ADAPTIVE POLICIES AND APPROXIMATION SCHEMES FOR DYNAMIC MATCHING

Ali Aouad (MIT)

Alireza AmaniHamedani (LBS)

Amin Saberi (Stanford)

Dynamic matching markets

Dynamic matching markets (with reneging)

Deceased donor transplant
(NHS, 2022-2023)

	England	
Organ	N	(pmp)
Kidney		
Deceased donors	1071	(18.9)
Transplants	2018	(35.7)
Transplant list	4945	(87.5)
Pancreas		
Deceased donors	276	(4.9)
Transplants	110	(1.9)
Transplant list	239	(4.2)
Heart		
Deceased donors	159	(2.8)
Transplants[3]	155	(2.7)
Transplant list[3]	250	(4.4)

Dynamic matching markets (with reneging)

Deceased donor transplant
(NHS, 2022-2023)

	England	
Organ	N	(pmp)
Kidney		
Deceased donors	1071	(18.9)
Transplants	2018	(35.7)
Transplant list	4945	(87.5)
Pancreas		
Deceased donors	276	(4.9)
Transplants	110	(1.9)
Transplant list	239	(4.2)
Heart		
Deceased donors	159	(2.8)
Transplants[3]	155	(2.7)
Transplant list[3]	250	(4.4)

- Market thickness vs. risk of abandonment
- Kidney exchanges (deceased donors), ridehailing, emergency response.
- Online matching problem with queueing agents

Dynamic matching markets (with reneging)

Deceased donor transplant
(NHS, 2022-2023)

Organ	England N	(pmp)
Kidney		
Deceased donors	1071	(18.9)
Transplants	2018	(35.7)
Transplant list	4945	(87.5)
Pancreas		
Deceased donors	276	(4.9)
Transplants	110	(1.9)
Transplant list	239	(4.2)
Heart		
Deceased donors	159	(2.8)
Transplants[3]	155	(2.7)
Transplant list[3]	250	(4.4)

- Market thickness vs. risk of abandonment
- Ride-hailing, kidney exchanges (deceased donors), emergency response.
- Online matching problem with queueing agents

Dynamic matching markets (with reneging)

Deceased donor transplant
(NHS, 2022-2023)

Organ	England N	(pmp)
Kidney		
Deceased donors	1071	(18.9)
Transplants	2018	(35.7)
Transplant list	4945	(87.5)
Pancreas		
Deceased donors	276	(4.9)
Transplants	110	(1.9)
Transplant list	239	(4.2)
Heart		
Deceased donors	159	(2.8)
Transplants[3]	155	(2.7)
Transplant list[3]	250	(4.4)

- Market thickness vs. risk of abandonment
- Ride-hailing, kidney exchanges (deceased donors), emergency response.
- Online matching problem with queueing agents

Related work

- Online matching: Karp et al. [‘90], Mehta et al. [‘07], Manshadi et al. [‘12], Jaillet and Lu [‘14], Huang & Xu [‘21], Papadimitriou et al. [‘21], …

- Matching queues: Caldentey and Kaplan [‘02], Bušić et al. [‘13], Gurvitch and Ward [‘14], Tsisiklis and Xu [‘17], Anderson et al. [‘17], Afèche et al. [‘19], Chen et al. [‘22], etc.

Related work

- Online matching: Karp et al. ['90], Mehta et al. ['07], Manshadi et al. ['12], Jaillet and Lu ['14], Huang & Xu ['21], Papadimitriou et al. ['21], …

  constant-factor sub-optimality

- Matching queues: Caldentey and Kaplan ['02], Bušić et al. ['13], Gurvitch and Ward ['14], Tsisiklis and Xu ['17], Anderson et al. ['17], Afèche et al. ['19], Chen et al. ['22], etc.

  asymptotic optimality

Related work

- Online matching: Karp et al. ['90], Mehta et al. ['07], Manshadi et al. ['12], Jaillet and Lu ['14], Huang & Xu ['21], Papadimitriou et al. ['21], …

- Matching queues: Caldentey and Kaplan ['02], Bušić et al. ['13], Gurvitch and Ward ['14], Tsisiklis and Xu ['17], Anderson et al. ['17], Afèche et al. ['19], Chen et al. ['22],…

- Dynamic matching: Aouad & Saritaç ['20], Collina et al. ['20], Kessel et al. ['22], Li et al. ['23], Kohlenberg and Gurvich ['24], Yu and Vossen ['24], Patel & Wajc ['24],…

Related work

- Online matching: Karp et al. ['90], Mehta et al. ['07], Manshadi et al. ['12], Jaillet and Lu ['14], Huang & Xu ['21], Papadimitriou et al. ['21], …
- Matching queues: Caldentey and Kaplan ['02], Bušić et al. ['13], Gurvitch and Ward ['14], Tsisiklis and Xu ['17], Anderson et al. ['17], Afèche et al. ['19], Chen et al. ['22],…
- Dynamic matching: Aouad & Saritaç ['20], Collina et al. ['20], Kessel et al. ['22], Li et al. ['23], Kohlenberg and Gurvich ['24], Yu and Vossen ['24], Patel & Wajc ['24],…

Recent literature focuses on simple, static policies

Related work

- Online matching: Karp et al. ['90], Mehta et al. ['07], Manshadi et al. ['12], Jaillet and Lu ['14], Huang & Xu ['21], Papadimitriou et al. ['21], …
- Matching queues: Caldentey and Kaplan ['02], Bušić et al. ['13], Gurvitch and Ward ['14], Tsisiklis and Xu ['17], Anderson et al. ['17], Afèche et al. ['19], Chen et al. ['22],…
- Dynamic matching: Aouad & Saritaç ['20], Collina et al. ['20], Kessel et al. ['22], Li et al. ['23], Kohlenberg and Gurvich ['24], Yu and Vossen ['24], Patel & Wajc ['24],…

Research question: How to design **adaptive** policies?

This talk

I Designing adaptive policies: Primal-dual interpretation

This talk

1. Designing adaptive policies: Primal-dual interpretation
2. Near-optimal algorithms for small networks or Euclidean networks

This talk

1. Designing adaptive policies: Primal-dual interpretation
2. Near-optimal algorithms for small networks or Euclidean networks
3. New LP relaxation framework: Hybrid of dynamic programming and fluid

Dynamic matching model

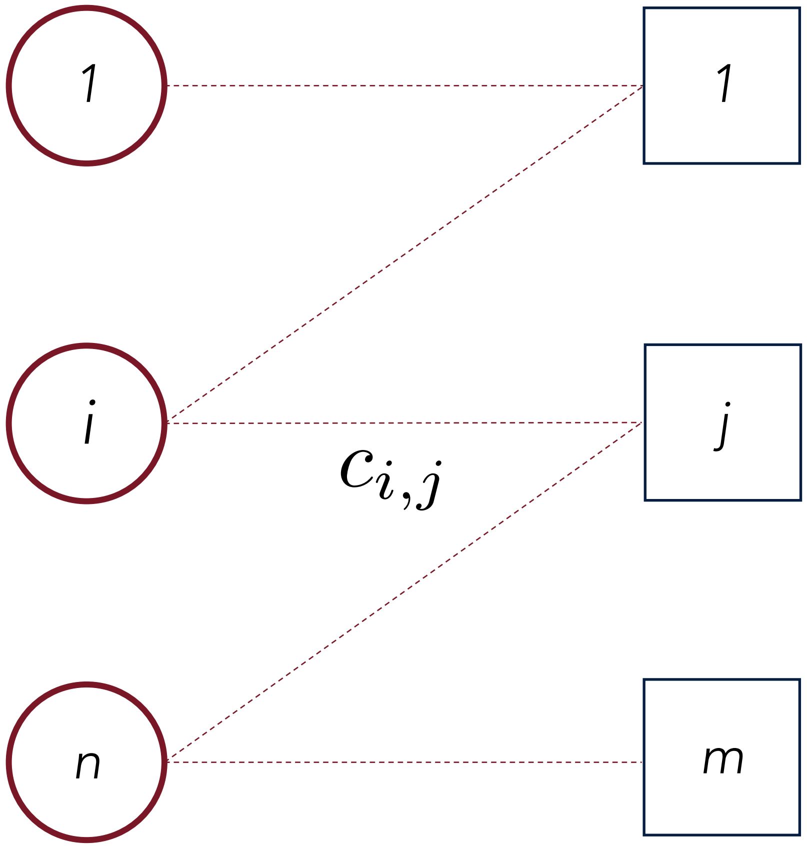

"Types" in edge-weighted bipartite graph

1 Network

Dynamic matching model

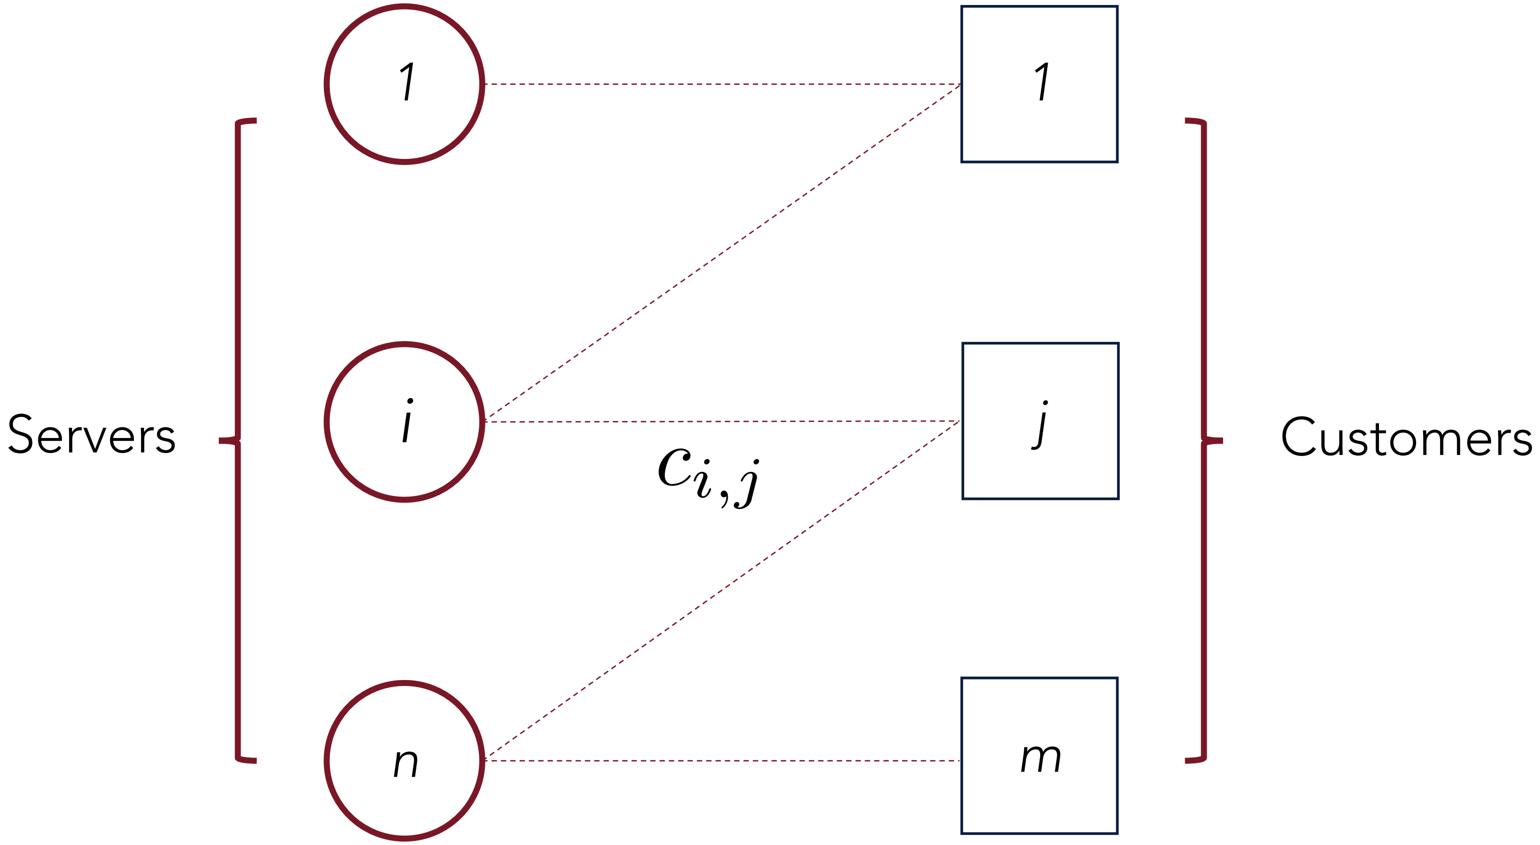

e.g. Euclidean graph ⇒ cost is distance
(pickup time in ridehailing)

1 Network

Dynamic matching model

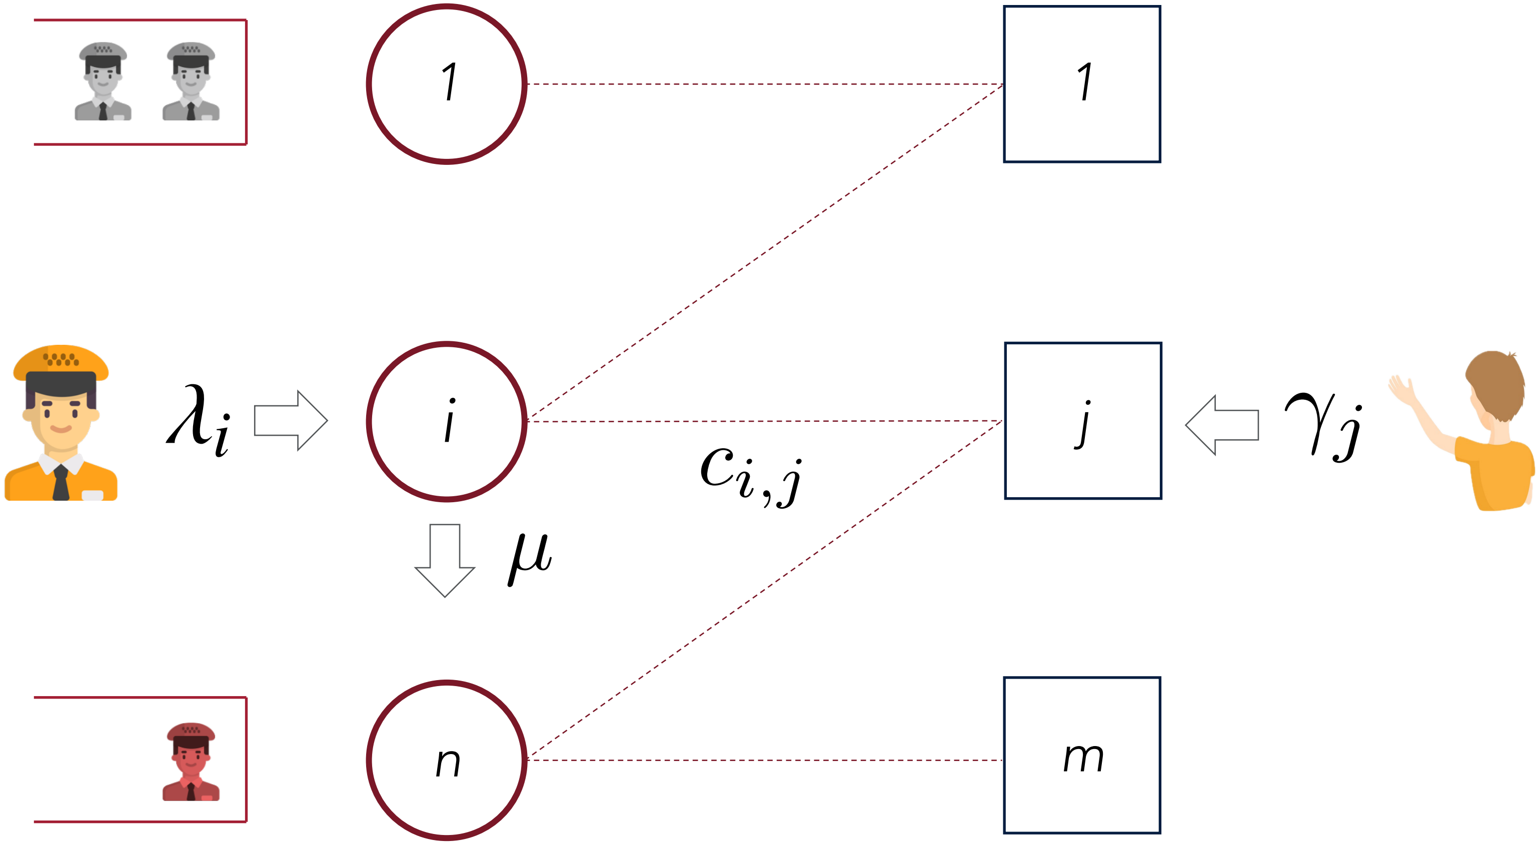

1 Network

Dynamic matching model

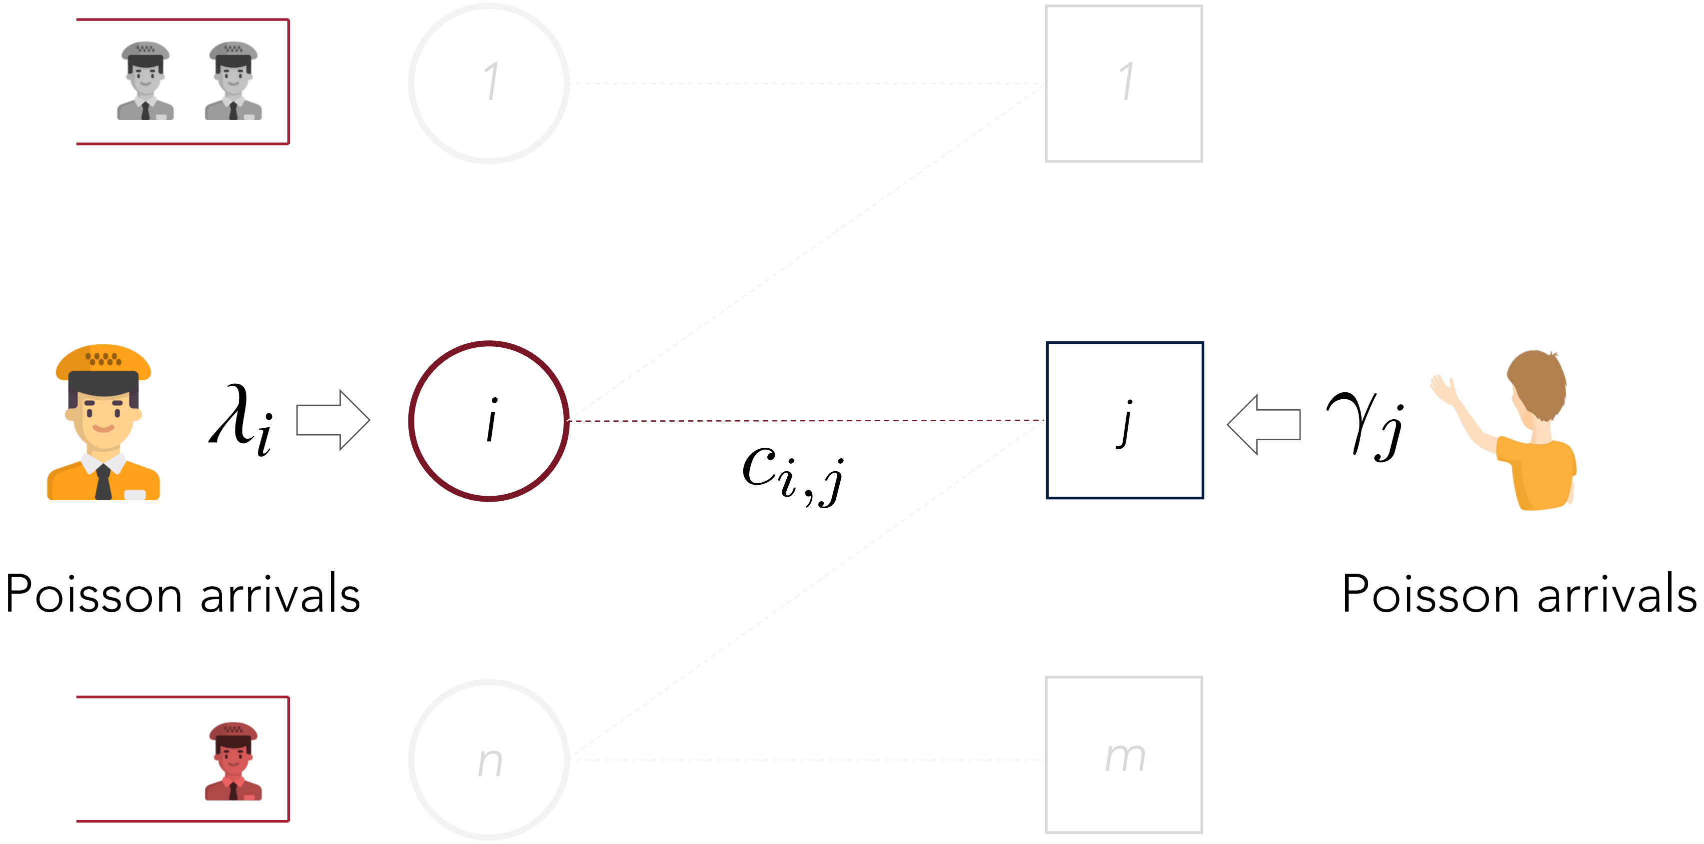

2 Stochastic process

Dynamic matching model

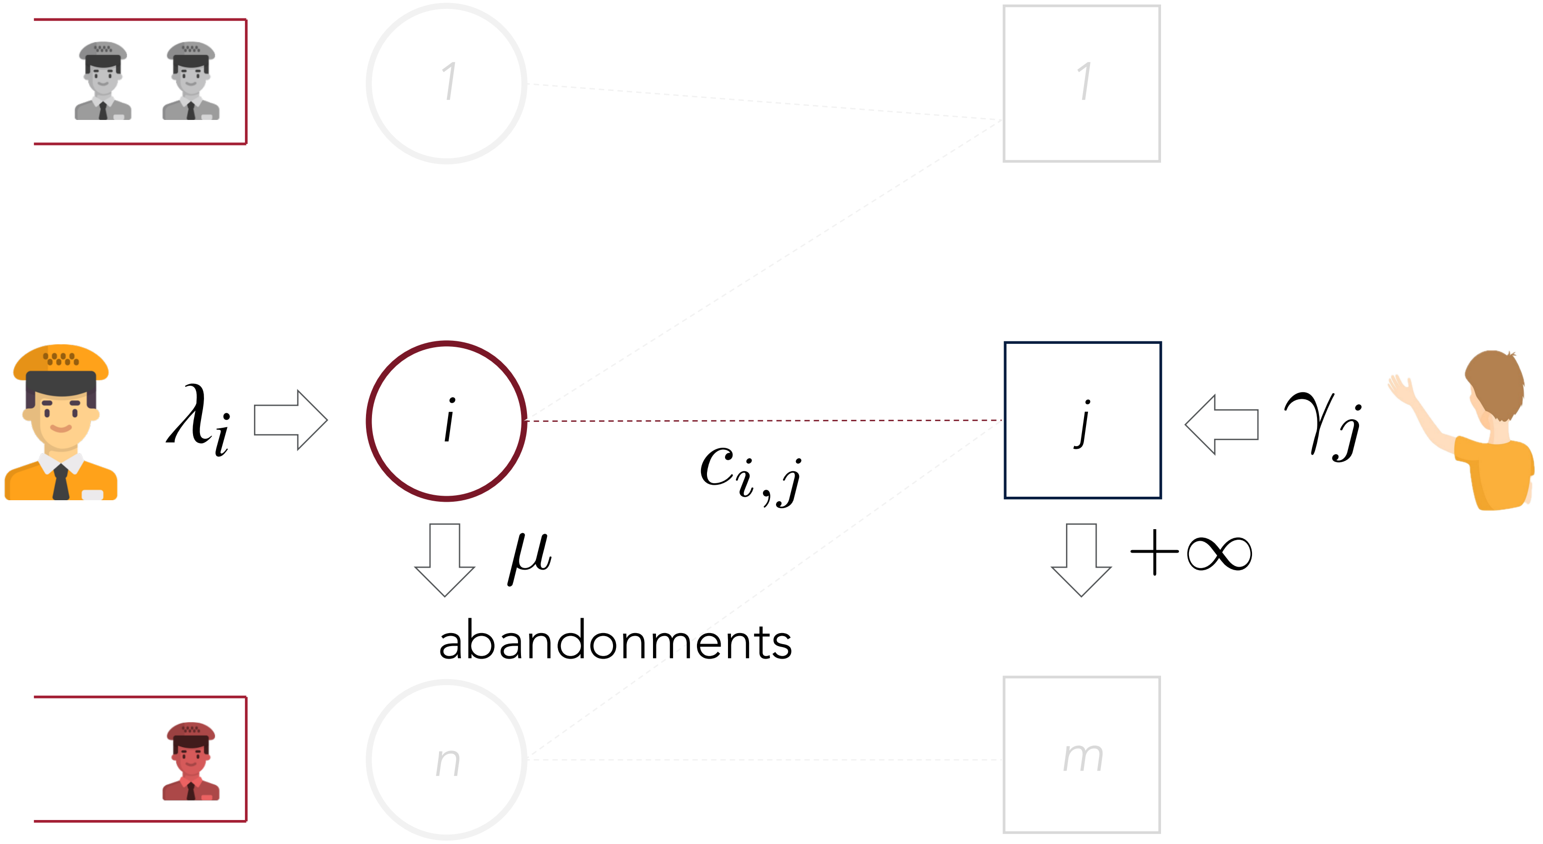

2 Stochastic process

Dynamic matching model

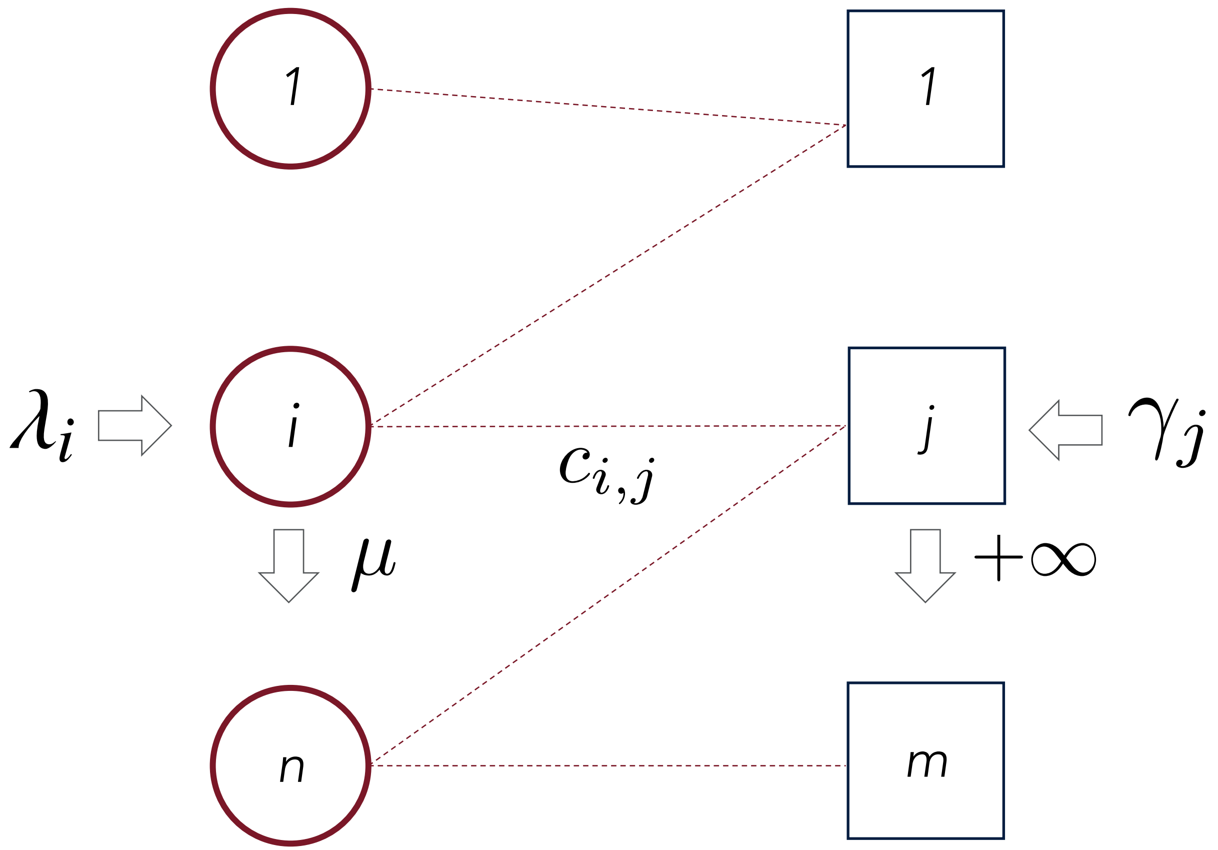

3 Optimality criterion

throughput target τ^*

$$\limsup_{t \to \infty} \frac{\mathbb{E}[T^\pi(t)]}{t} \geq \tau^*$$

Dynamic matching model

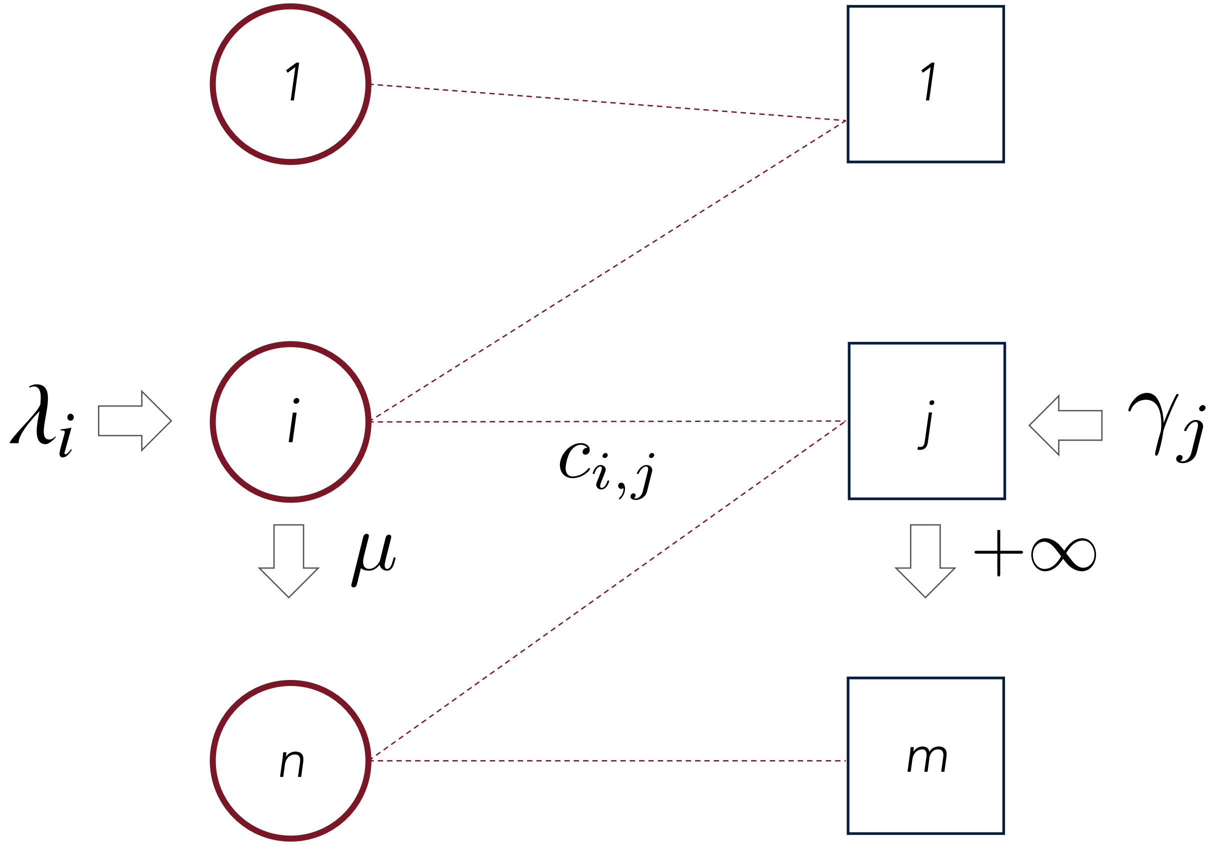

3 Optimality criterion

cost-throughput target (c^*, τ^*)

$$\limsup_{t \to \infty} \frac{\mathbb{E}[C^\pi(t)]}{t} \leq c^*$$

$$\limsup_{t \to \infty} \frac{\mathbb{E}[T^\pi(t)]}{t} \geq \tau^*$$

Dynamic matching model

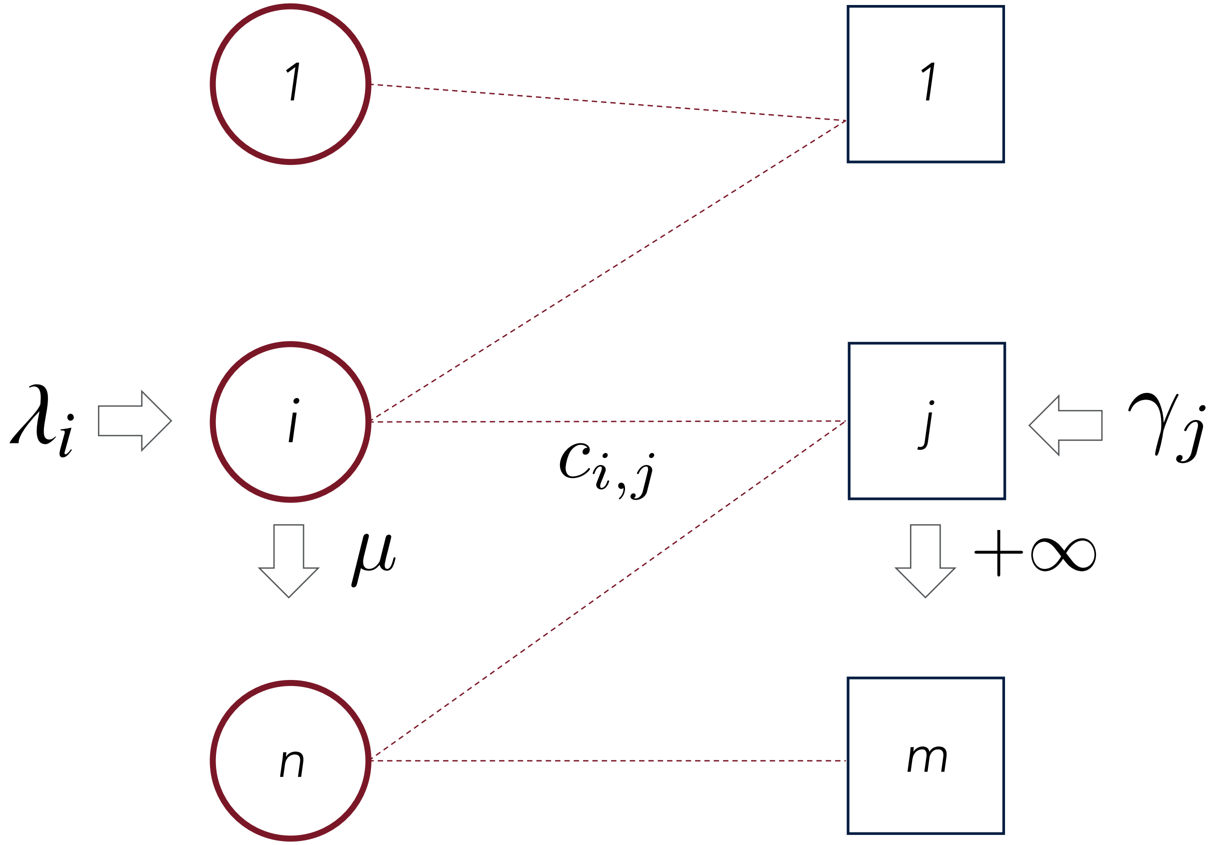

3 Optimality criterion

cost-throughput target (c^*, τ^*)

$$\limsup_{t\to\infty} \frac{\mathbb{E}[C^\pi(t)]}{t} \leq c^*$$

$$\limsup_{t\to\infty} \frac{\mathbb{E}[T^\pi(t)]}{t} \geq \tau^*$$

Dynamic matching model

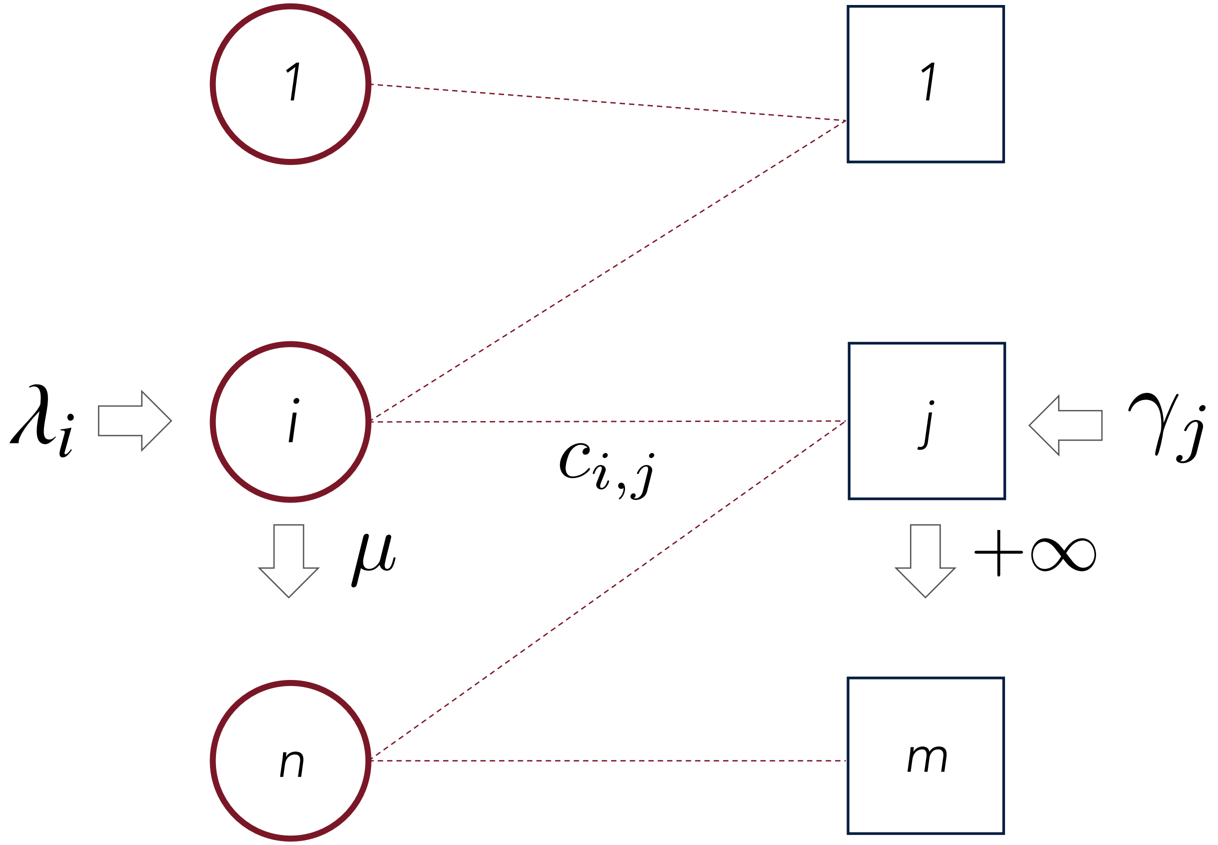

3 Optimality criterion

cost-throughput target (c^*, τ^*)

$$\limsup_{t \to \infty} \frac{\mathbb{E}[C^\pi(t)]}{t} \leq c^* \cdot (1 + \epsilon)$$

$$\limsup_{t \to \infty} \frac{\mathbb{E}[T^\pi(t)]}{t} \geq \tau^* \cdot (1 - \epsilon)$$

Some challenges

- Average-cost infinite-dimensional MDP: formulation
- “Endogenous” market thickness: steady-state induced by policy

Some challenges

- Average-cost infinite-dimensional MDP: formulation
- "Endogenous" market thickness: steady-state induced by policy
- No asymptotic scaling & thin market ⚠️: unlike O(1)-regret dynamic matching Ashlagi et al. ['23], Gupta ['22], Wei et al. ['23]

Static LP relaxation [AS, ’20]

Static LP relaxation [AS, '20]

$$
\begin{aligned}
\text{(SLP)} \quad & \min_{x_{i,j}, x_{i,a} \geq 0} && \sum_{(i,j)} c_{i,j} \cdot x_{i,j} && \\
& \text{s.t.} && \sum_{j} x_{i,j} + x_{i,a} = \lambda_i \ , && \forall i \\
& && \sum_{(i,j)} x_{i,j} \geq \tau^* \ , && \\
& && \frac{\mu_i}{\lambda_j} \cdot x_{i,j} \leq x_{i,a} \ , && \forall (i,j)
\end{aligned}
$$

Static LP relaxation [AS, '20]

$$(\text{SLP}) \quad \min_{x_{i,j}, x_{i,a} \geq 0} \quad \sum_{(i,j)} c_{i,j} \cdot x_{i,j}$$

$$\text{s.t.} \quad \sum_{j} x_{i,j} + x_{i,a} = \lambda_i \ , \qquad \forall i$$

$$\sum_{(i,j)} x_{i,j} \geq \tau^* \ ,$$

$$\frac{\mu_i}{\lambda_j} \cdot x_{i,j} \leq x_{i,a} \ , \qquad \forall (i,j)$$

The value of adaptive policies

The value of adaptive policies

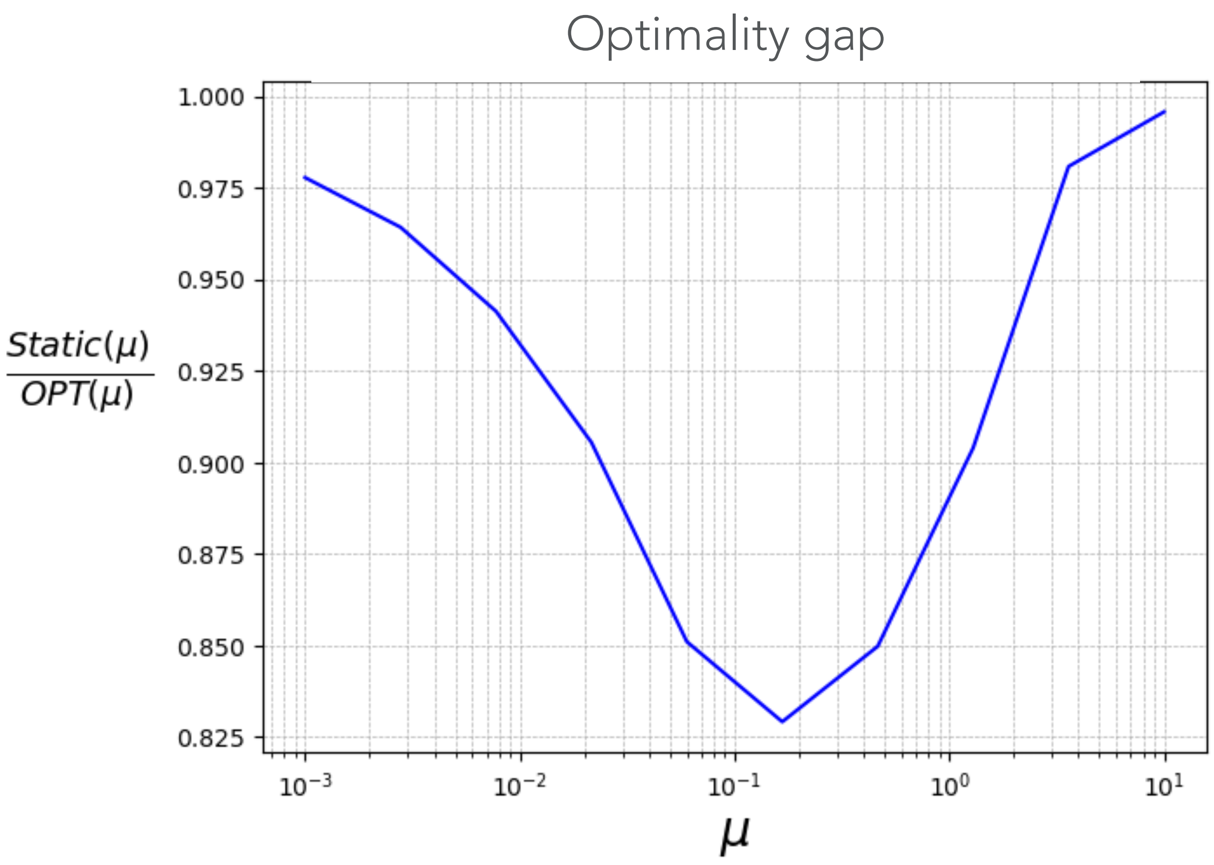

3 customer types, one server type with $\lambda = 1$

The value of adaptive policies

Optimality gap

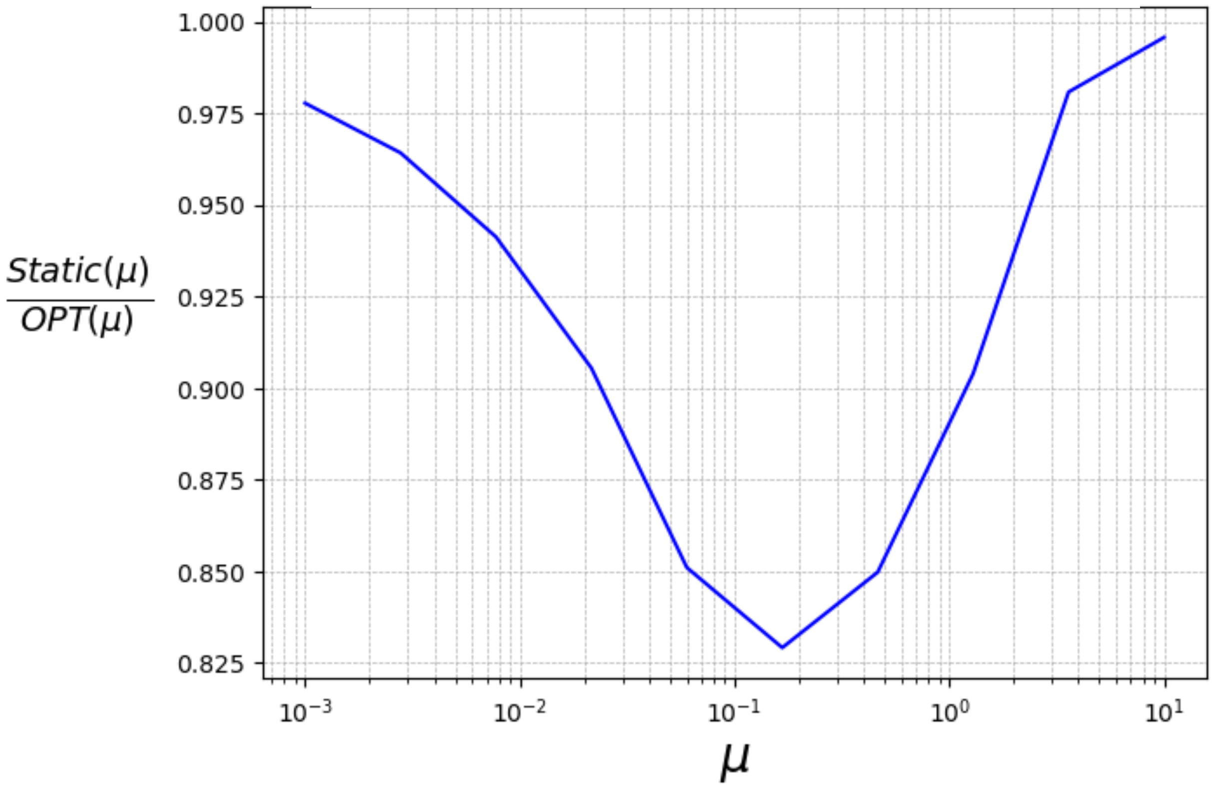

Static LP integrality gap

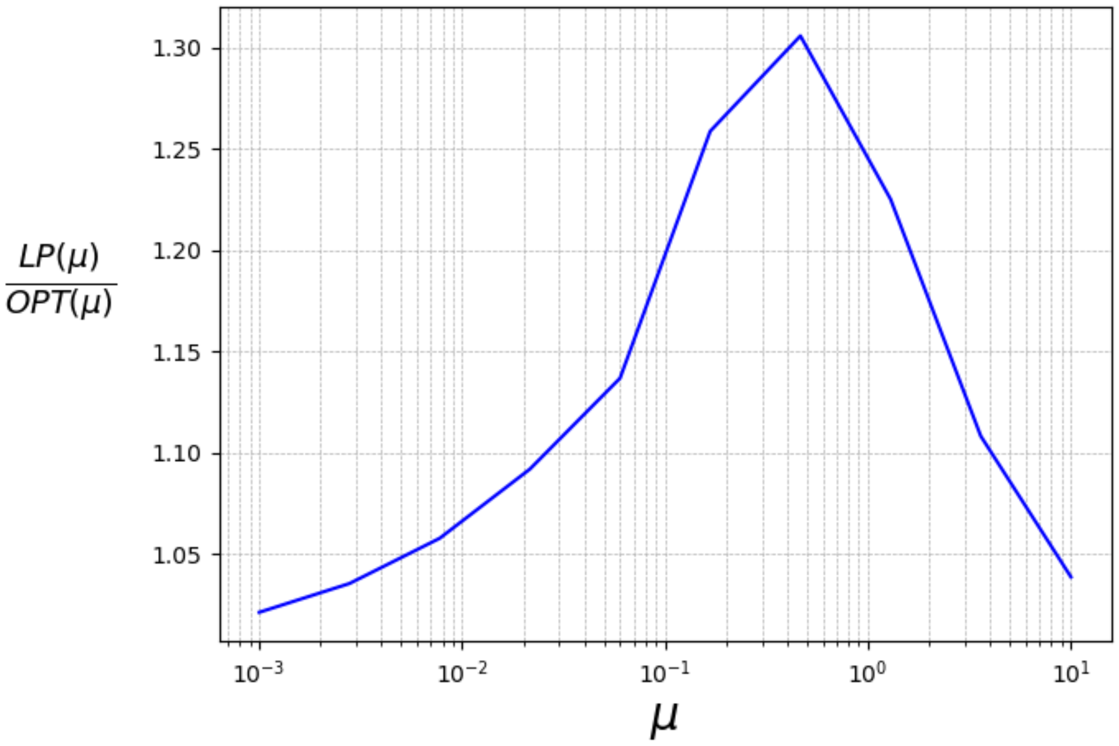

3 customer types, one server type with $\lambda = 1$

Single queue: An exact Dynamic LP

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell$$

$$\text{s.t.} \quad (x_M^\ell)_{M,\ell} \in \mathcal{K}$$

birth-death process

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell$$

$$\text{s.t.} \quad (x_M^\ell)_{M,\ell} \in \mathcal{K}$$

birth-death process

$(\mathcal{K})$:

- detail balance
- probas-sum-to-1

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell$$

$$\text{s.t.} \quad (x_M^\ell)_{M,\ell} \in \mathcal{K}$$

birth-death process

$$(\mathcal{K}) \begin{cases} \lambda \cdot \sum_M x_{i,S}^{\ell-1} = \sum_M x_M^\ell \cdot (\gamma(M) + \mu \cdot \ell) \\ \text{probas-sum-to-1} \end{cases}$$

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell$$

$$\text{s.t.} \quad (x_M^\ell)_{M,\ell} \in \mathcal{K}$$

birth-death process

$$(\mathcal{K}) \begin{cases} \lambda \cdot \sum_M x_{i,S}^{\ell-1} = \sum_M x_M^\ell \cdot (\gamma(M) + \mu \cdot \ell) \\ \sum_{M,\ell} x_M^\ell = 1 \end{cases}$$

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

$$(\text{DLP})$$

$$\begin{aligned} \min_{\boldsymbol{x} \geq 0} \quad & \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell \\ \text{s.t.} \quad & (x_M^\ell)_{M,\ell} \in \mathcal{K} \end{aligned}$$

throughput target .

$$(\mathcal{K}) \begin{cases} \lambda \cdot \sum_M x_{i,S}^{\ell-1} = \sum_M x_M^\ell \cdot (\gamma(M) + \mu \cdot \ell) \\ \sum_{M,\ell} x_M^\ell = 1 \end{cases}$$

Single queue: An exact Dynamic LP

- Decision variable x_M^ℓ : stationary probability that $\ell \in \mathbb{N}$ servers are waiting, and the optimal policy is about to match them with a customer in $\forall M, \ell$

(PASTA property [Wolff '82])

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{\ell, M} \sum_{j \in M} \gamma_j \cdot c_j \cdot x_M^\ell$$

$$\text{s.t.} \quad (x_M^\ell)_{M,\ell} \in \mathcal{K}$$

$$\sum_{\ell} \sum_{M} \gamma(M) \cdot x_M^\ell \geq \tau^* \ .$$

$$(\mathcal{K}) \begin{cases} \lambda \cdot \sum_{M} x_{i,S}^{\ell-1} = \sum_{M} x_M^\ell \cdot (\gamma(M) + \mu \cdot \ell) \\ \sum_{M,\ell} x_M^\ell = 1 \end{cases}$$

Single queue: Primal-dual solution

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy has queue length-dependent thresholds δ^ℓ (increasing concave).

Single queue: Primal-dual solution

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy has queue length-dependent thresholds δ^ℓ (increasing concave).

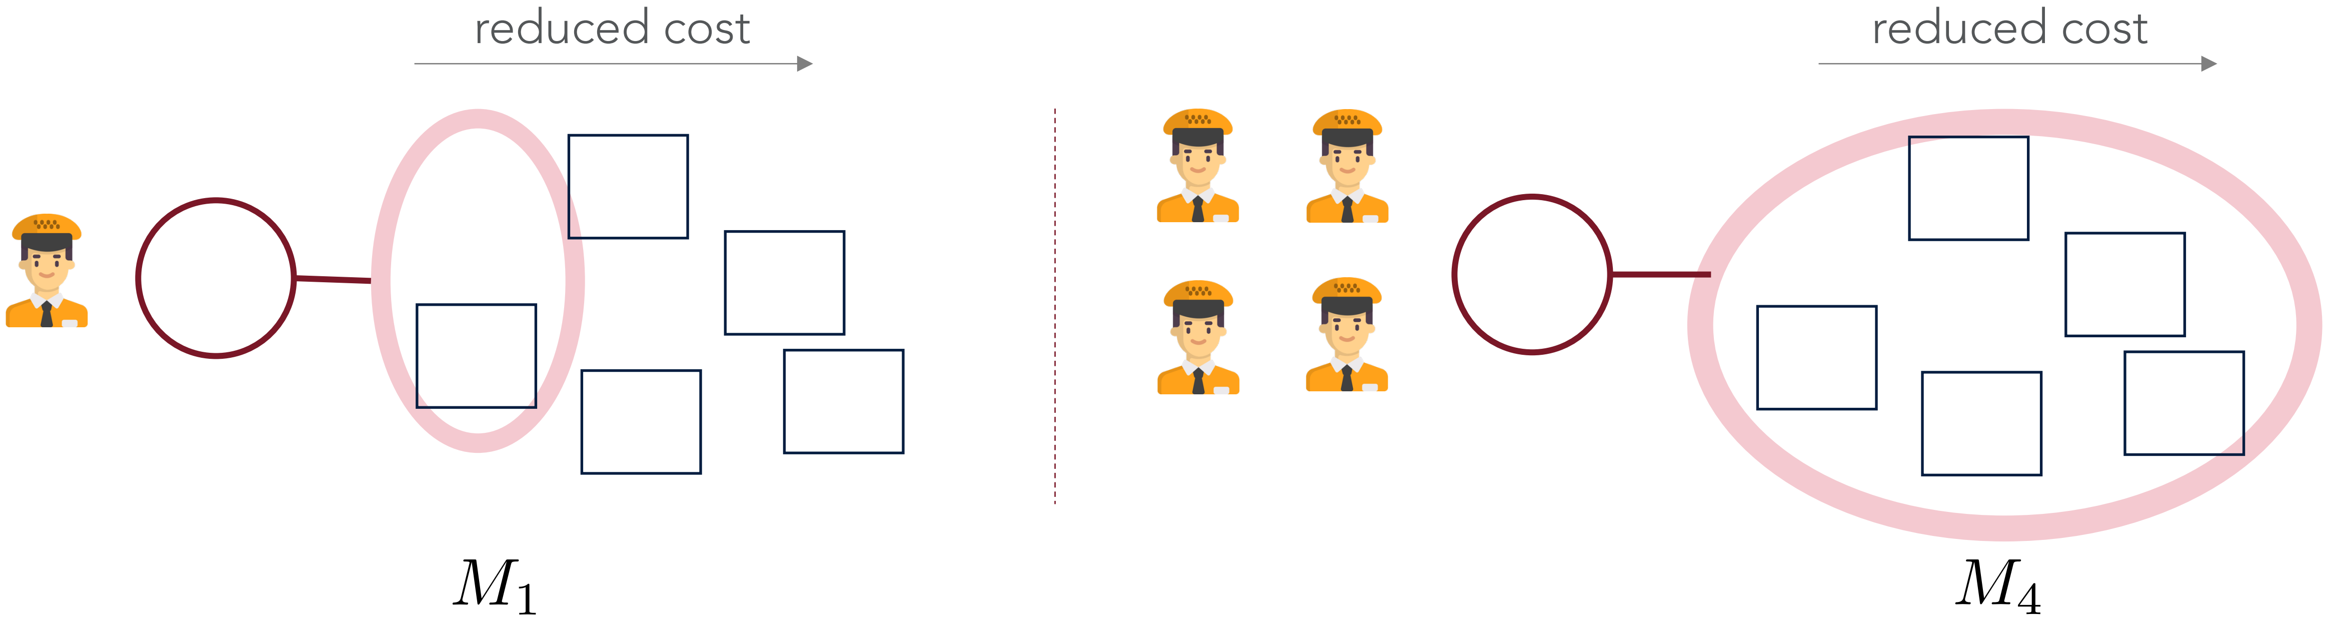

$$M_\ell = \{j \in M : c_j - \theta \le \delta^\ell\}$$

Single queue: Primal-dual solution

Theorem 1 [AAS'24]: There exists a Fully Polynomial Time Approximation Scheme for the bi-criteria dynamic matching problem with a single queue. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(\epsilon^{-O(1)} \cdot |\mathcal{I}|^{O(1)}\right)$.

Single queue: Primal-dual solution

Theorem 1 [AAS'24]: There exists a Fully Polynomial Time Approximation Scheme for the bi-criteria dynamic matching problem with a single queue. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(\epsilon^{-O(1)} \cdot |\mathcal{I}|^{O(1)}\right)$.

Proof: Primal-dual algorithm

- Exponential queue lengths $\Longrightarrow$ Polynomial truncation (sensitivity analysis)
- Exponential matching sets $\Longrightarrow$ Separation in the dual

Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$

Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$ (*m* arbitrary)

Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$
- "pigeons fly with pigeons, hawkes fly with hawkes"

Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$
- “pigeons fly with pigeons, hawkes fly with hawkes”


Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$
- "pigeons fly with pigeons, hawkes fly with hawkes"

scarce servers ⇔ infrequent customers

abundant servers ⇔ frequent customers

Small networks: A tale of two timescales

- Small networks: $n \leq \Upsilon$
- "pigeons fly with pigeons, hawkes fly with hawkes"

scarce servers ⇔ infrequent customers > abundant servers ⇔ frequent customers

Small networks: The hybrid LP

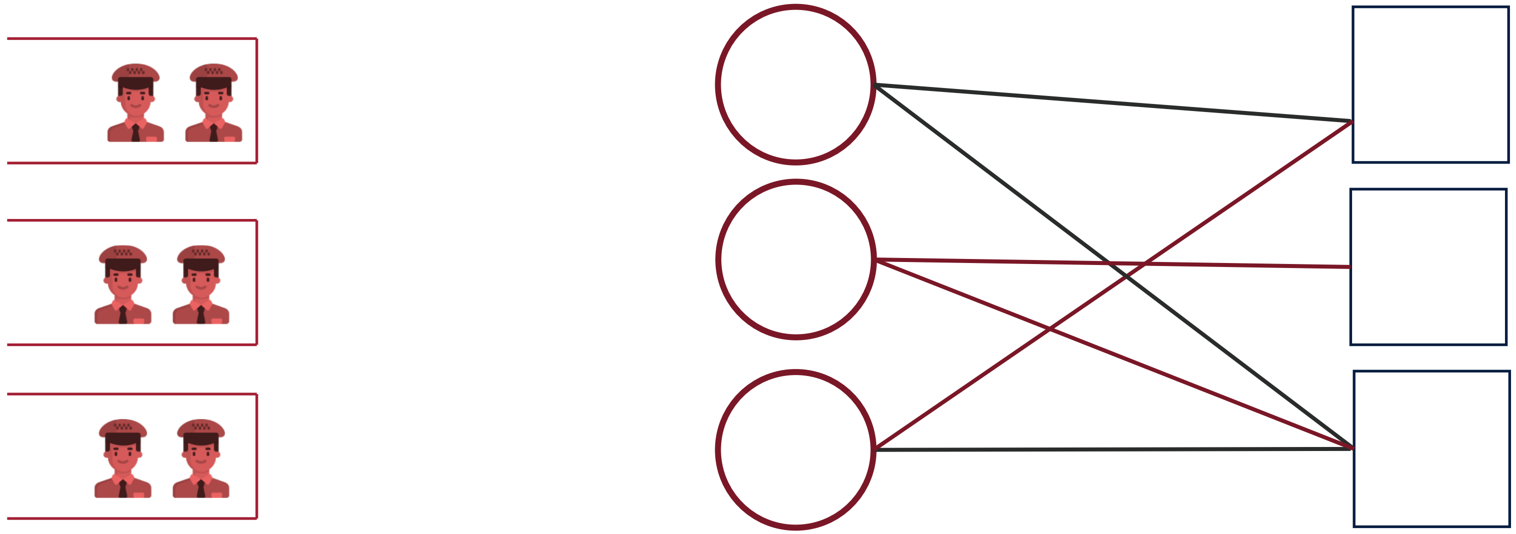

Small networks: The hybrid LP

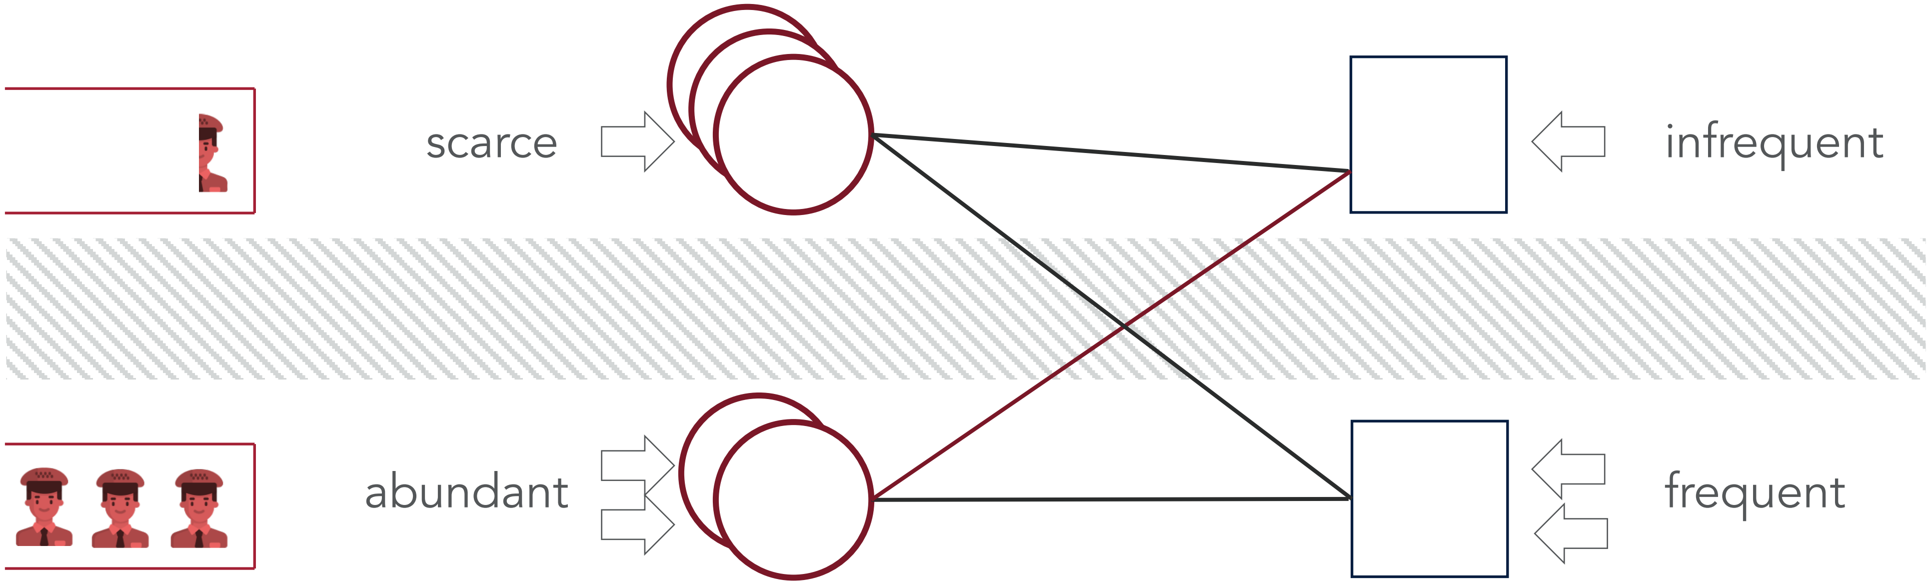

Small networks: The hybrid LP

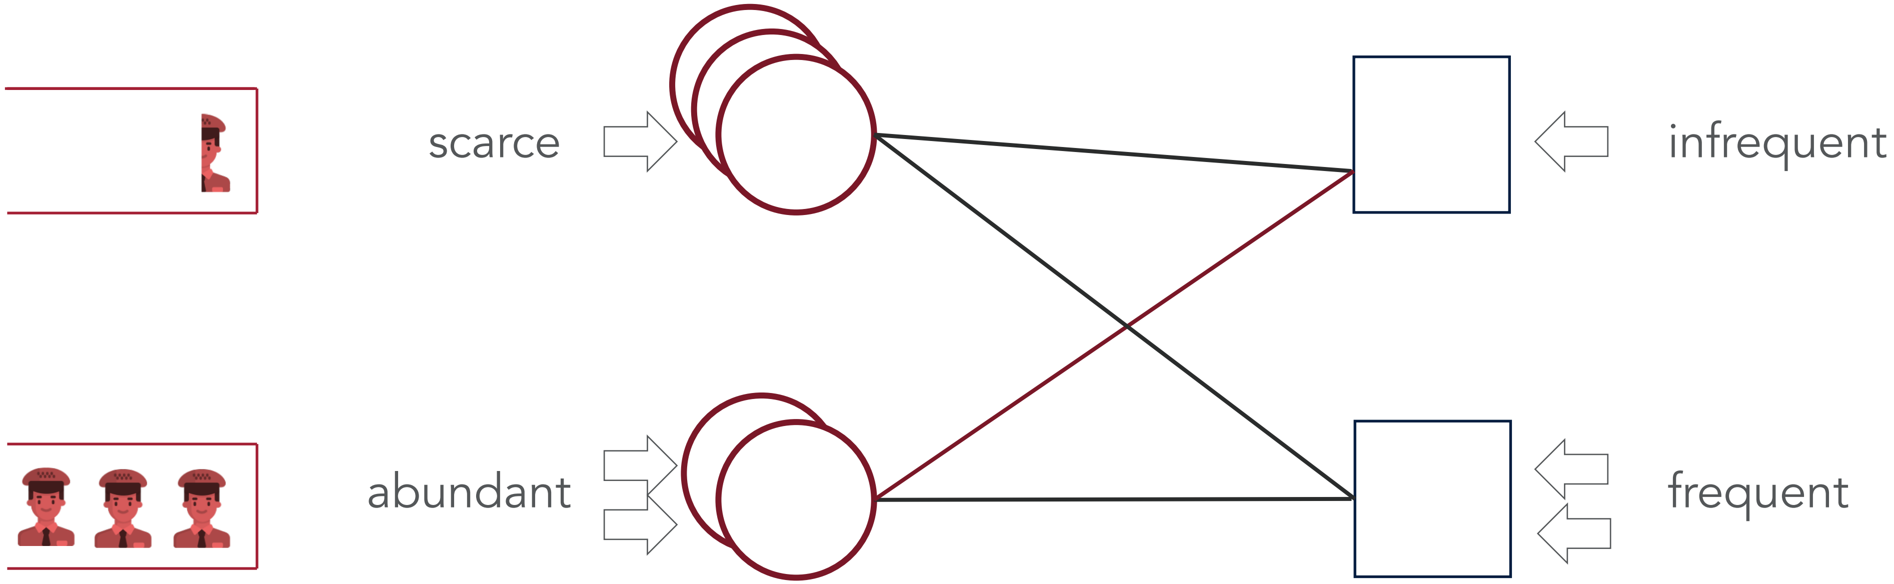

Small networks: The hybrid LP

$$\min_{\boldsymbol{x},\boldsymbol{y}} \quad \underbrace{\sum_{i\in\mathcal{S}^{\mathrm{s}}} \sum_{\boldsymbol{M},j\in M_i} \gamma_j c_{i,j} \cdot x^{\boldsymbol{\ell}}_{\mathbf{M}}}_{\text{“dynamic” variables}} + \underbrace{\sum_{i\in\mathcal{S}^{\mathrm{a}}} \sum_{j\in[m]} \gamma_j c_{i,j} \cdot y_{i,j}}_{\text{“static” variables}}$$

$$s.t. \quad \ldots$$

A tale of two timescales

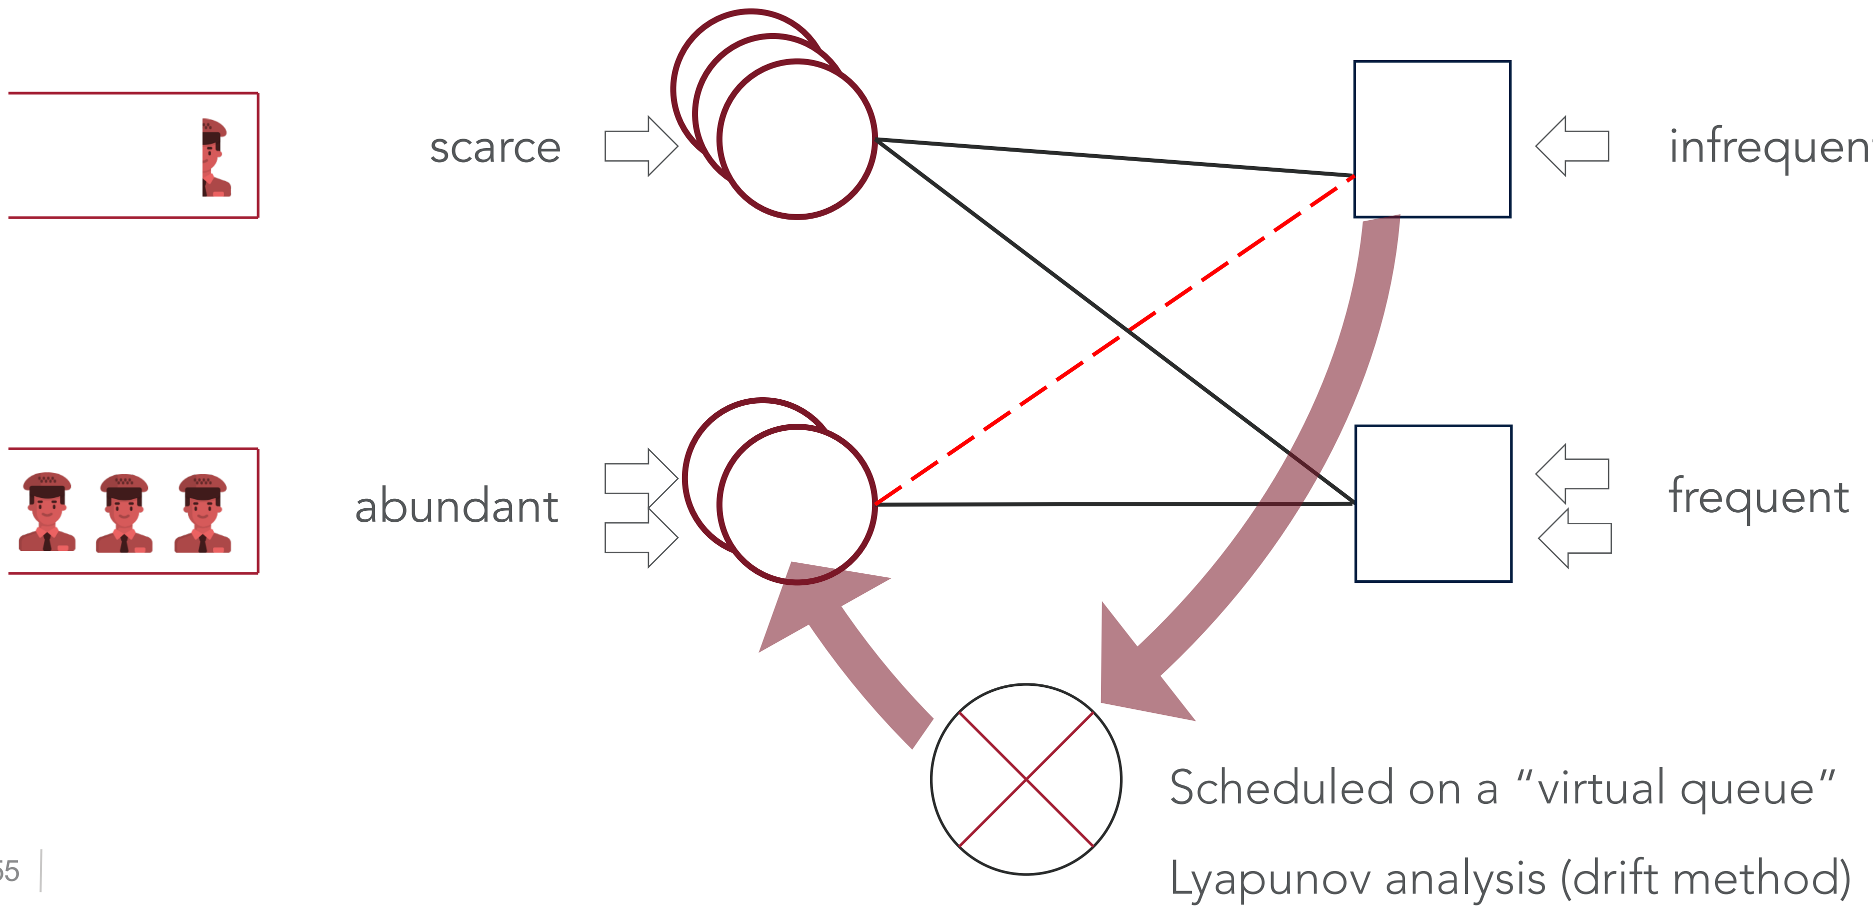

Small networks

Theorem 2 [AAS'24]: There exists an FPTAS for the **bi-criteria dynamic matching problem** for small networks $n \leq \Upsilon$. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(\epsilon^{-\Upsilon} \cdot |\mathcal{I}|^{O(1)}\right)$.

Euclidean networks (main result)

Theorem 3 [AAS'24]: There exists an FPTAS for the bi-criteria dynamic matching problem for Euclidean networks. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(g(\epsilon, d) \cdot |\mathcal{I}|^{O(1)}\right)$.

Euclidean networks (main result)

Theorem 3 [AAS'24]: There exists an FPTAS for the bi-criteria dynamic matching problem for Euclidean networks. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(g(\epsilon, d) \cdot |\mathcal{I}|^{O(1)}\right)$.

- Euclidean networks: embedded in fixed-dimensional Euclidean space

Euclidean networks (main result)

- Euclidean networks: embedded in fixed-dimensional Euclidean space

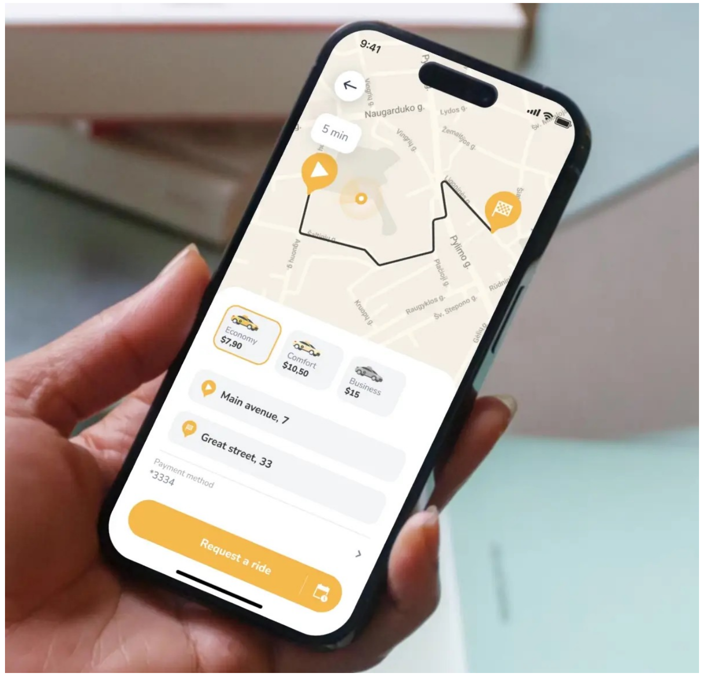

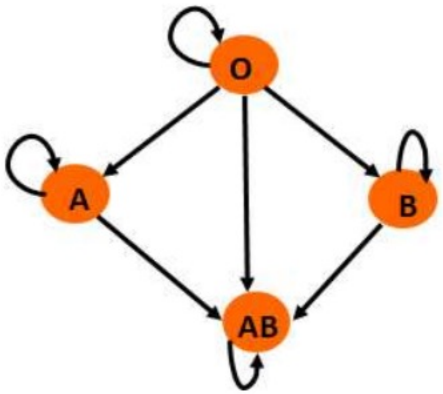

Kidney transplants: ABO-compatibility
[AR, '21]

Euclidean networks (main result)

Theorem 3 [AAS'24]: There exists an FPTAS for the bi-criteria dynamic matching problem for Euclidean networks. For each $\epsilon \in (0,1)$, we compute a $(1+\epsilon)$-approximate policy in time $O\left(g(\epsilon,d) \cdot |\mathcal{I}|^{O(1)}\right)$.

- Euclidean graphs: embedded in fixed-dimensional Euclidean space
- Dimension can be "small": ride-hailing (d~2-4), kidney exchange (d~10)

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS
Small network	-	near-optimal FPTAS

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS
Small network	-	near-optimal FPTAS
Spatial network	3-approx. (metric) [AS, '22]	near-optimal FPTAS (Euclidean)

rare in the matching literature ☺

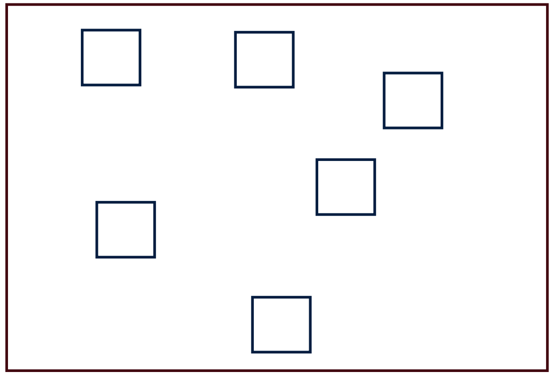

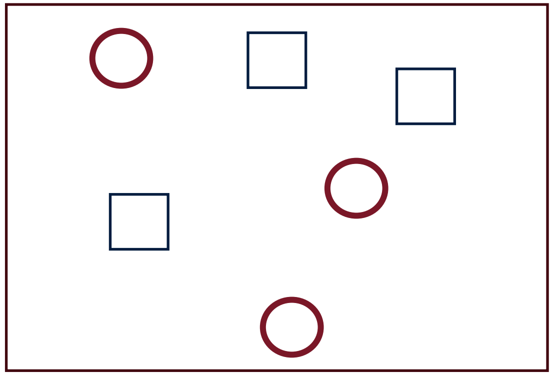

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS
Small network	-	near-optimal FPTAS
Cost network	3-approx. (metric) [AS, '22]	near-optimal PTAS (Euclidean)

	Competitive Ratio	Approximation Ratio
Reward network	$(1-1/\sqrt{e})\approx 0.393$ [PW,'24]	(1-1/e)-approx. [AS, '22]

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS
Small network	-	near-optimal FPTAS
Cost network	3-approx. (metric) [AS, '22]	near-optimal PTAS (Euclidean)

	Competitive Ratio	Approximation Ratio
Reward network	$(1-1/\sqrt{e}) \approx 0.393$ [PW,'24]	(1-1/e)-approx.

Our contributions

	Static	Adaptive
Single queue	0.656-approx. [KSSW, '22]	near-optimal FPTAS
Small network	-	near-optimal FPTAS
Cost network	3-approx. (metric) [AS, '22]	near-optimal PTAS (Euclidean)

	Competitive Ratio	Approximation Ratio
Reward network	$(1-1/\sqrt{e})+\epsilon\approx0.393+\epsilon$	$(1-1/e+\epsilon)$-approx.

Amanihamedani, Aouad, Pollner, and Saberi ['24]

Take-aways

- Dynamic matching with abandonment for thick/thin markets (⚠️ no scaling)
- Surprising tractability: Euclidian networks, small networks
- LPs for adaptive policies: a hybrid LP framework
- Simpler policies? In follow-up work, more fine-grained analysis of correlations

Thank you, questions?

maouad@mit.edu

Soon on arXiv

Appendix

Open questions

- o Breaching 1-1/e for (single-unit) matching in AS [‘22]?
- o The hardness of approximation bounds?
- o Cost-minimization is hard…
  - Heterogeneous server patience with a single customer?
  - Approximations for matching rewards - waiting costs?

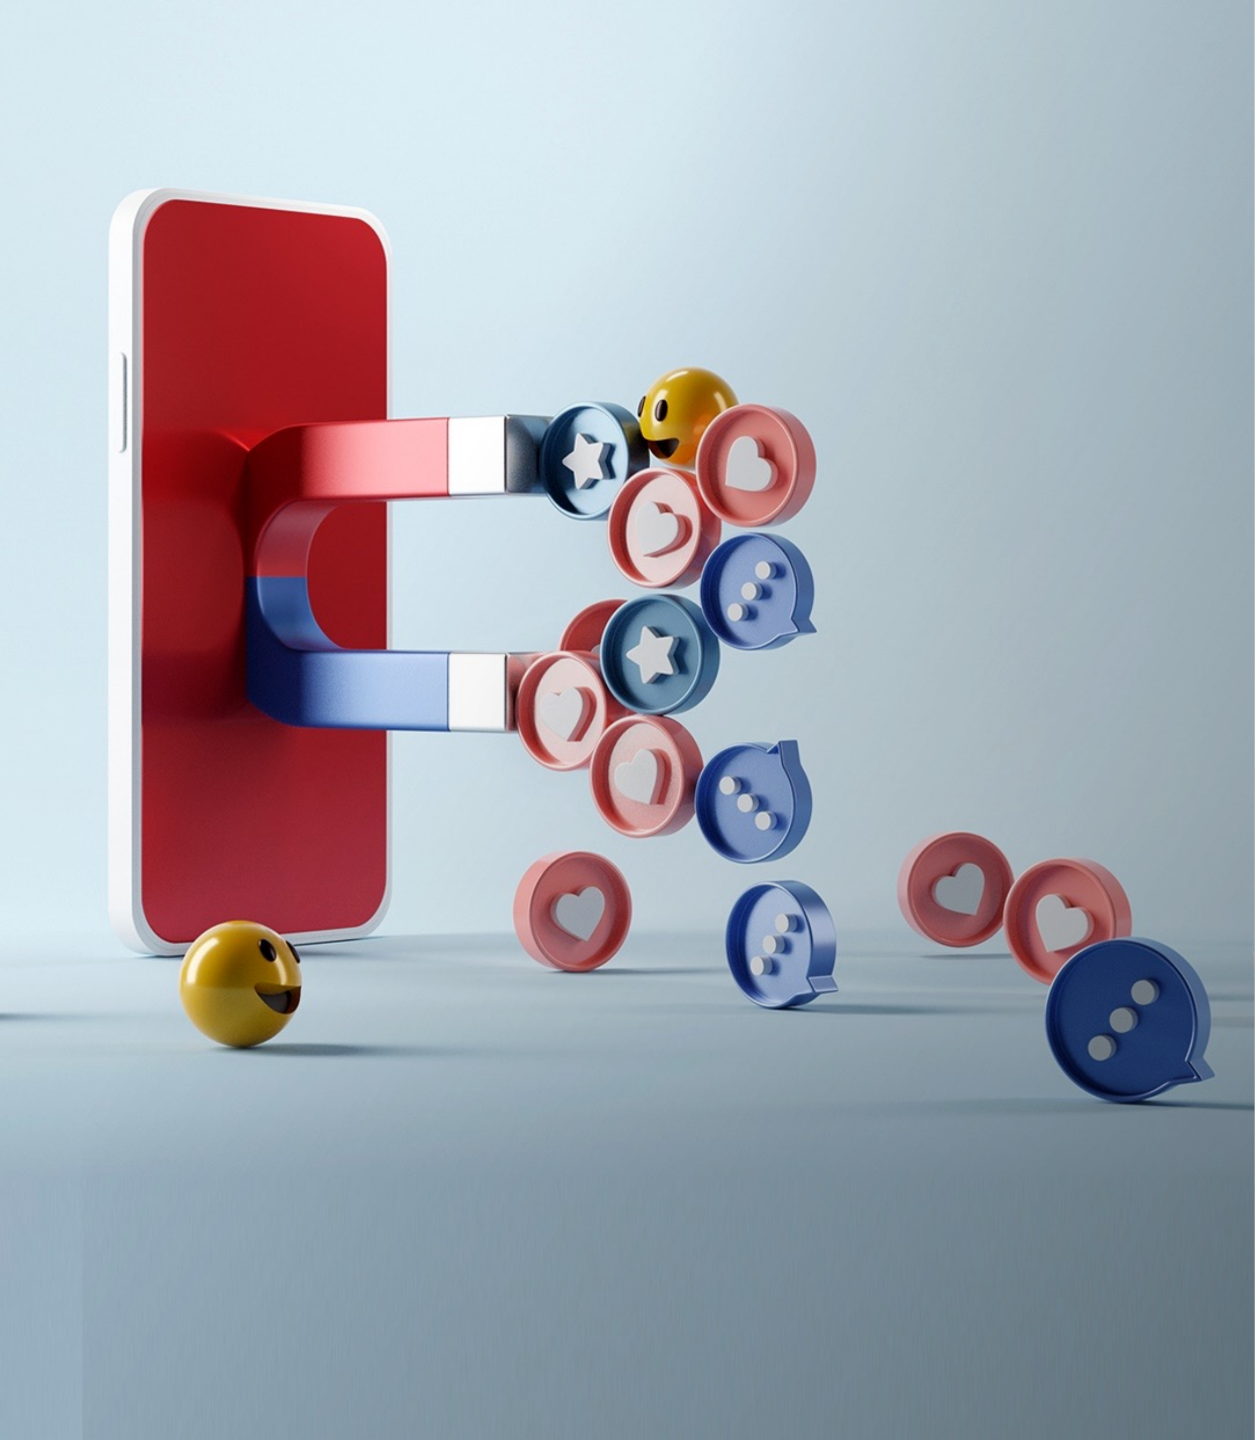

PROOF OUTLINE THEOREM 3

Step 1 (outline): Localized matching

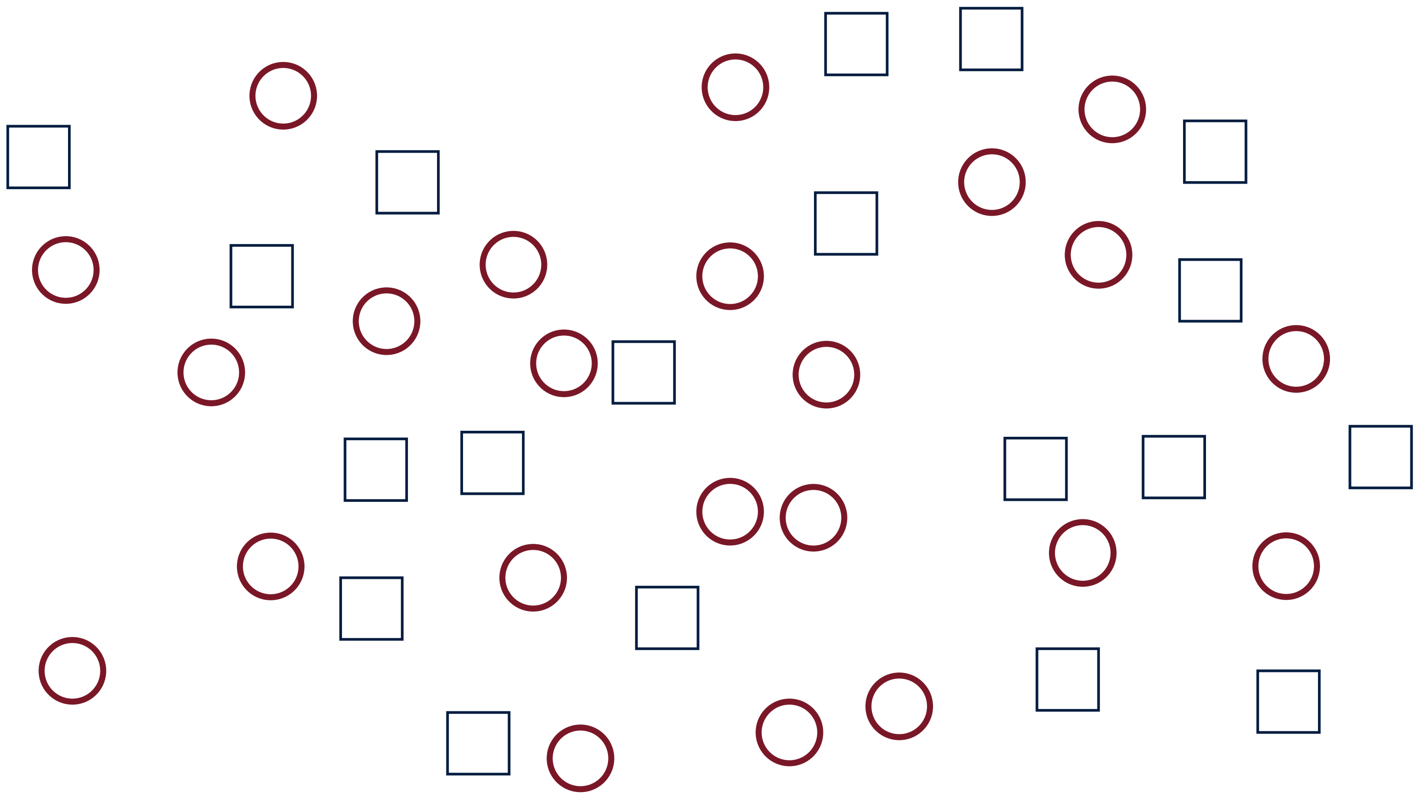

Step 1 (outline): Localized matching

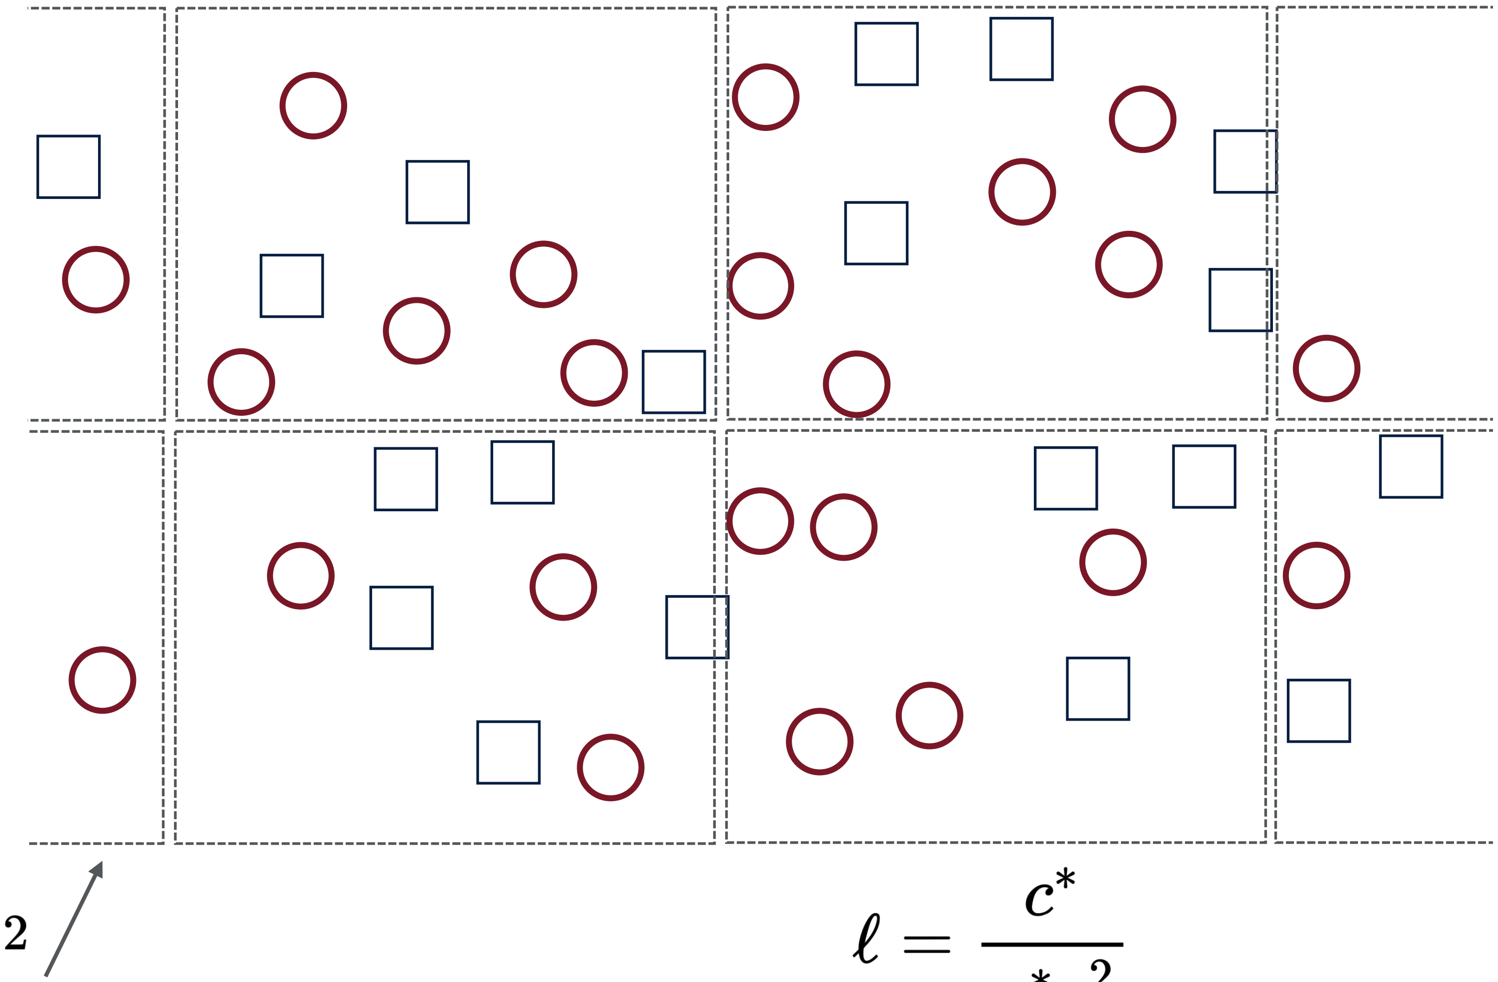

$\vec{\delta} \sim U[0, \ell]^2$

$$\ell = \frac{c^*}{\tau^* \varepsilon^2}$$

Step 1 (outline): Localized matching

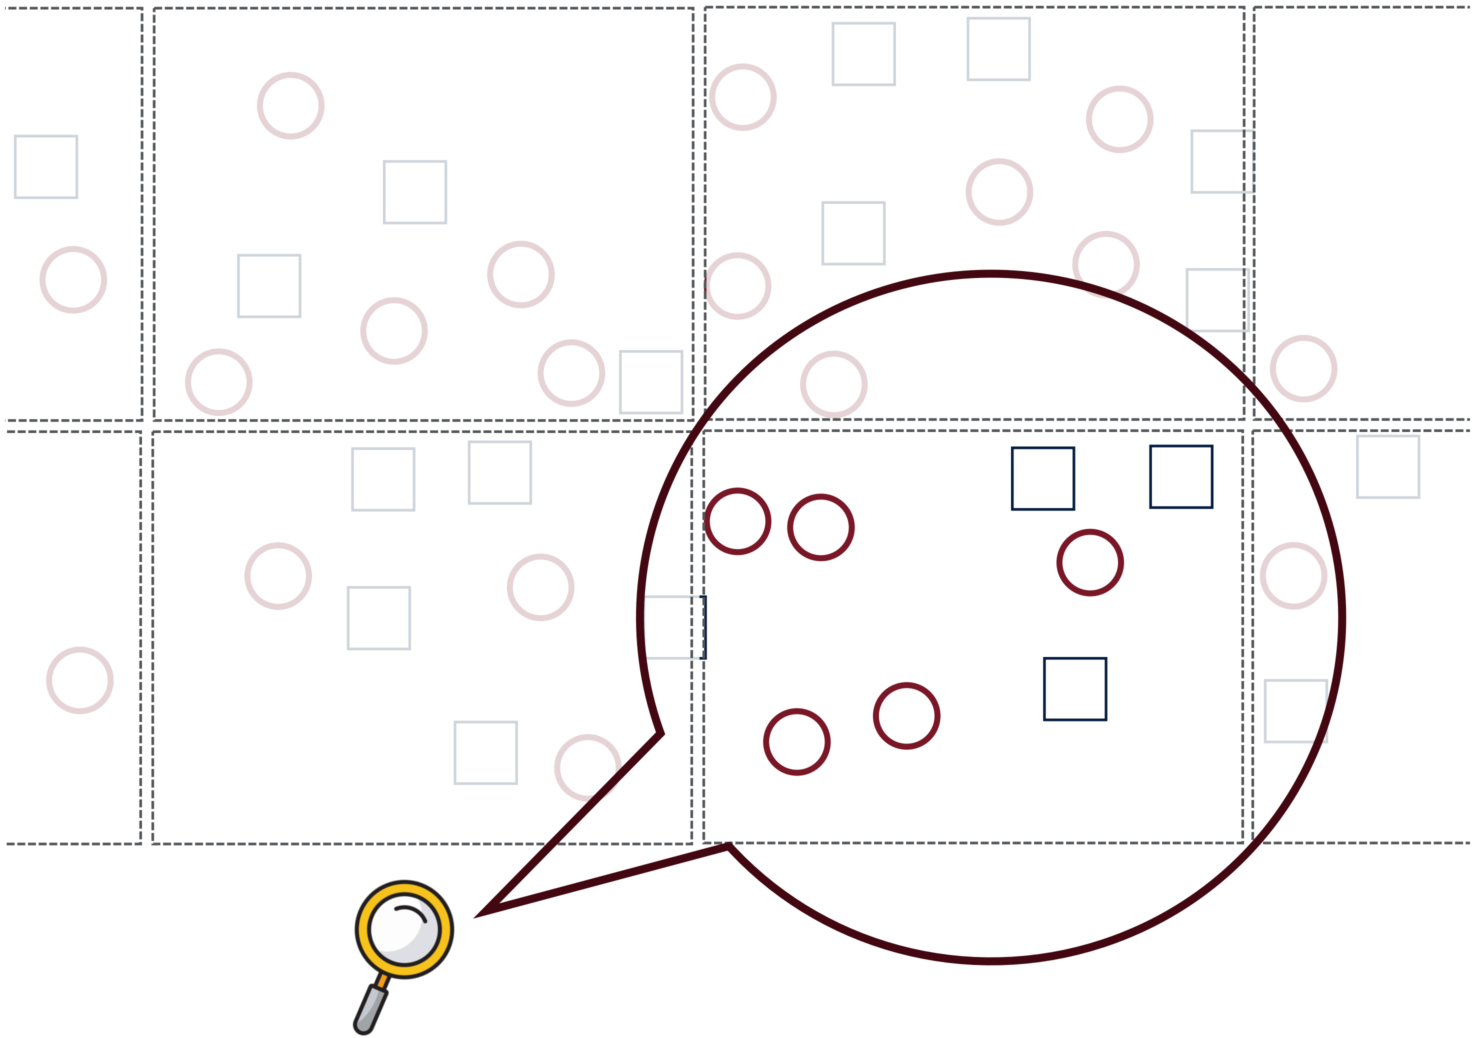

Step 1 (outline): Localized matching

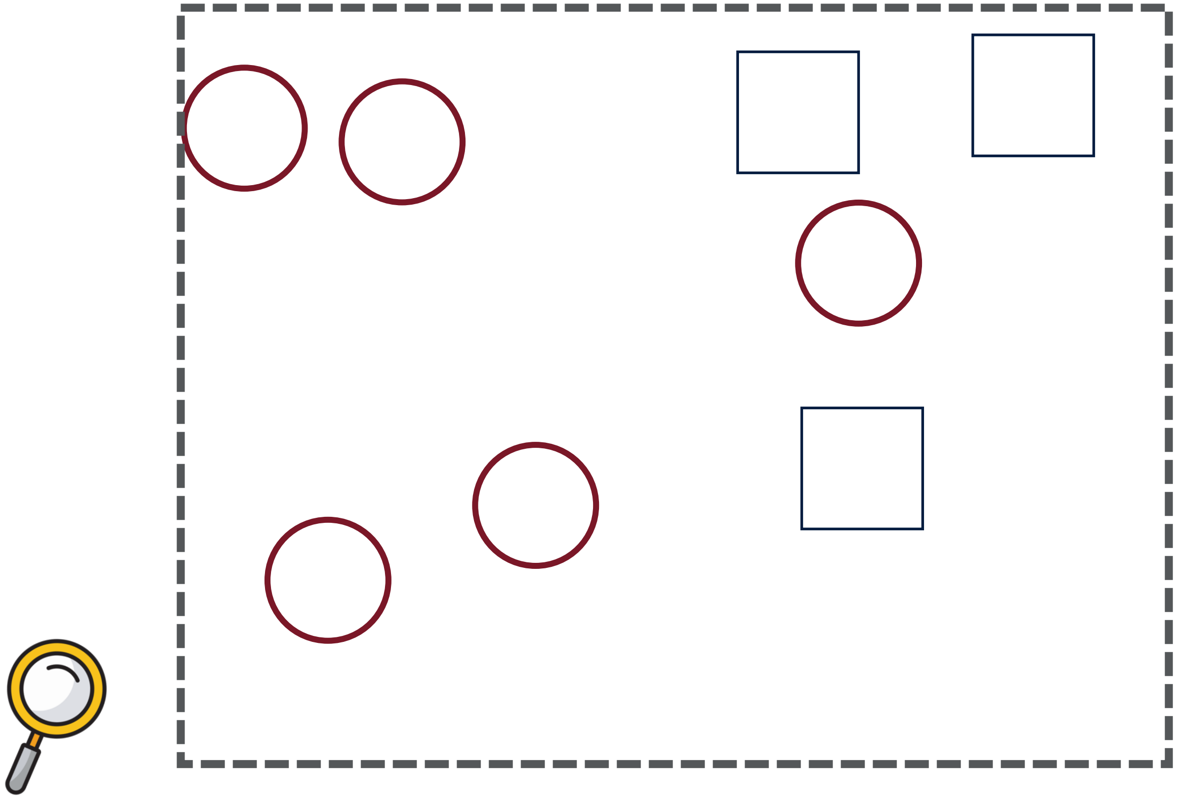

Step 2 (outline): Reduction to DLP via clustering

Cluster servers and customers: Each cell has at most $\frac{1}{\epsilon}$ types (≈ small # types)

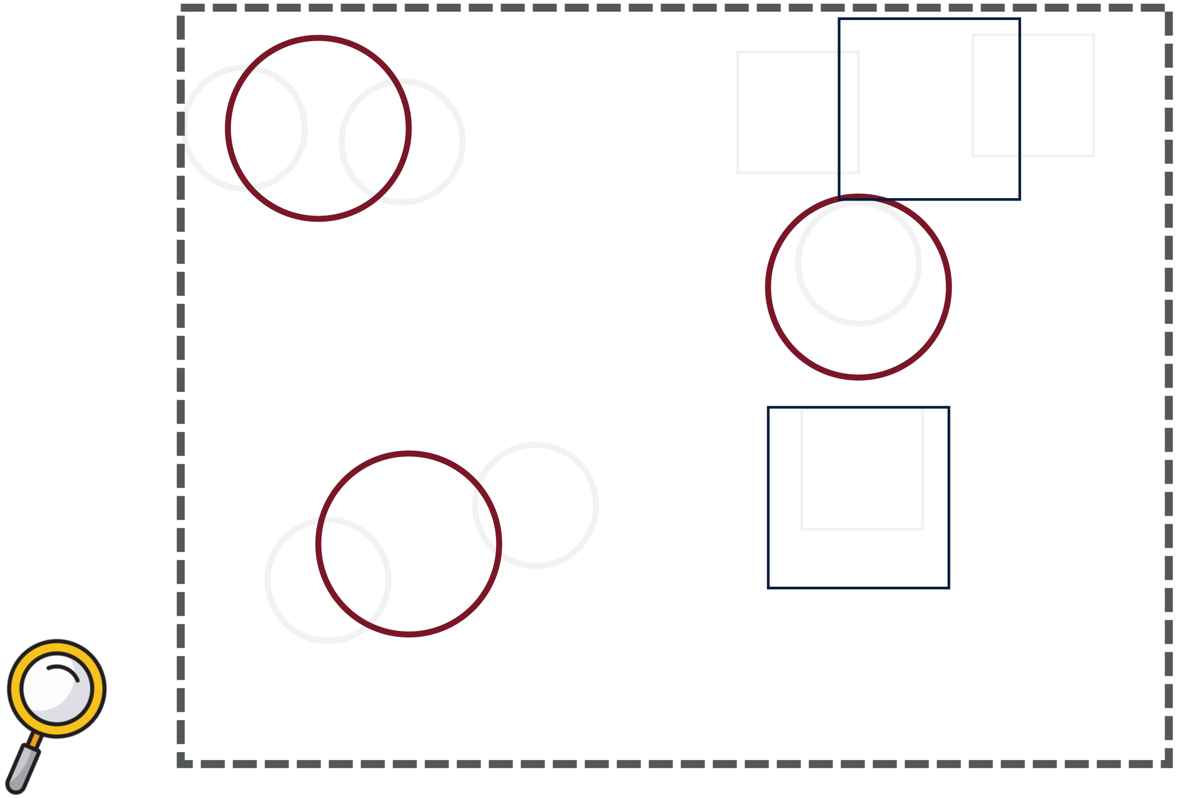

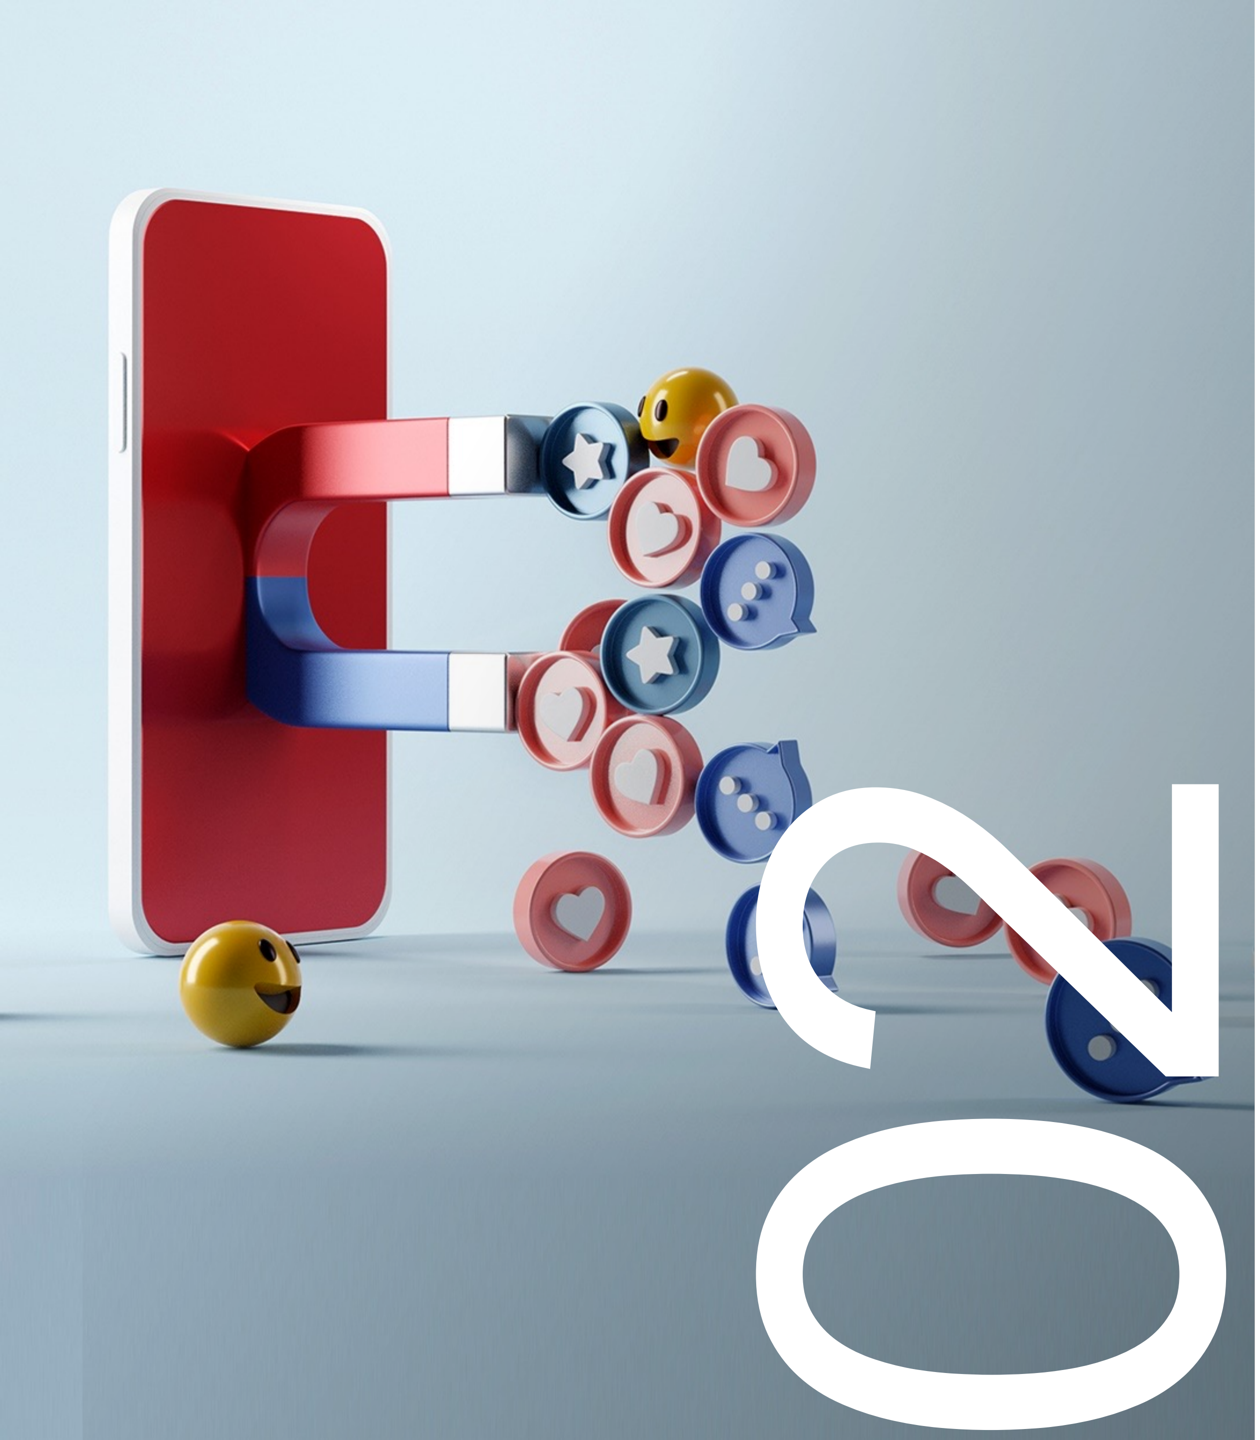

02

MULTIVARIATE DLP

Tighter relaxations: Dynamic-LP

- **PASTA property [Wolff '82]:** the arrivals are independent of the state

Tighter relaxations: Dynamic-LP

- **PASTA property [Wolff '82]:** the arrivals are independent of the state
- **Decision variable** $x_{i,S}^{q}$: stationary probability that ***q* servers** of **type *i*** are waiting and the optimal policy is **about to match them with a customer in *S***

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S} \gamma_j \cdot c_{i,j} \cdot x^q_{i,S}$$

$$\text{s.t.} \quad (x^q_{i,S})_{S,q} \in \mathcal{P}_i$$

queue-adapted process

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x_{i,S}^q$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x} \geq 0} \quad \sum_{i,S,q} \sum_{j \in S} \gamma_j \cdot c_{i,j} \cdot x_{i,S}^q$$

$$\text{s.t.} \quad (x_{i,S}^q)_{S,q} \in \mathcal{P}_i$$

queue-adapted process

$(\mathcal{P}_i)$

- detail balance
- probas-sum-to-1

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that q servers of type i are waiting and the optimal policy is about to match them with a customer in S

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S}\gamma_j \cdot c_{i,j}\cdot x^q_{i,S}$$

$$\text{s.t.} \quad (x^q_{i,S})_{S,q}\in \mathcal{P}_i$$

queue-adapted process

$$(\mathcal{P}_i)\begin{cases} \lambda_i \cdot \sum_S x^{q-1}_{i,S} = \sum_S x^q_{i,S}\cdot(\gamma(S)+\mu_i\cdot q) \quad \forall S,q \\ \sum_{S,q} x^q_{i,S} = 1 \end{cases}$$

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x_{i,S}^q$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S}\gamma_j \cdot c_{i,j}\cdot x_{i,S}^q$$

$$\text{s.t.} \quad (x_{i,S}^q)_{S,q} \in \mathcal{P}_i$$

no contention

$$(\mathcal{P}_i)\begin{cases} \lambda_i \cdot \sum_S x_{i,S}^{q-1} = \sum_S x_{i,S}^q \cdot (\gamma(S) + \mu_i \cdot q) \quad \forall S, q \\ \sum_{S,q} x_{i,S}^q = 1 \end{cases}$$

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S}\gamma_j\cdot c_{i,j}\cdot x^q_{i,S}$$

$$\text{s.t.} \quad (x^q_{i,S})_{S,q}\in\mathcal{P}_i$$

$$\sum_{i,q}\sum_{\substack{S\subseteq[m]:\\ j\in S}} x^q_{i,S}\leq 1\ , \quad \forall j$$

$$(\mathcal{P}_i)\begin{cases}\lambda_i\cdot\sum_{S}x^{q-1}_{i,S}=\sum_{S}x^q_{i,S}\cdot(\gamma(S)+\mu_i\cdot q) & \forall S,q\\ \sum_{S,q}x^q_{i,S}=1\end{cases}$$

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S}\gamma_j \cdot c_{i,j}\cdot x^q_{i,S}$$

$$\text{s.t.} \quad (x^q_{i,S})_{S,q} \in \mathcal{P}_i$$

$$\sum_{i,q}\sum_{\substack{S\subseteq[m]:\\ j\in S}} x^q_{i,S} \leq 1\ , \quad \forall j$$

throughput target .

$$(\mathcal{P}_i)\begin{cases}\lambda_i \cdot \sum_S x^{q-1}_{i,S} = \sum_S x^q_{i,S}\cdot(\gamma(S)+\mu_i\cdot q) \quad \forall S,q \\ \sum_{S,q} x^q_{i,S} = 1\end{cases}$$

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \sum_{i,S,q} \sum_{j\in S} \gamma_j \cdot c_{i,j} \cdot x^q_{i,S}$$

s.t.

$$(x^q_{i,S})_{S,q} \in \mathcal{P}_i$$

$$\sum_{i,q} \sum_{\substack{S\subseteq[m]:\\ j\in S}} x^q_{i,S} \leq 1 \ , \quad \forall j$$

$$\sum_{i,q} \sum_{S\subseteq[m]} \gamma(S) \cdot x^q_{i,S} \geq \tau^* \ .$$

$$(\mathcal{P}_i) \begin{cases} \lambda_i \cdot \sum_{S} x^{q-1}_{i,S} = \sum_{S} x^q_{i,S} \cdot (\gamma(S) + \mu_i \cdot q) & \forall S, q \\ \sum_{S,q} x^q_{i,S} = 1 \end{cases}$$

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x_{i,S}^q$: stationary probability that q servers of type i are waiting and the optimal policy is about to match them with a customer in S

(DLP)

$$\min_{\boldsymbol{x}\geq 0} \quad \sum_{i,S,q}\sum_{j\in S}\gamma_j\cdot c_{i,j}\cdot x_{i,S}^q$$

$$\text{s.t.} \quad (x_{i,S}^q)_{S,q}\in\mathcal{P}_i$$

$$\sum_{i,q}\sum_{\substack{S\subseteq[m]:\\ j\in S}} x_{i,S}^q \leq 1\ , \quad \forall j$$

$$\sum_{i,q}\sum_{S\subseteq[m]}\gamma(S)\cdot x_{i,S}^q \geq \tau^*\ .$$

$$(\mathcal{P}_i)\begin{cases}\lambda_i\cdot\sum_S x_{i,S}^{q-1} = \sum_S x_{i,S}^q\cdot(\gamma(S)+\mu_i\cdot q) \quad \forall S,q\\ \sum_{S,q} x_{i,S}^q = 1\end{cases}$$

+ some other polymatroid constraints! ☺

Tighter relaxations: Dynamic-LP

- PASTA property [Wolff '82]: the arrivals are independent of the state
- Decision variable $x^q_{i,S}$: stationary probability that *q* servers of type *i* are waiting and the optimal policy is about to match them with a customer in *S*

$$\text{(DLP)} \quad \begin{aligned} \min_{\boldsymbol{x}\geq 0} \quad & \sum_{i,S,q}\sum_{j\in S}\gamma_j \cdot c_{i,j}\cdot x^q_{i,S} \\ \text{s.t.} \quad & (x^q_{i,S})_{S,q} \in \mathcal{P}_i \\ & \sum_{i,q}\sum_{\substack{S\subseteq[m]:\\ j\in S}} x^q_{i,S} \leq 1\ , \quad \forall j \\ & \sum_{i,q}\sum_{S\subseteq[m]} \gamma(S)\cdot x^q_{i,S} \geq \tau^*\ . \end{aligned}$$

$$(\mathcal{P}_i) \begin{cases} \lambda_i \cdot \sum_S x^{q-1}_{i,S} = \sum_S x^q_{i,S}\cdot(\gamma(S)+\mu_i\cdot q) \quad \forall S,q \\ \sum_{S,q} x^q_{i,S} = 1 \end{cases}$$

Fact: DLP is tighter [KSSW'22] [YV'24]

Dual formulation and properties

Romeijn et al. ['92] (transversality)

$$\alpha_i + \sum_{j \in S} \gamma_j \cdot \left(\delta_i^q + \theta - c_{i,j} - \frac{\beta_j}{\gamma_j} \right) \leq \lambda_i \cdot \delta_i^{q+1} - \mu_i \cdot q \cdot \delta_i^q \qquad \forall i, S, q$$

$$\alpha_i \leq \lambda_i \cdot \delta_i^1 \ , \quad \delta_i^q \leq 0 \ , \quad \theta, \beta_j \geq 0$$

Dual formulation and properties

Lemma [AAS'24]: DLP dual describes weakly coupled average-cost MDPs, where an optimal policy is characterized by queue length-dependent thresholds δ_i^{q*} .

Dual formulation and properties

Lemma [AAS'24]: DLP dual describes weakly coupled average-cost MDPs, where an optimal policy is characterized by queue length-dependent thresholds δ_i^{q*}.

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by **queue length-dependent thresholds** δ_i^{q*}.

- Matching set: $S_i^{-q} \subseteq S_{i,q}^* \subseteq S_i^{+q}$

- Thresholds on reduced costs:

$$S_i^{-q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* < \delta_i^{q*} \right\} \text{ and } S_i^{+q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* \leq \delta_i^{q*} \right\}$$

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by **queue length-dependent thresholds** δ_i^{q*}.

- Matching set: $S_i^{-q} \subseteq S_{i,q}^* \subseteq S_i^{+q}$
- Thresholds on reduced costs:

$$S_i^{-q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* < \delta_i^{q*} \right\} \text{ and } S_i^{+q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* \leq \delta_i^{q*} \right\}$$

reduced cost reduced cost

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by **queue length-dependent thresholds** δ_i^{q*}.

- Matching set: $S_i^{-q} \subseteq S_{i,q}^{*} \subseteq S_i^{+q}$

- Thresholds on reduced costs:

$$S_i^{-q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* < \delta_i^{q*} \right\} \text{ and } S_i^{+q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* \leq \delta_i^{q*} \right\}$$

thresholds
=marginal costs

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by **queue length-dependent thresholds** δ_i^{q*}.

- Matching set: $S_i^{-q} \subseteq S_{i,q}^* \subseteq S_i^{+q}$

- Thresholds on reduced costs:

$$S_i^{-q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* < \delta_i^{q*} \right\} \text{ and } S_i^{+q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* \leq \delta_i^{q*} \right\}$$

throughput reward

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by **queue length-dependent thresholds** δ_i^{q*}.

- Matching set: $S_i^{-q} \subseteq S_{i,q}^* \subseteq S_i^{+q}$
- Thresholds on reduced costs:

$$S_i^{-q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* < \delta_i^{q*} \right\} \text{ and } S_i^{+q} = \left\{ j \in [m] : c_{i,j} + \frac{\beta_j^*}{\gamma_j} - \theta^* \leq \delta_i^{q*} \right\}$$

shadow price (contention)

Dual formulation and properties

Lemma [AAS'24]: DLP describes weakly coupled average-cost MDPs, where an optimal policy is characterized by queue length-dependent thresholds δ_i^{q*}.

Lemma [AAS'24]: Thresholds $(\delta_i^{q*})_q$ are monotone increasing and concave.

"Reasonable" adaptive policies

reduced costs $c_{i,j} + \beta_j - \theta$

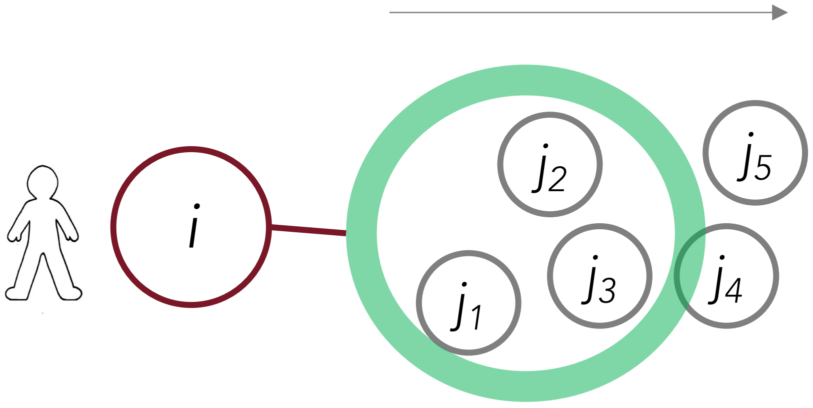

$S(\delta_i^{1*})$

reduced costs

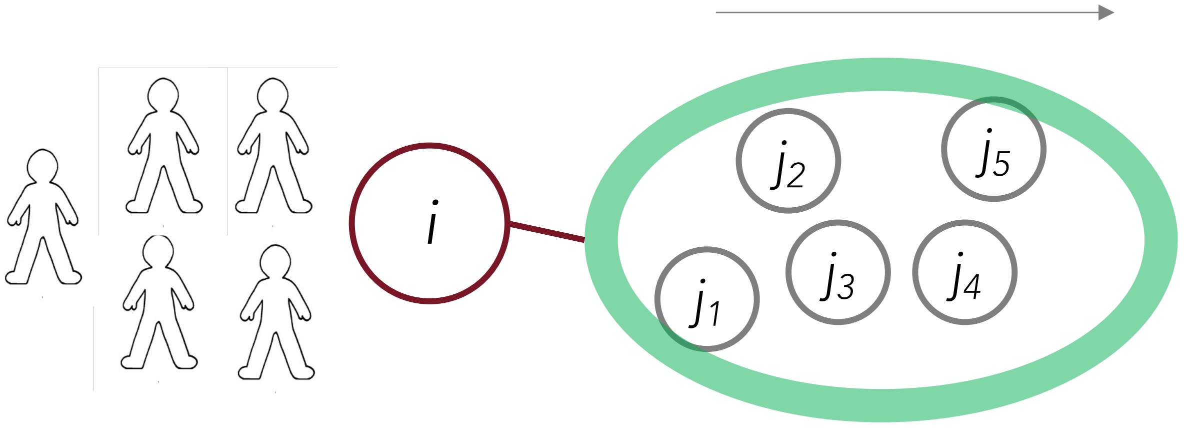

$S(\delta_i^{5*}) \supseteq S(\delta_i^{1*})$

Can we solve (DLP) efficiently?

Can we solve (DLP) efficiently?

Theorem 1 [AAS'24]: There is a fully polynomial-time approximation scheme for D-LP.

Can we solve (DLP) efficiently?

Theorem 1 [AAS'24]: There is a fully polynomial-time approximation scheme for D-LP.

Proof ideas:

1. State space collapse: limited adaptivity
2. Efficient separation oracle: sliding ellipsoid, using both primal & dual

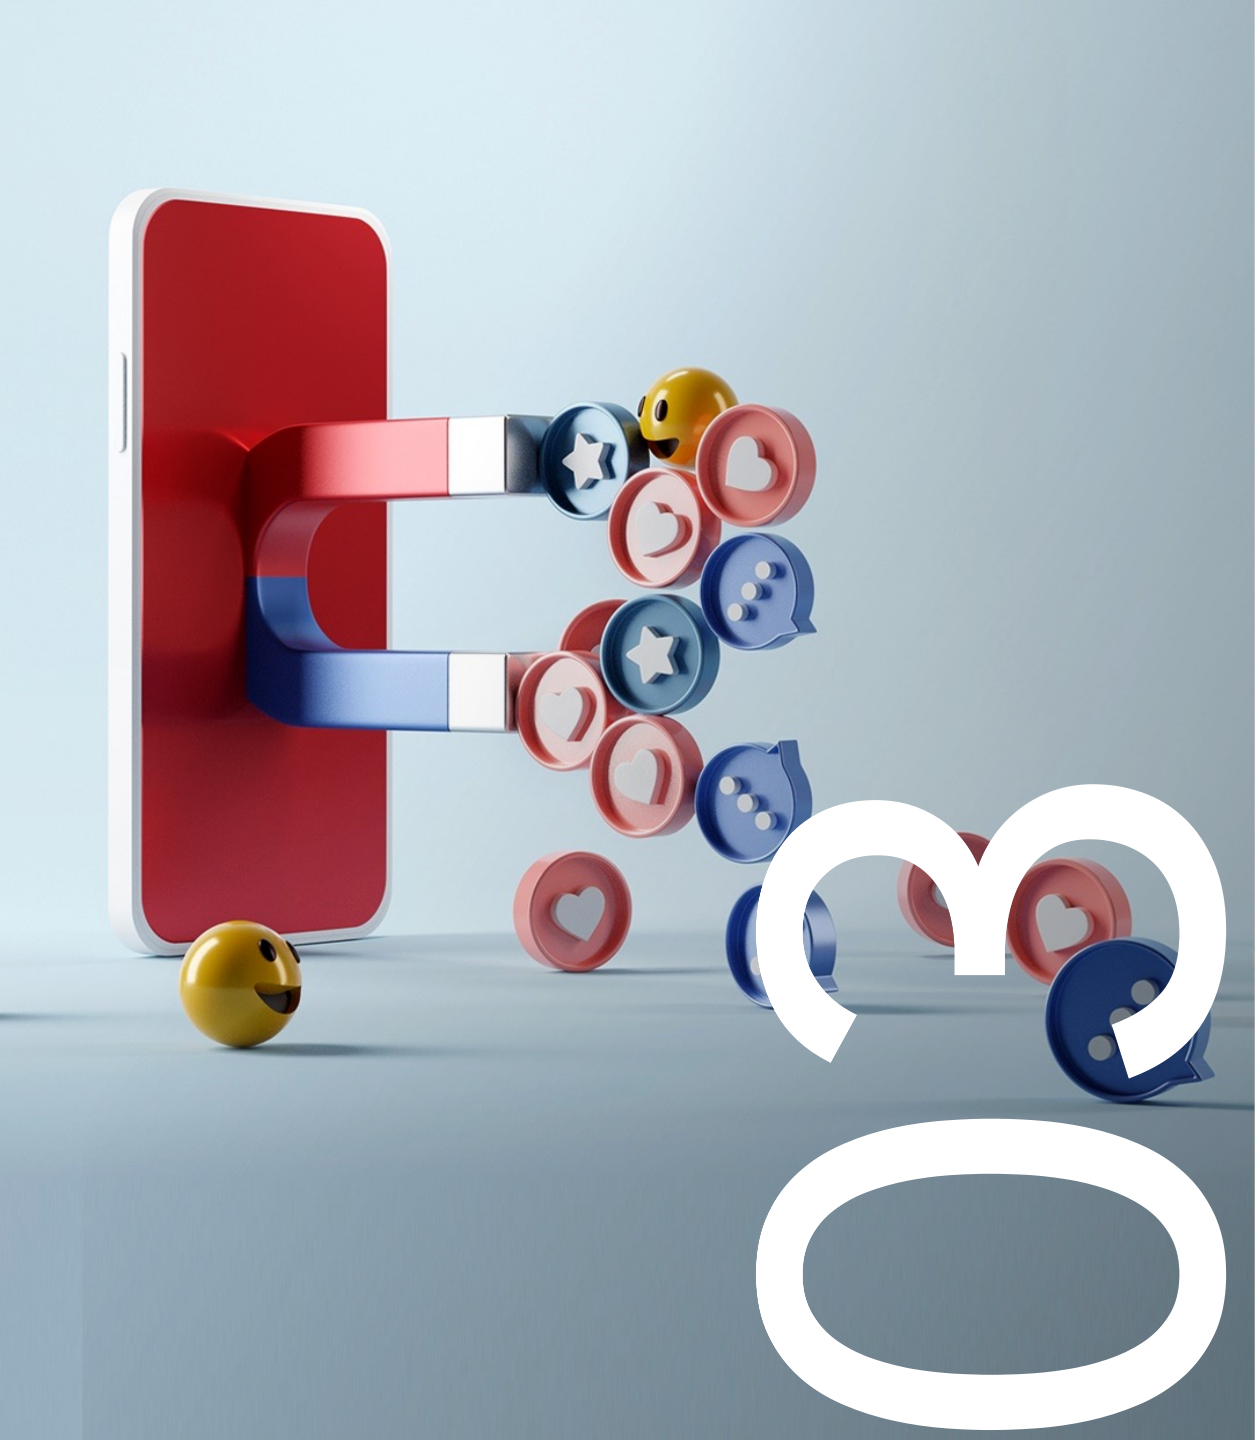

03

TRUNCATION LEMMA

On adaptivity

- **Definition:** A Q-**bounded** solution uses up to Q agents in each queue.

On adaptivity

- **Definition**: A Q-bounded solution uses up to Q agents in each queue.
- An instance is (Q, ϵ)-adapted if any solution's matching rates can be approximated by some Q-bounded solution within factor 1- ϵ.

On adaptivity

- **Definition**: A Q-bounded solution uses up to Q agents in each queue.
- An instance is (Q, ϵ)-adapted if any solution's matching rates can be approximated by some Q-bounded solution within factor 1- ϵ .
- Easy upper bound:

$$Q = O\left(\log\left(\frac{\tau_{\max}}{\epsilon\tau^*}\right) \cdot \left(\max_i \frac{\lambda_i}{\mu_i}\right)\right)$$

On adaptivity

- **Definition**: A Q-bounded solution uses up to Q agents in each queue.
- An instance is (Q, ϵ)-adapted if any solution's matching rates can be approximated by some Q-bounded solution within factor 1- ϵ.
- Easy upper bound:

$$Q = O\left(\log\left(\frac{\tau_{\max}}{\epsilon\tau^*}\right) \cdot \left(\max_i \frac{\lambda_i}{\mu_i}\right)\right)$$

Lemma [AAS'24]: Every instance is $Q = O\left(\log\left(\frac{\tau_{\max}}{\epsilon\tau^*}\right)\right)$, ϵ -adapted.

Step 1 - Sensitivity

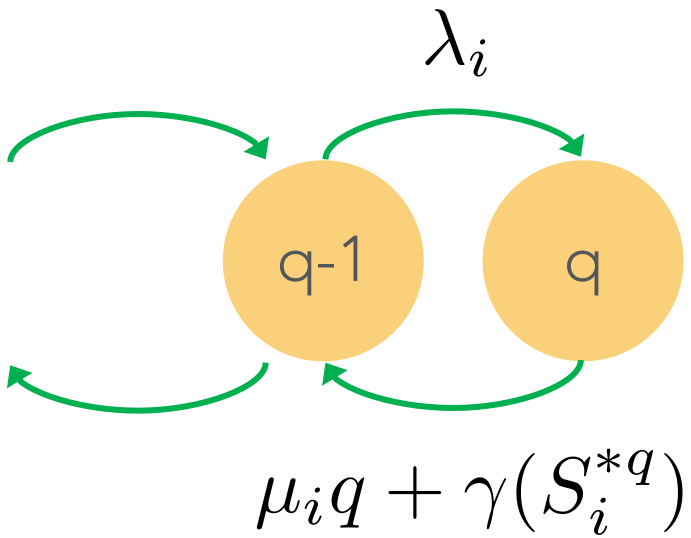

Sensitivity: we can inflate, deflate each rates by $O(\epsilon)$-fraction with small loss.

Step 2 – Distribution design problem

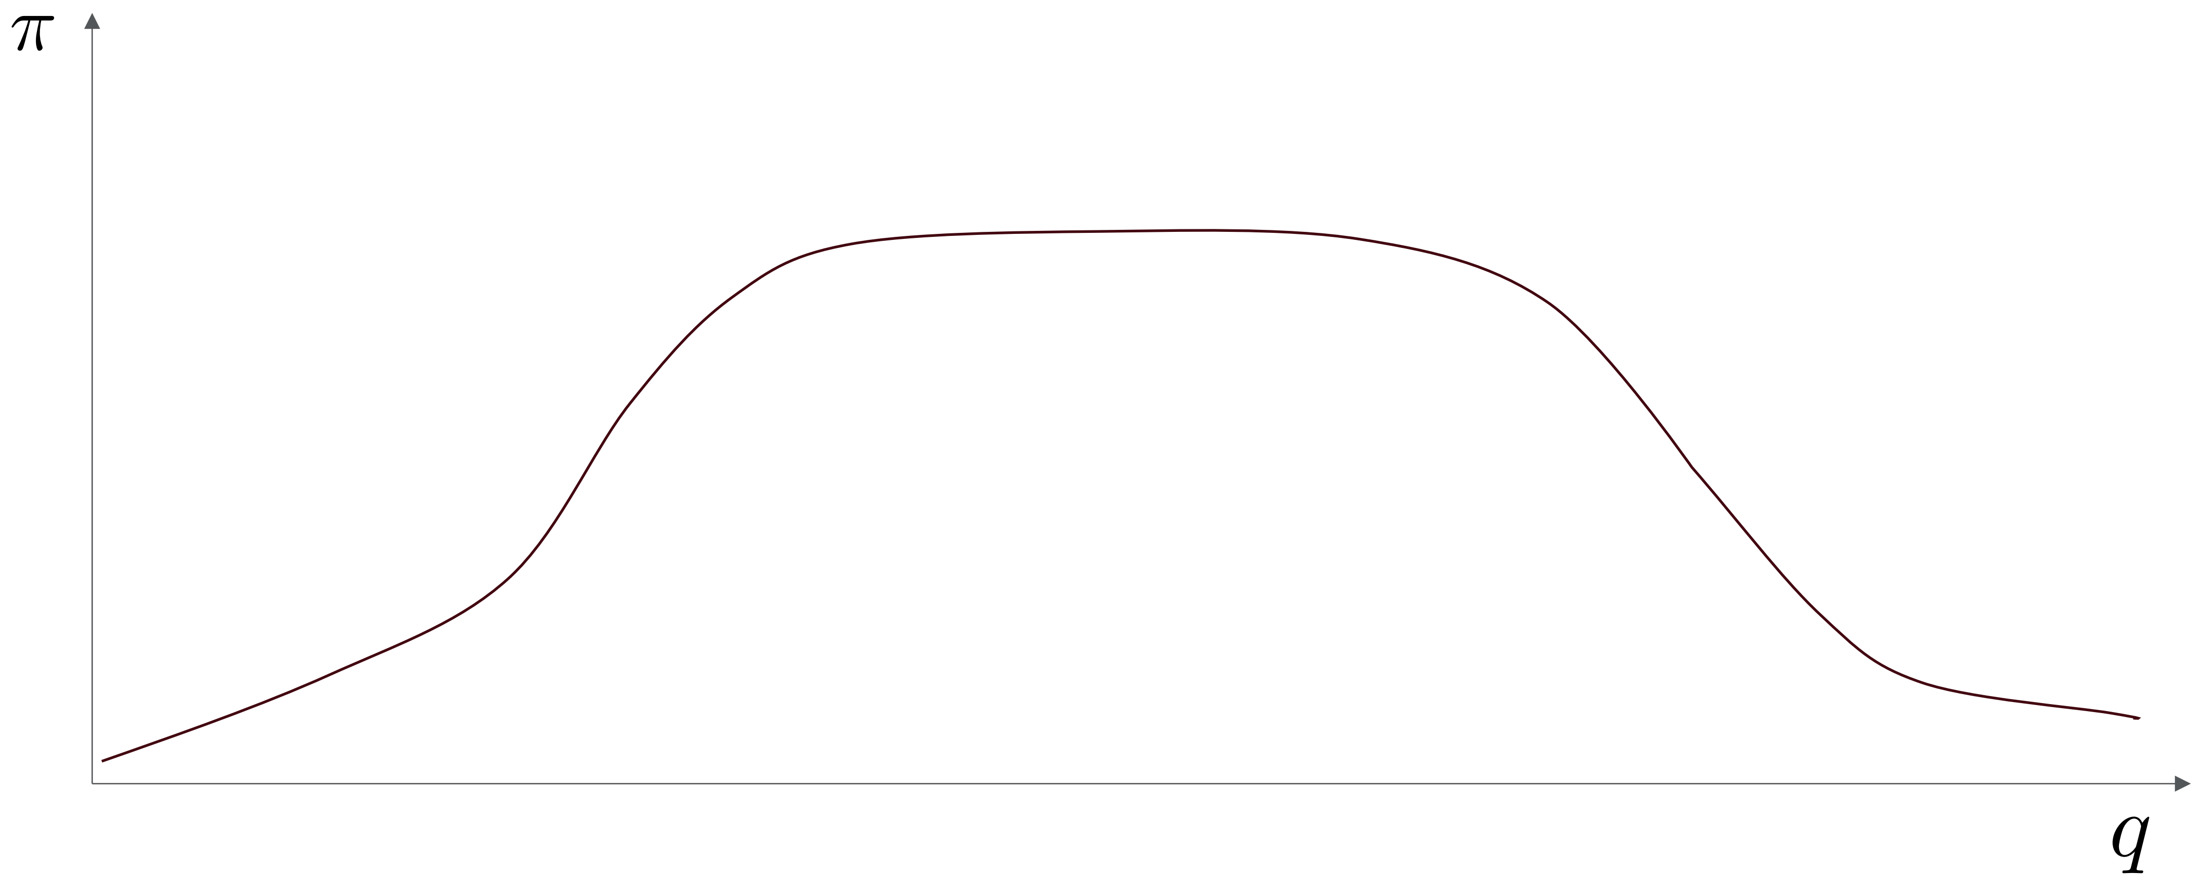

Step 2 – Distribution design problem

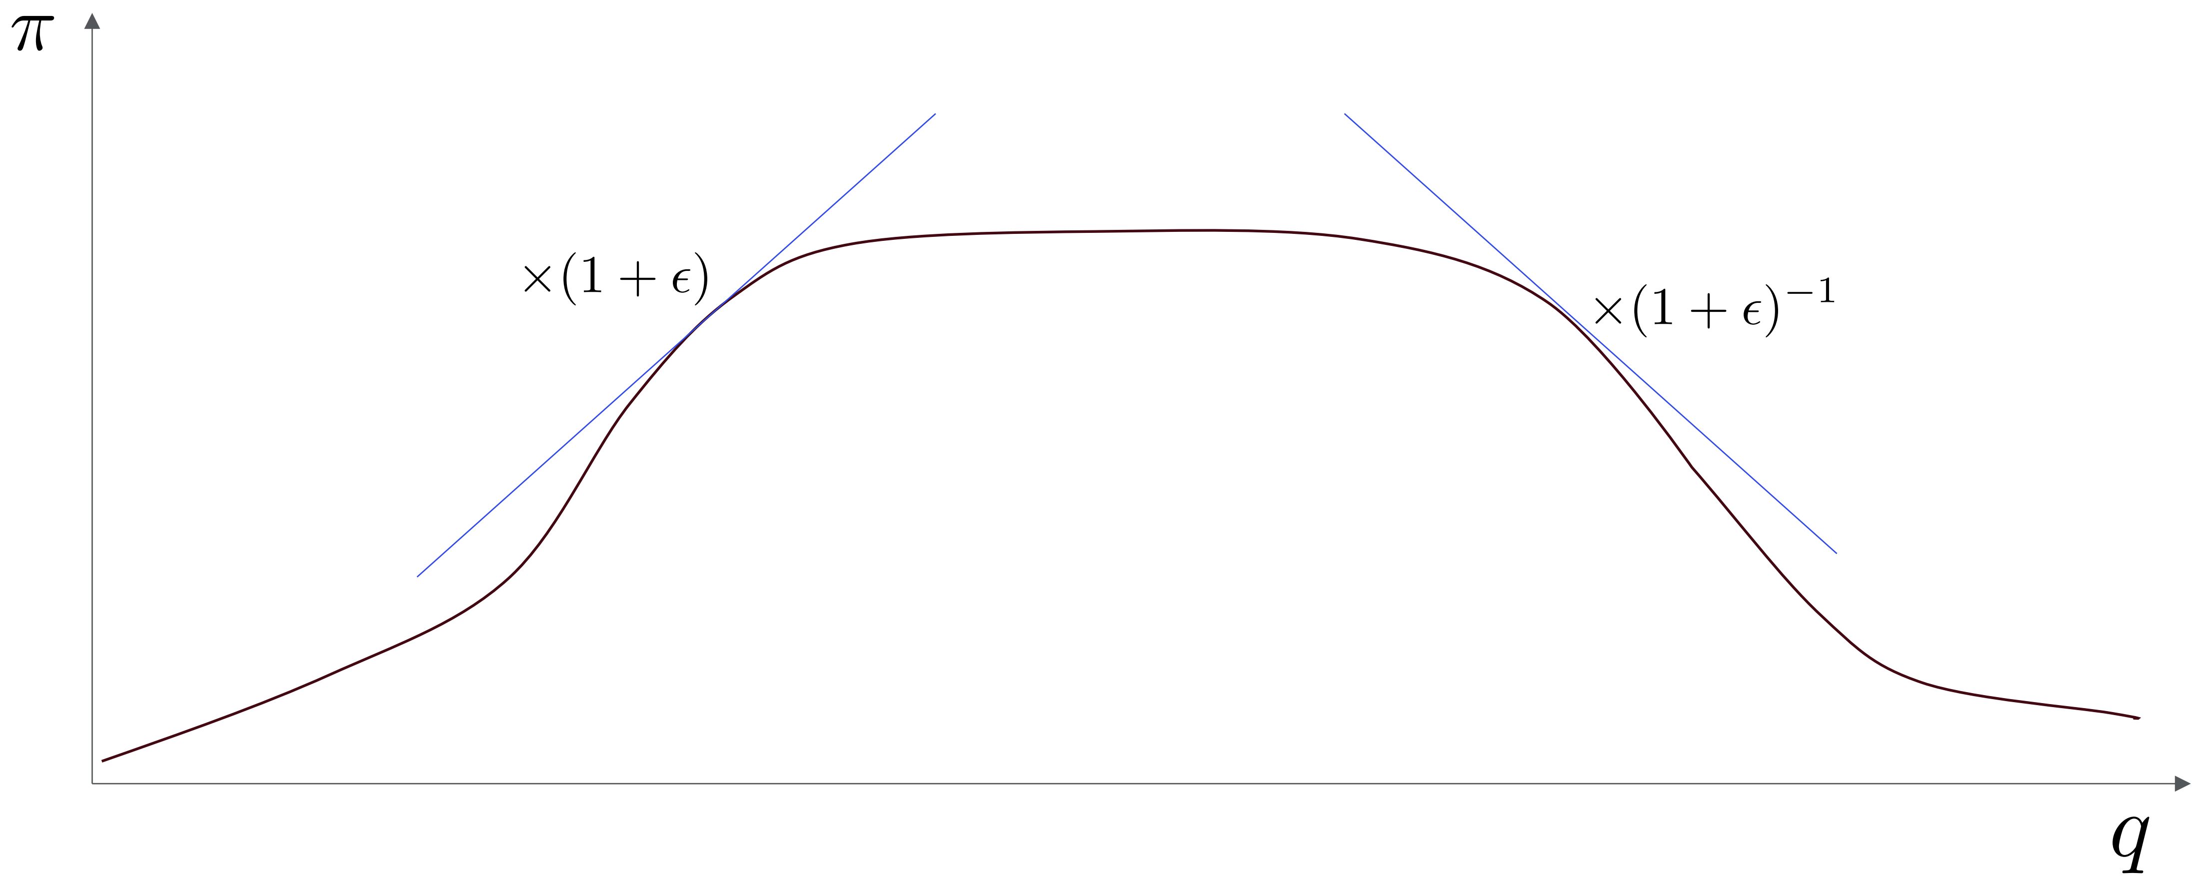

Step 2 – Distribution design problem

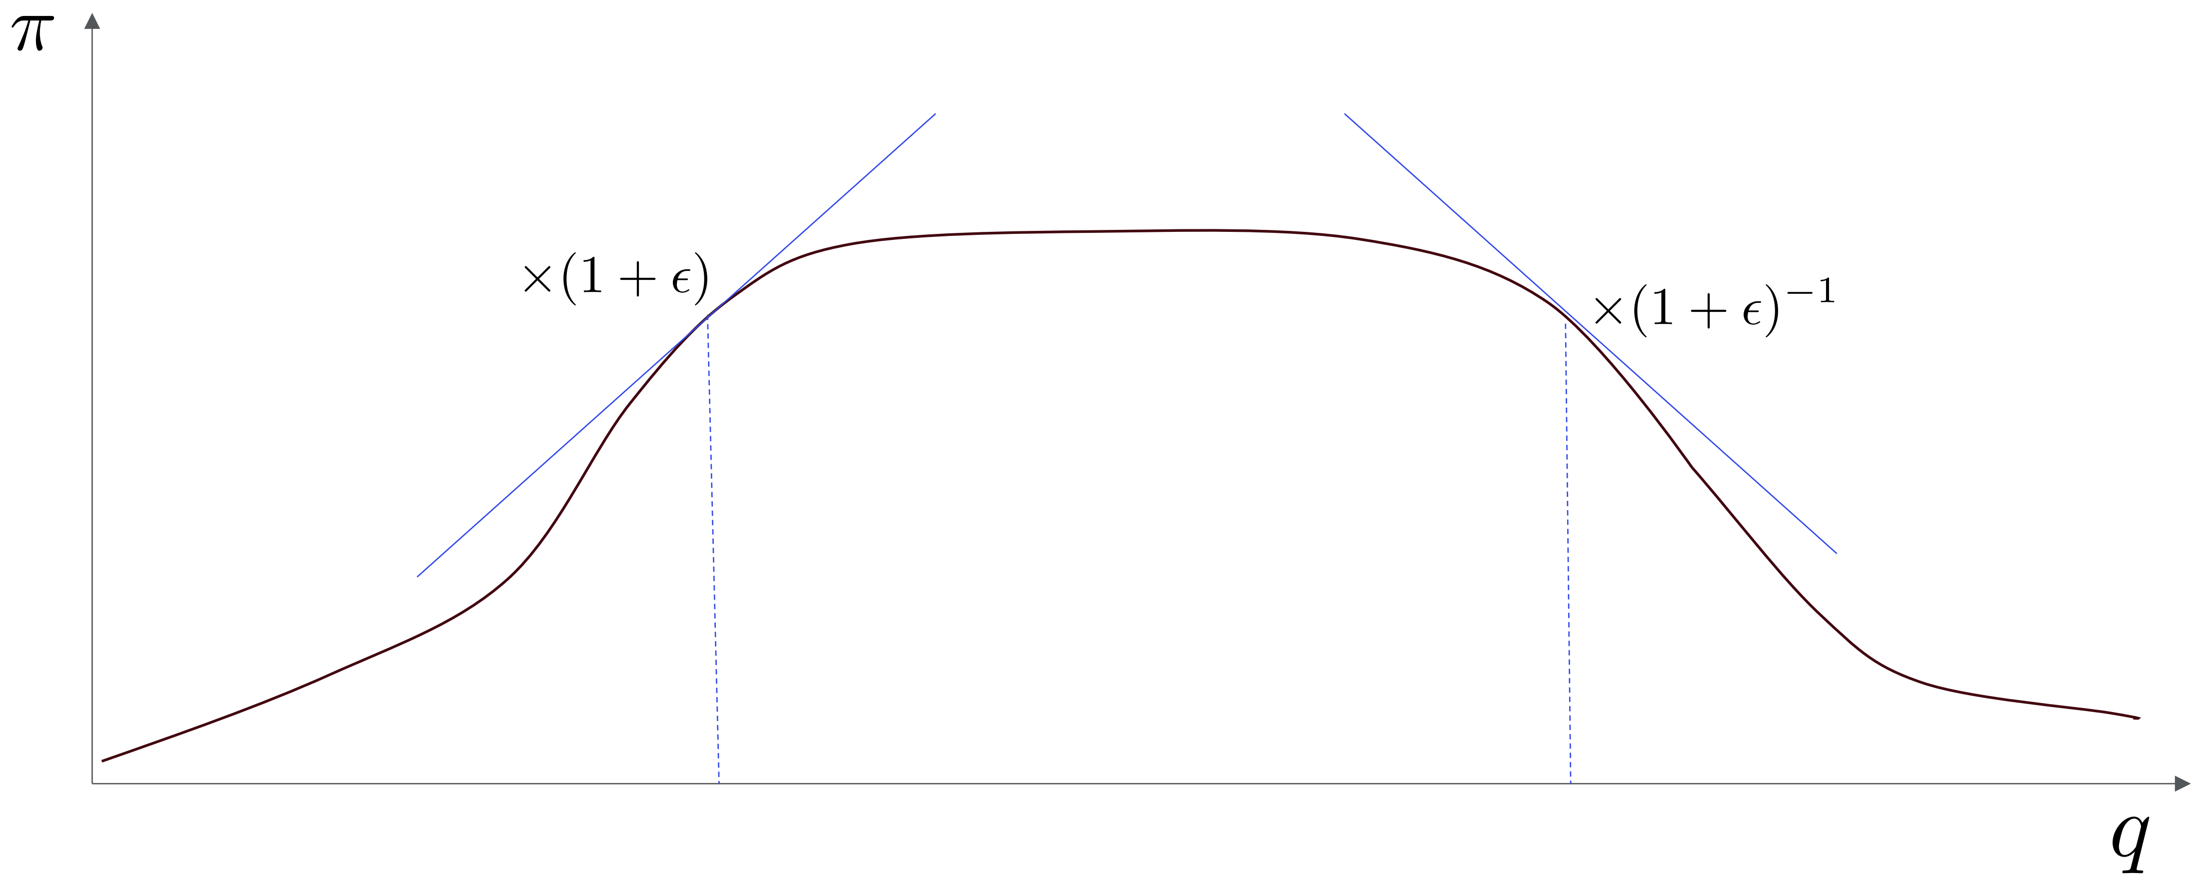

Step 2 – Distribution design problem

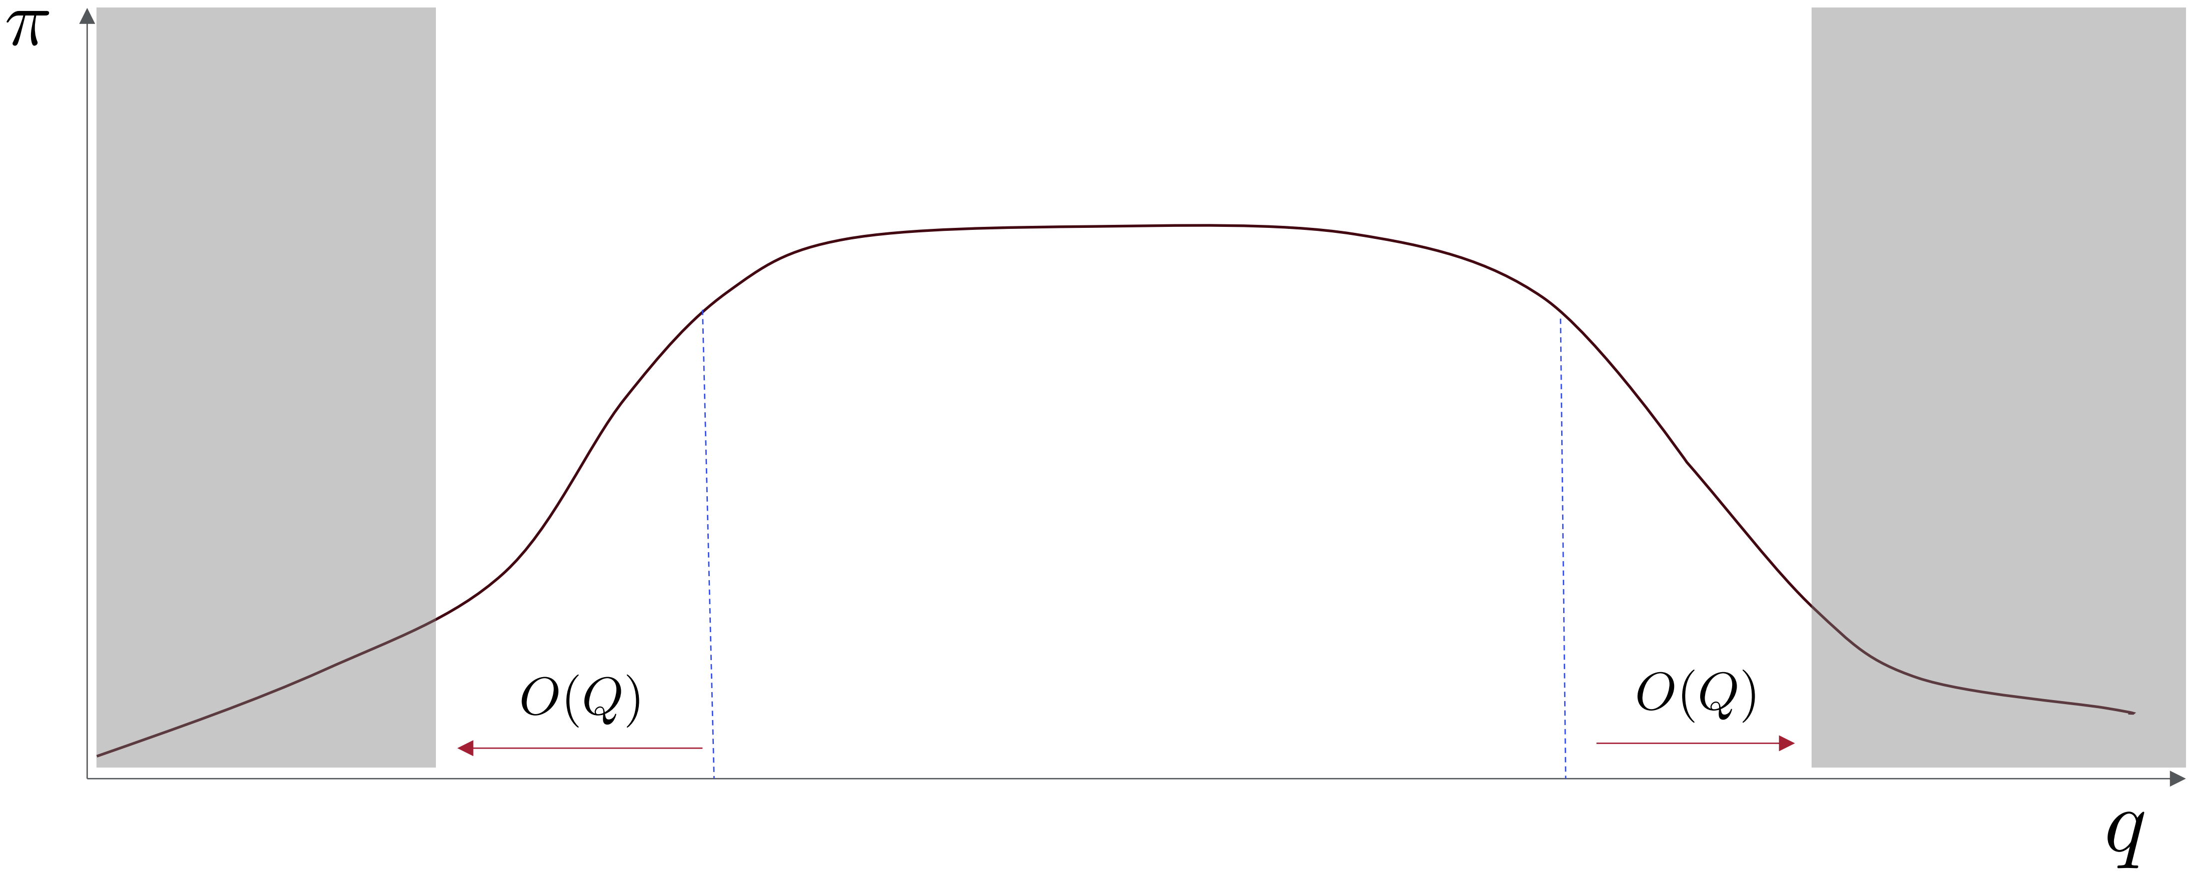

Step 2 – Distribution design problem

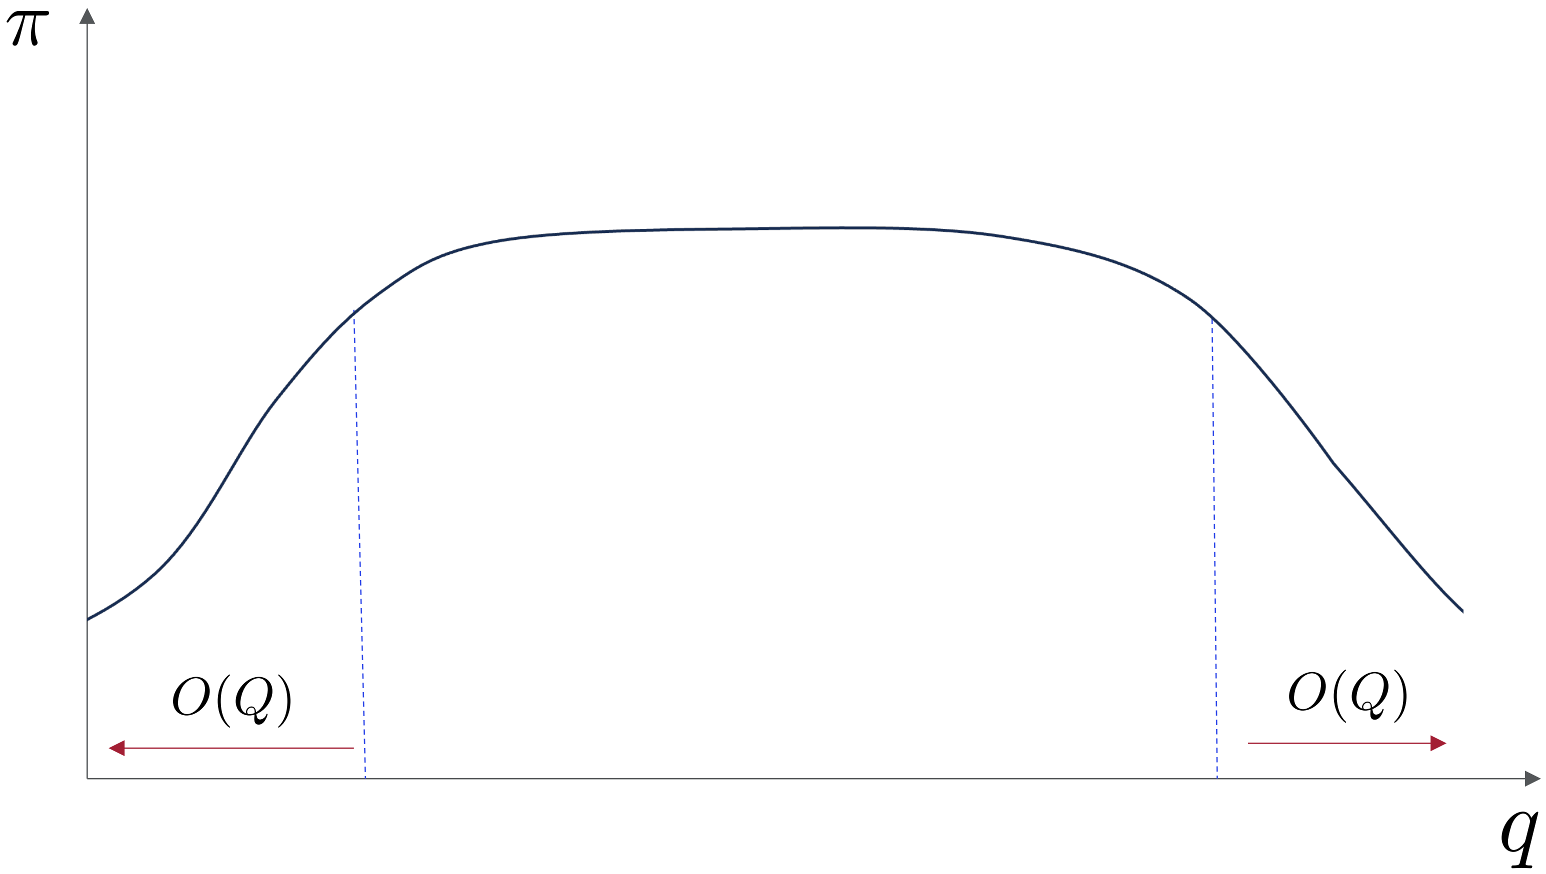

Step 2 – Distribution design problem

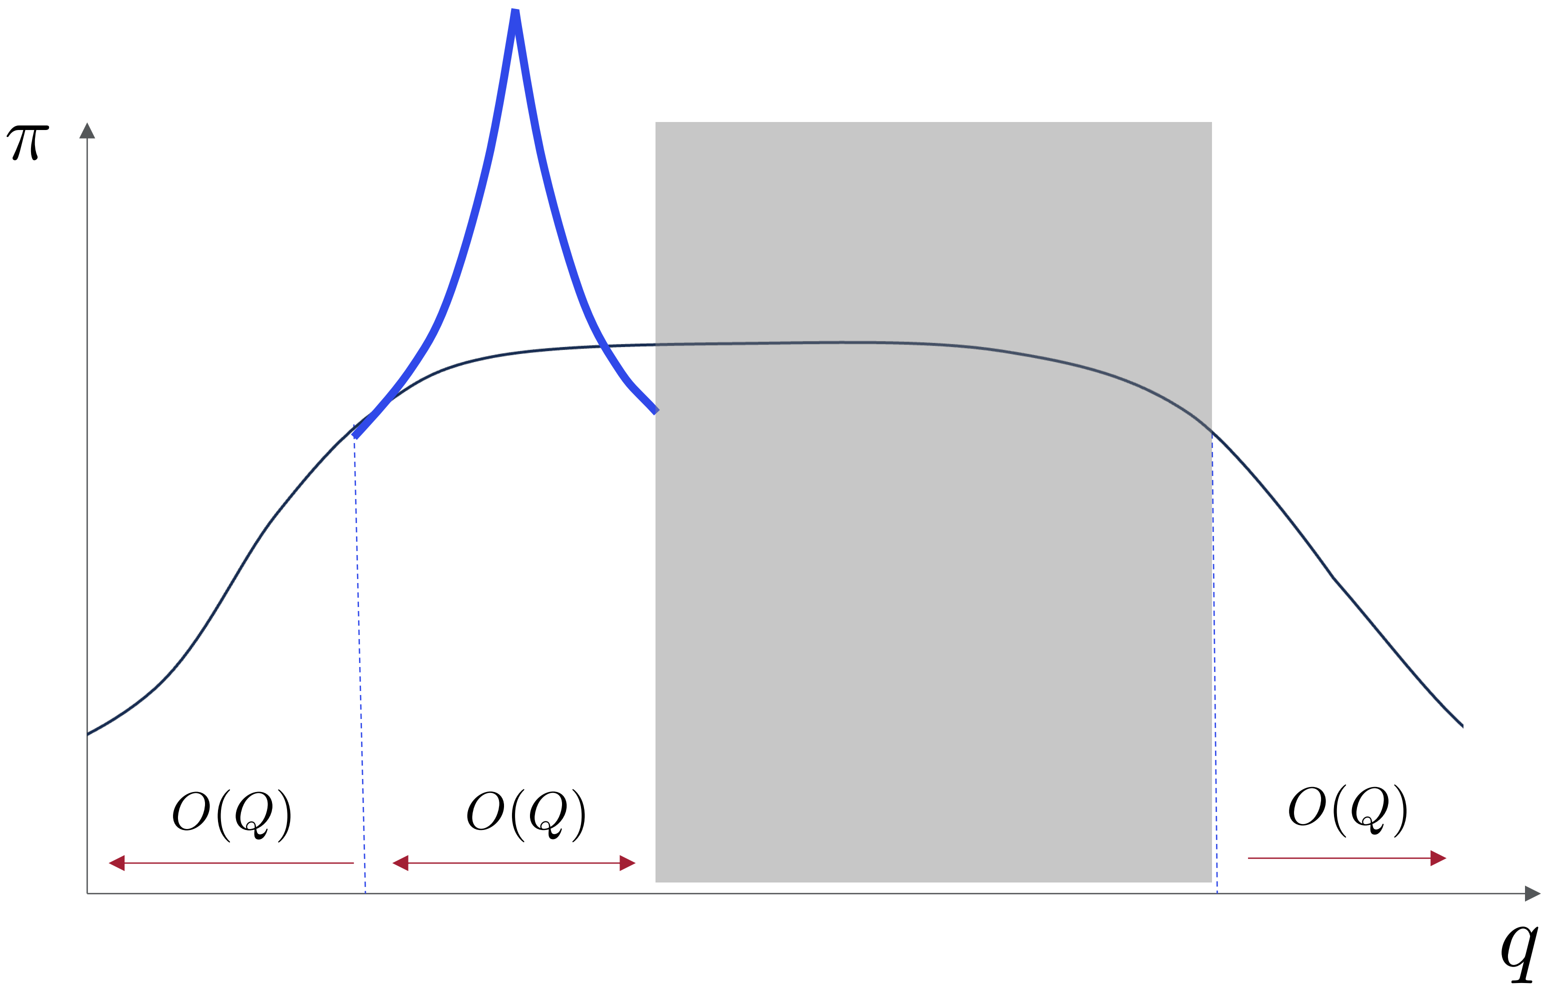

Step 2 – Distribution design problem

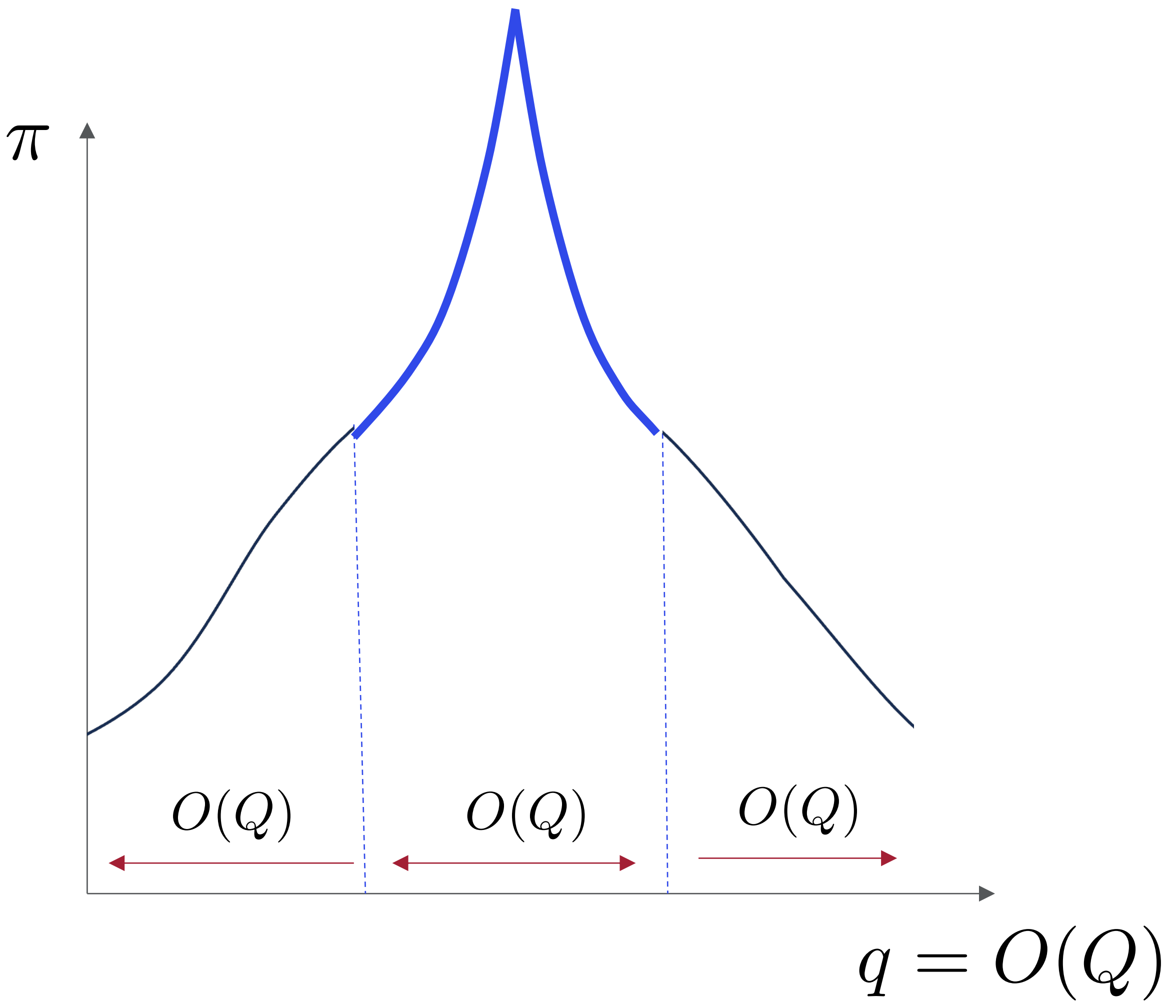

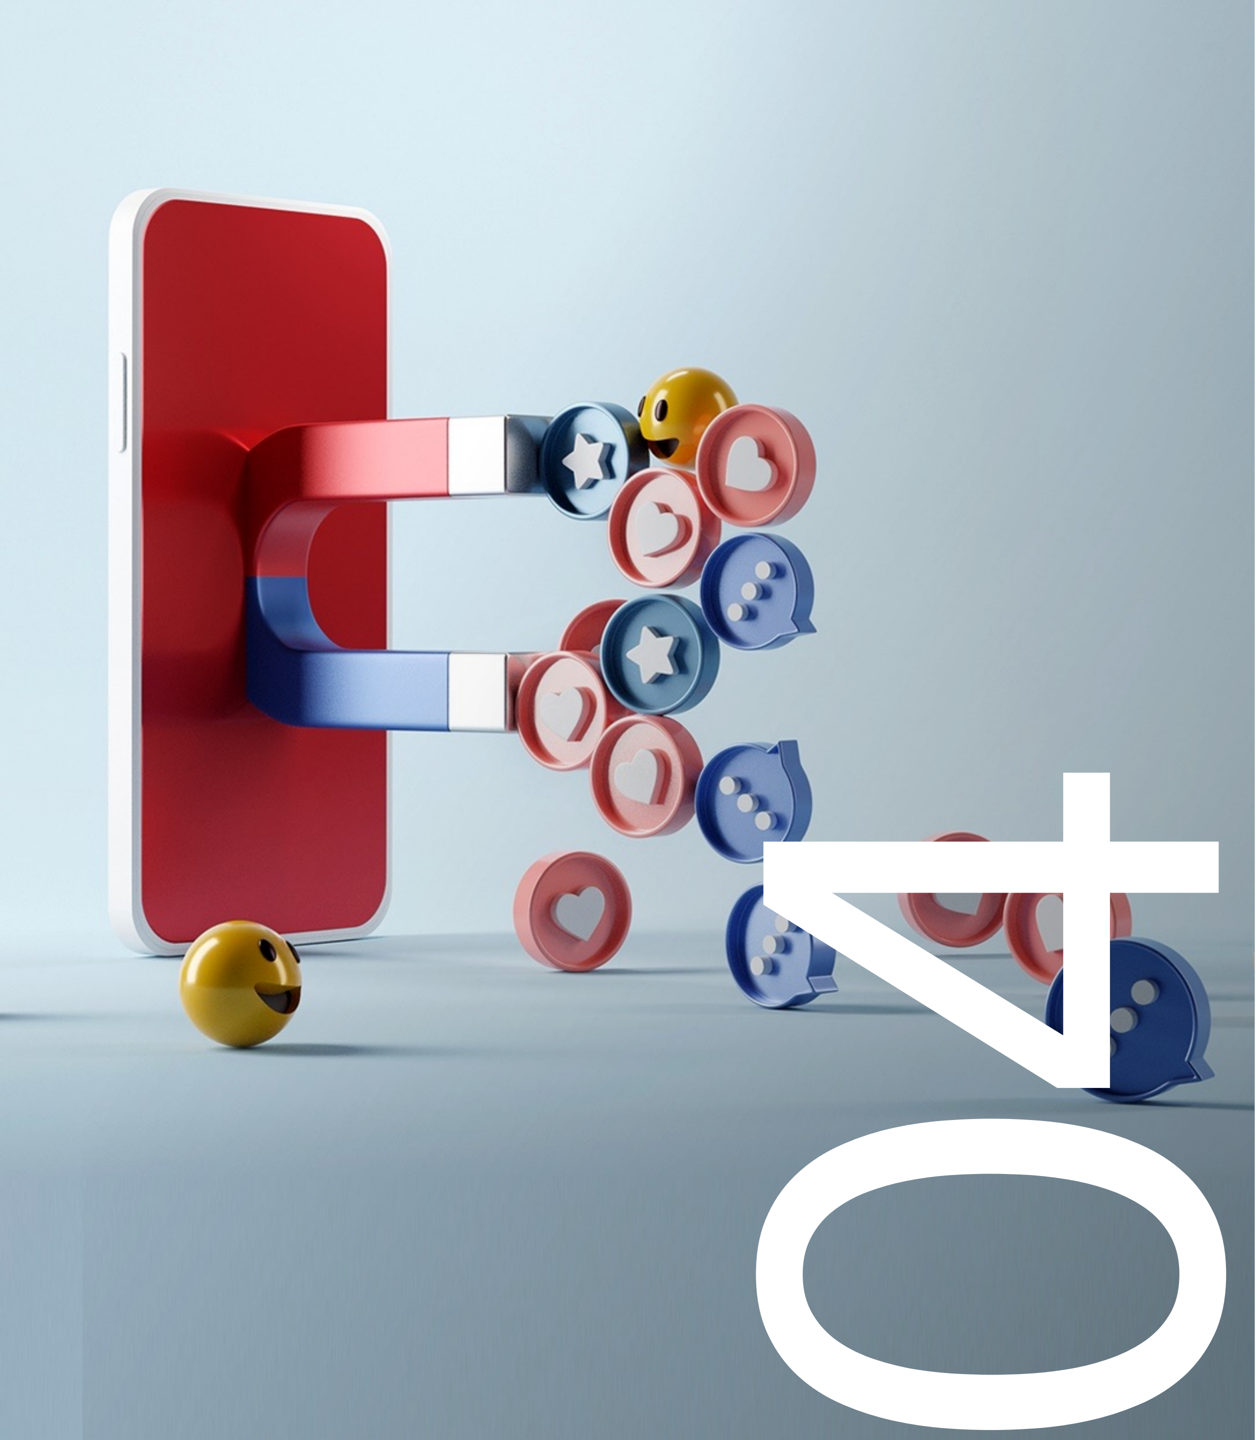

04

CRS REDUCTION– REWARDS SETTING

Main algorithmic results

Corollary 1 [AAS'24]: There is an FPTAS for the single-server setting.

Theorem 2 [AAS'24]: The cost-throughput problem on *d*-Euclidan graphs with uniform reneging rates can be approximated within factor $1-\epsilon$ in time $\mathrm{poly}(\epsilon^{-d} Q, |\mathcal{I}|)$.

Theorem 3 [AAS'24]: There exists an online rounding of DLP that is (1-1/e)-approximate (lossless reduction to offline contention resolution).

Correlated LP-rounding approach [AS, '22]

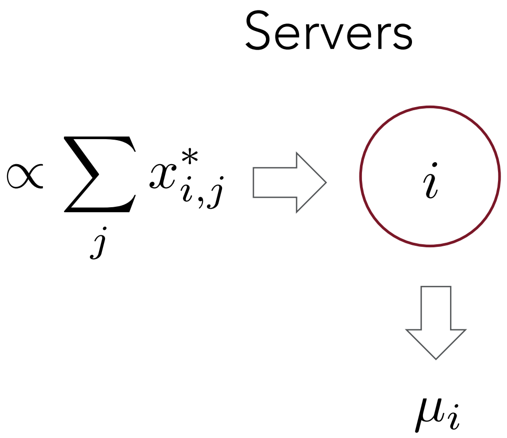

Correlated LP-rounding approach [AS, '22]

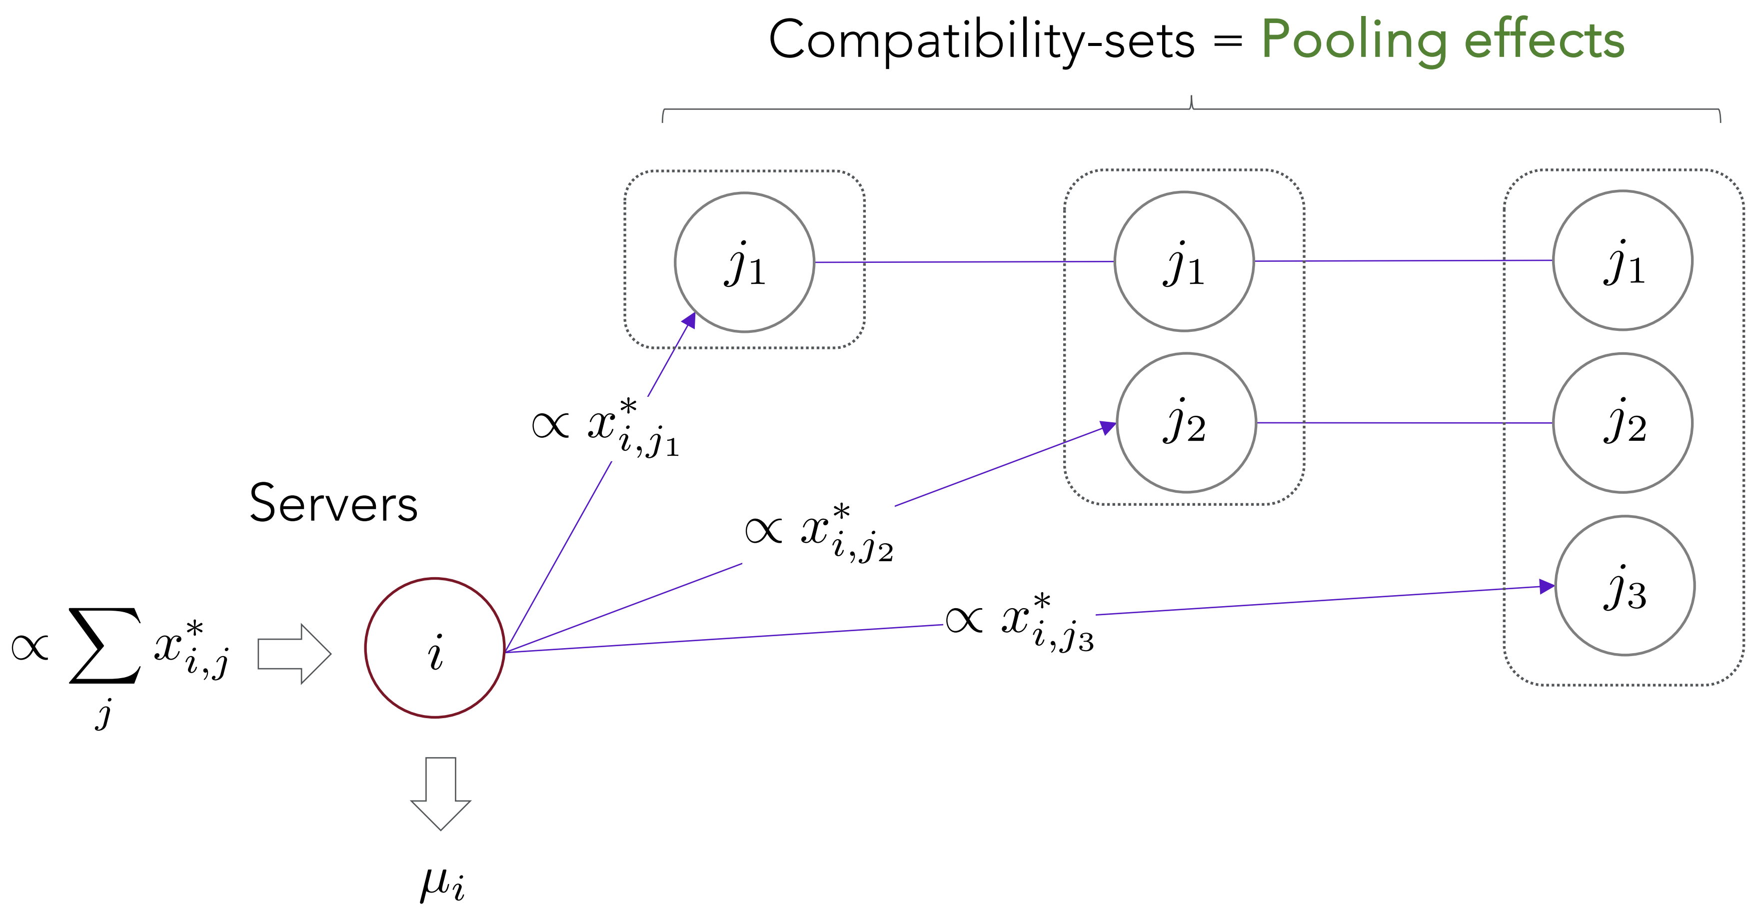

Correlated LP-rounding approach [AS, '22]

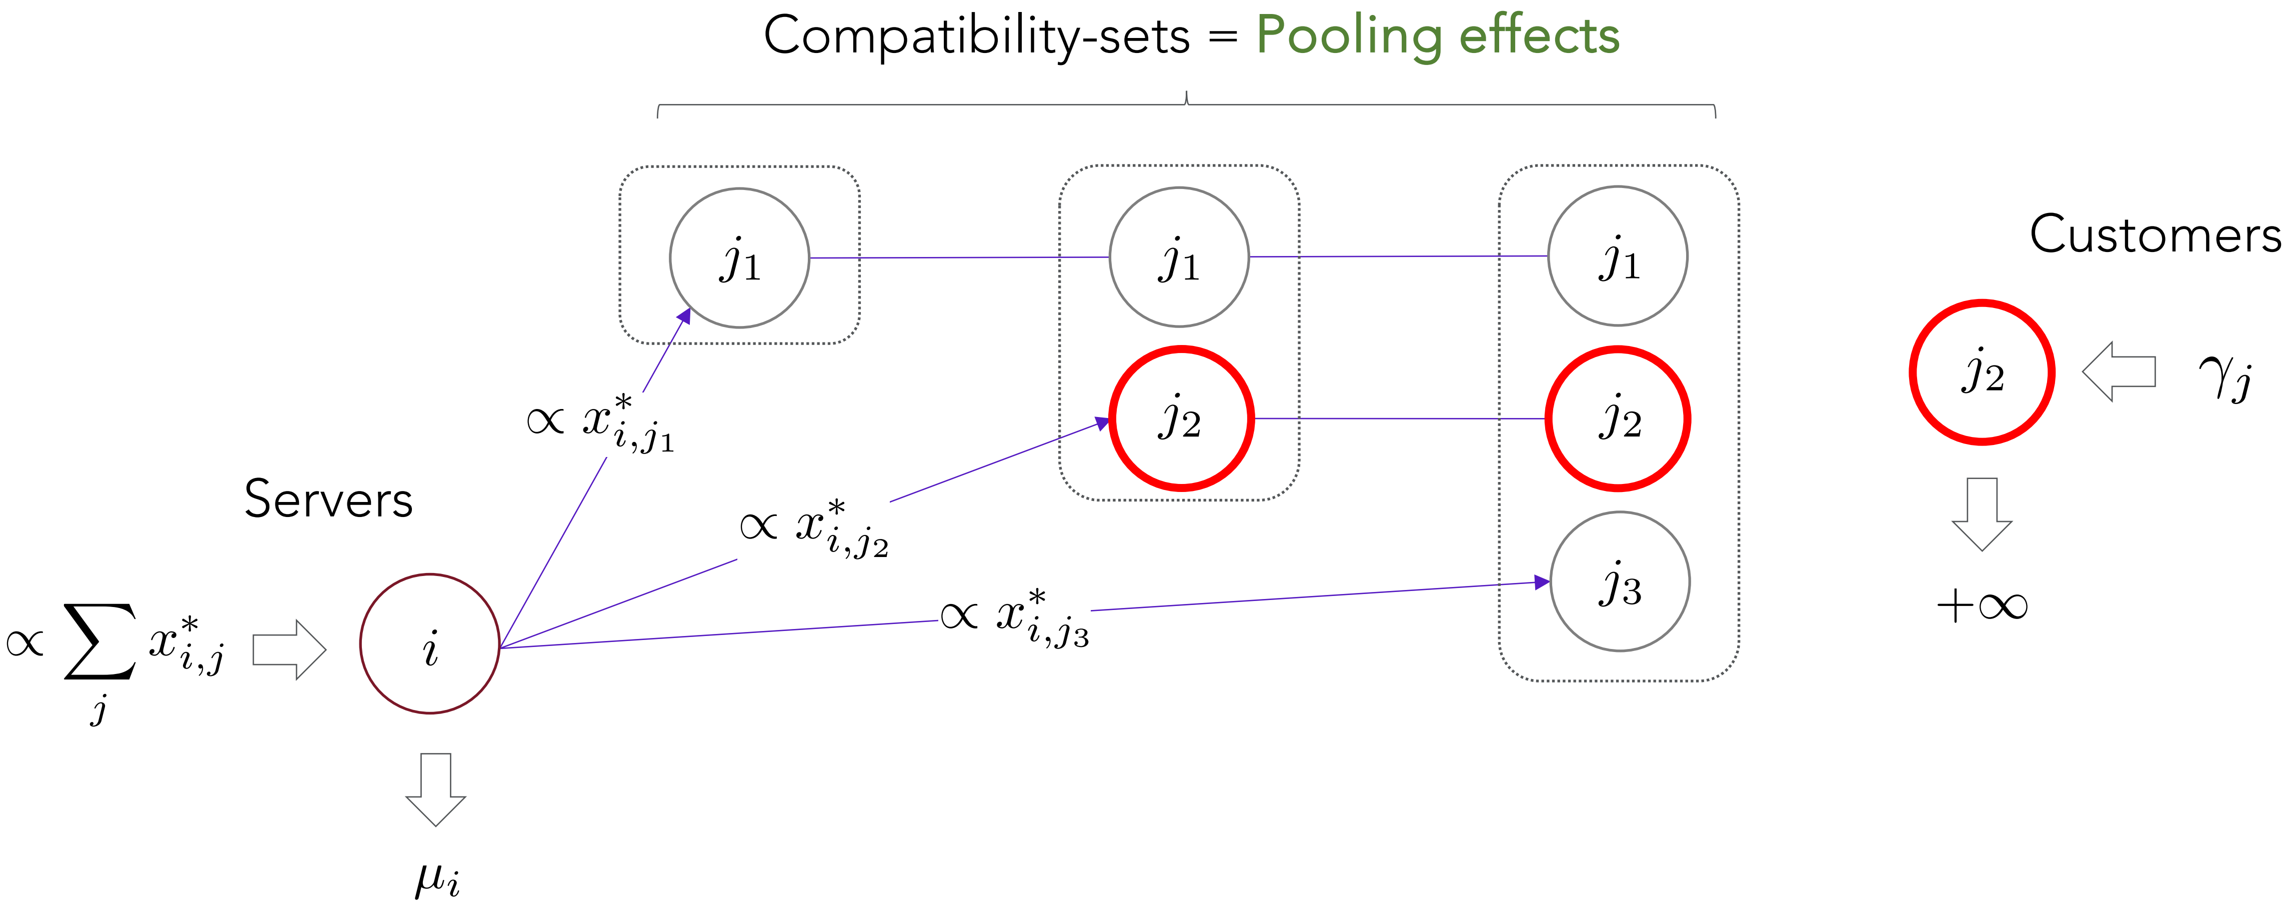

Reduction to contention resolution scheme

Vondrák et al. ['11]: For any matroid and any feasible $x \in \mathcal{P}_x$, there exists an efficient (1-1/e) – balanced contention resolution scheme.

(choosing item i with proba 1-1/e conditional on being independently sampled with proba x_i*)*

Reduction to contention resolution scheme

Vondrák et al. ['11]: For any matroid and any feasible $x \in \mathcal{P}_x$, there exists an efficient (1-1/e) – balanced contention resolution scheme.

(choosing item i with proba 1-1/e conditional on being independently sampled with proba x_i)

Algorithmic recipe

1. Approximately solve DLP
2. Upon an arrival of type *j*, independently draw server requests with probability $(x_{i,S}^{Q_i(t)*})_{i,S}$
3. Run CRS(j) on requests to match
4. Discard the unused requests.

Reduction to contention resolution scheme

Vondrák et al. ['11]: For any matroid and any feasible $x \in \mathcal{P}_x$, there exists an efficient (1-1/e) – balanced contention resolution scheme.

(choosing item i with proba 1-1/e conditional on being independently sampled with proba x_i*)*

Algorithmic recipe

1. Approximately solve DLP
2. Upon an arrival of type *j*, independently draw server requests with probability $(x_{i,S}^{Q_i(t)*})_{i,S}$
3. Run CRS(j) on requests to match
4. **Discard the unused requests.**

Fact: Discarding induces correlations ☹

Continuous-time discarding

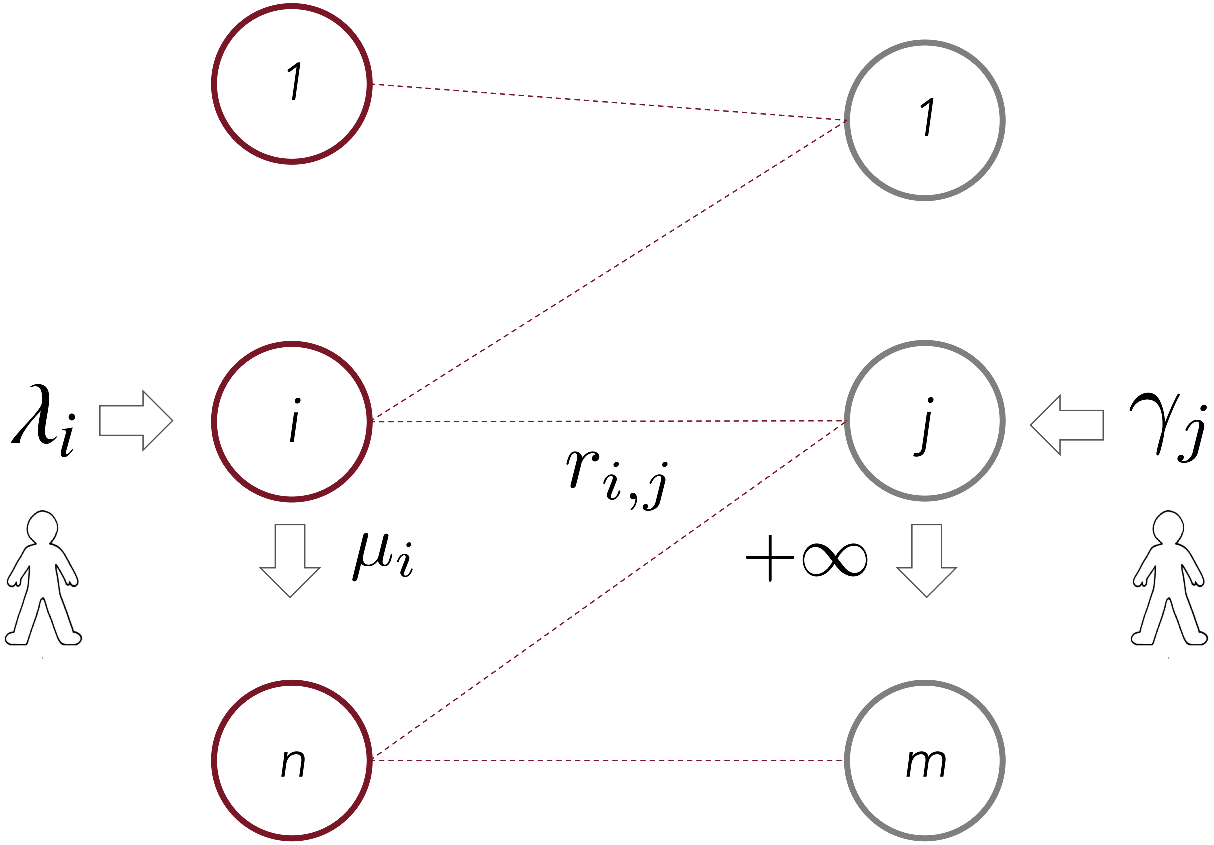

Continuous-time discarding

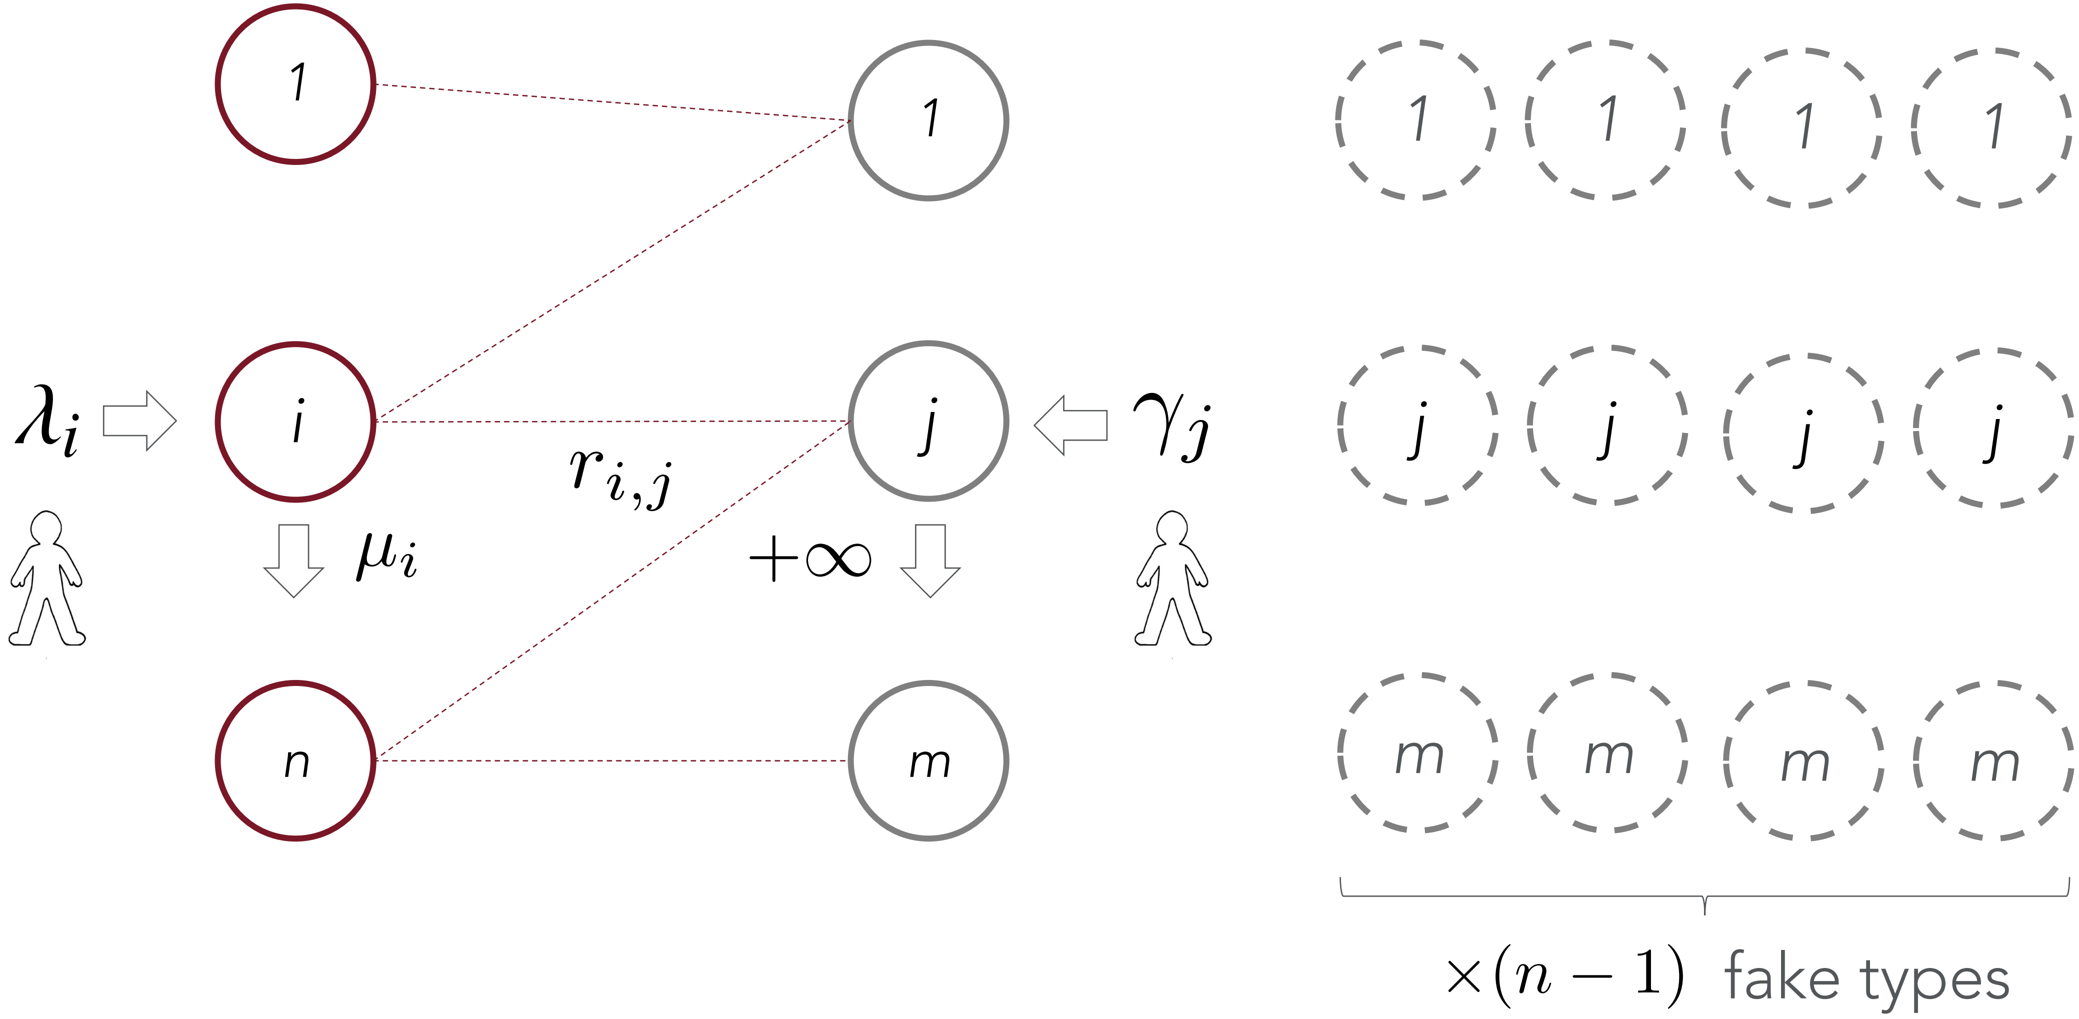

Continuous-time discarding