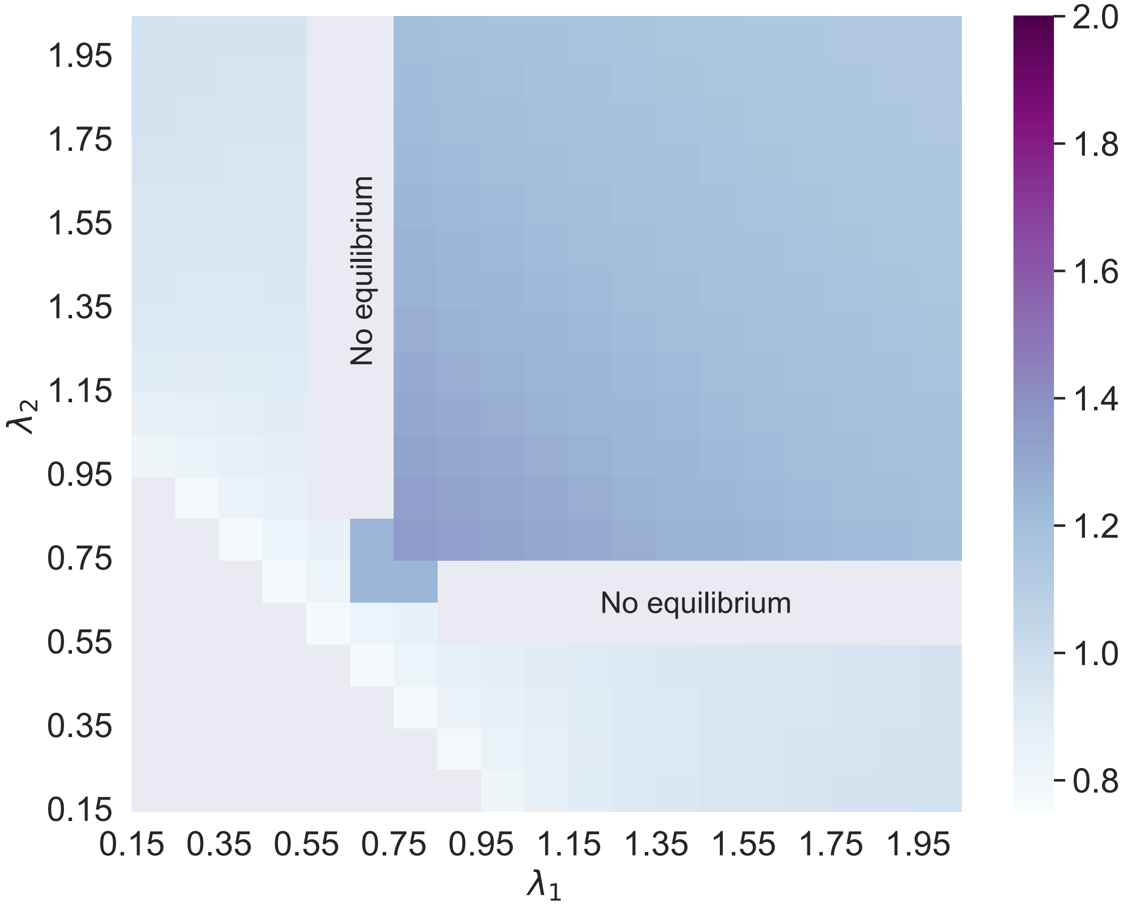

Efficiency ratio of competitive equilibrium with multihoming and fragmented market

($\Lambda = 5$, $c_D = 3$, $c_I = 0.02$)

Conclusion

- Multihoming + supply scarcity and demand imbalance ⇒ market unraveling & inefficiency of equilibria

  ! not addressed in the literature Kolkor et al. ['22], Allon et al. ['23]

- Implications for regulation policies and fragmentation

- Similar tragedy of commons for other online matching environments?


FINANCIAL TIMES

UK COMPANIES TECH MARKETS CLIMATE OPINION WORK & CAREERS LIFE & ARTS HTSI

Opinion **Disrupters**

Ride-hailing apps should work with taxis to reduce the price of anarchy

Uber's recent move to make peace with an old foe in New York must be a sign of things to come

Accessed on Dec. 2023

Questions:

aamanihamedani@london.edu