Integer Programming Games

Do You Really Need Them?

Gabriele Dragotto
Games, Network and Learning

IMS & NUS Singapore
April 3-6, 2023

DataX
Accelerating Scientific Discovery at Princeton

NUS
National University
of Singapore

PRINCETON
UNIVERSITY

Learn-and-Play

From Data to Complex Interventions

Gabriele Dragotto
Games, Network and Learning

IMS & NUS Singapore
April 3-6, 2023

Accelerating Scientific Discovery at Princeton

NUS National University of Singapore

PRINCETON UNIVERSITY

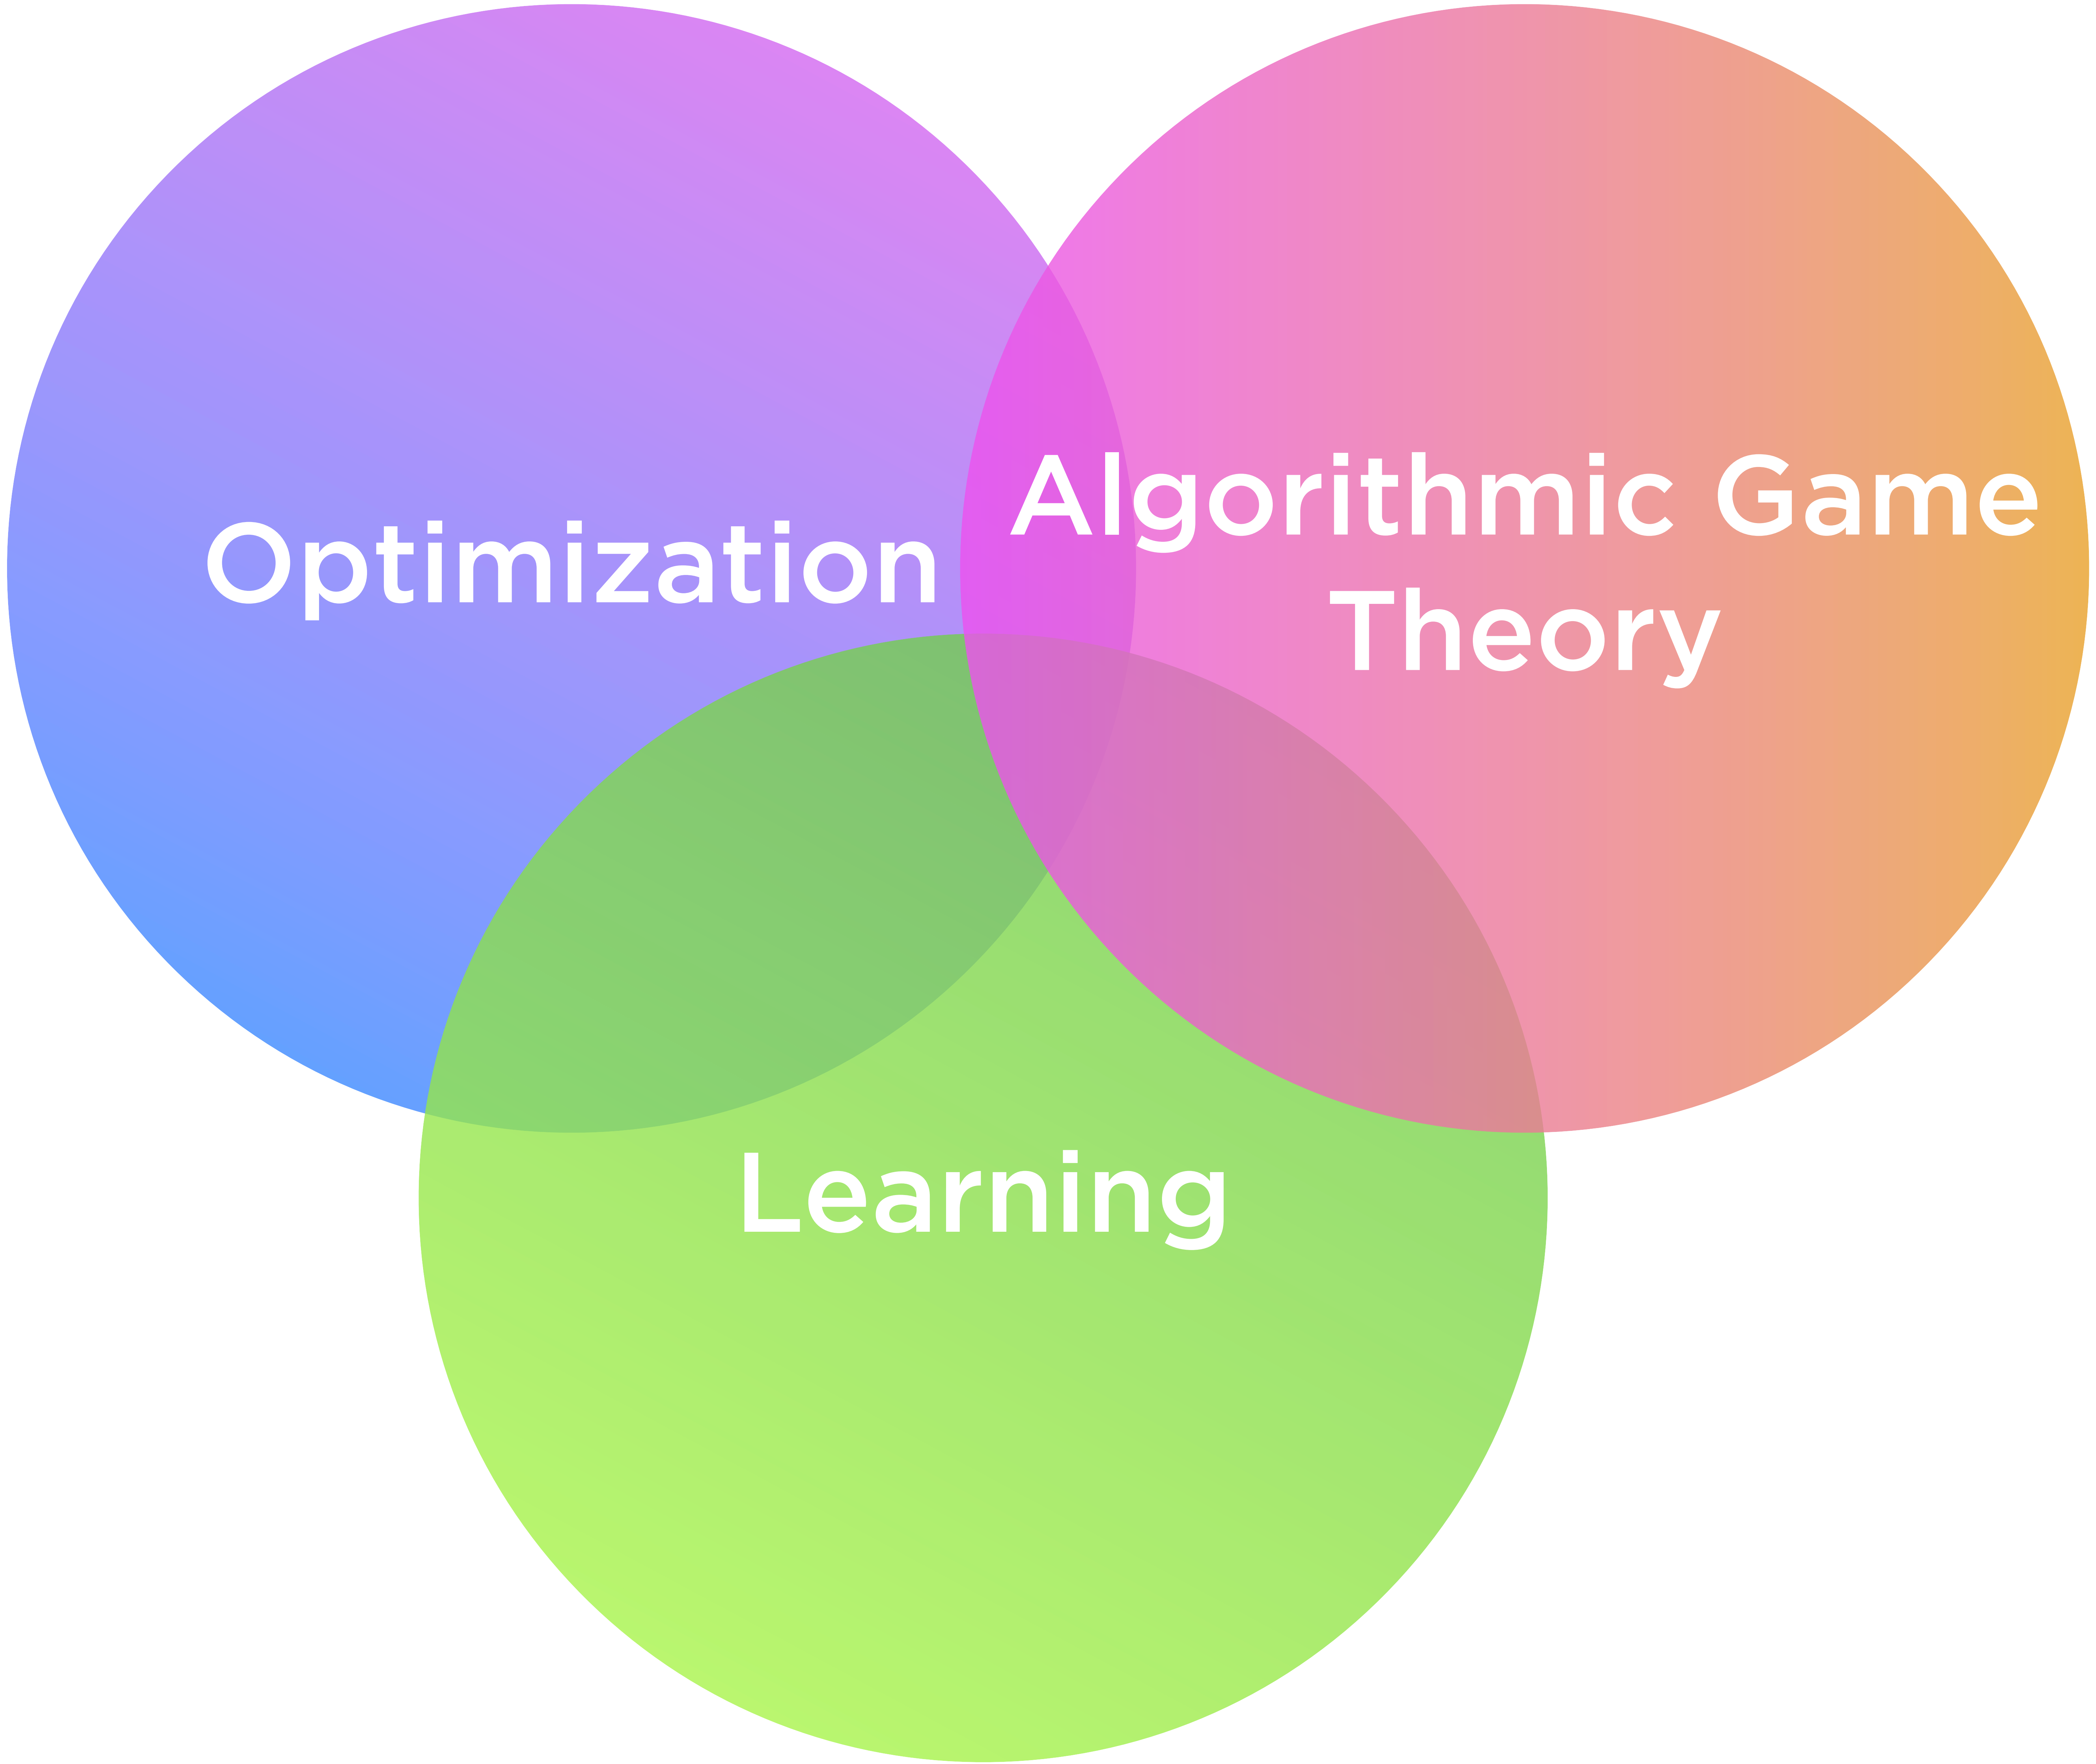
Optimization
Algorithmic Game Theory
Learning

Network Congestion

Network Congestion

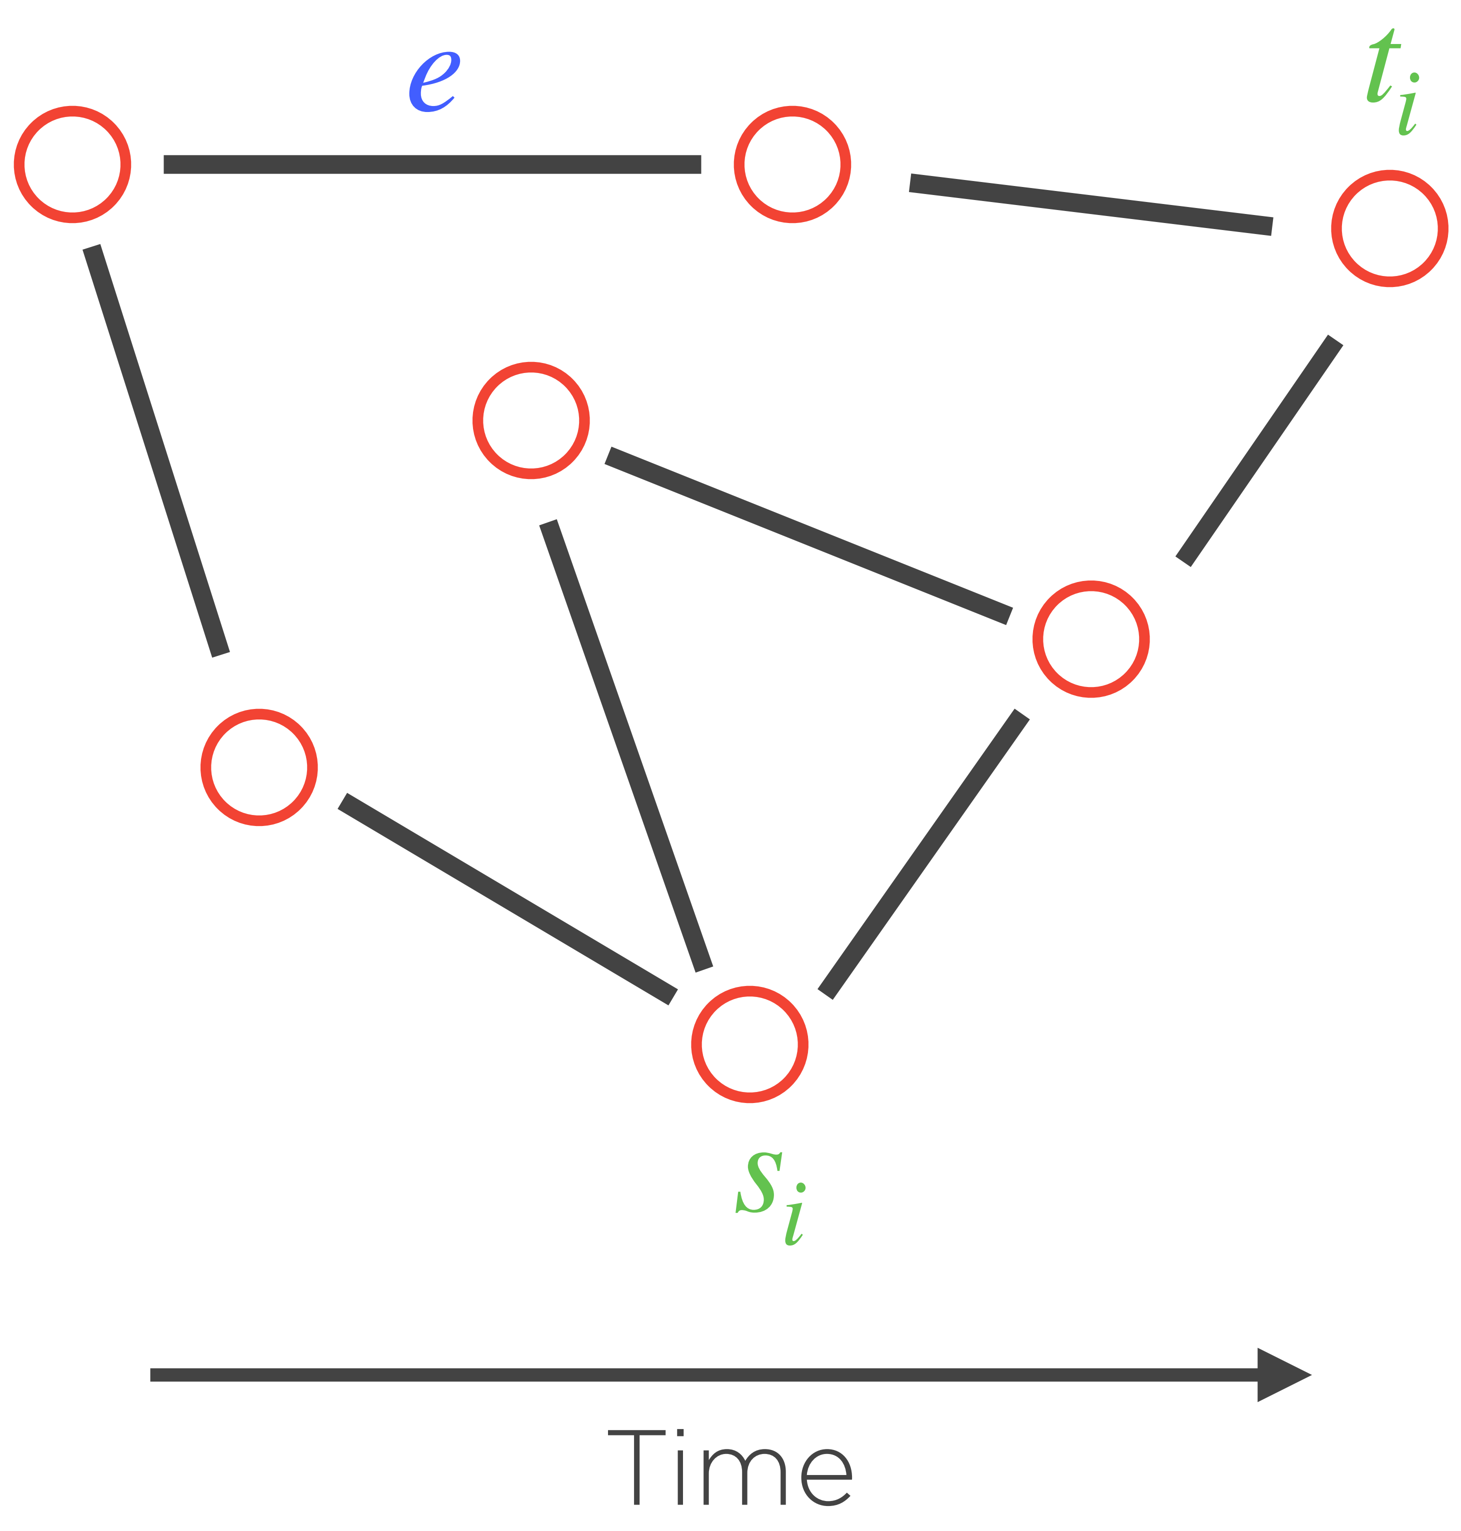

There are n players optimizing simultaneously the shortest path on a graph $G = (V, E)$ so that:

- The player i needs to go **from s_i to t_i**
- $x_{ie} = 1$ if the player i selects the edge $e \in E$

$$\min_{x_i} \{u_i(x_i; x_{-i}) : x_i \in \mathcal{X}_i\}$$

A regulator **observes the outcome** of the interaction but **does not know** some utilities/actions

A regulator wants to **intervene in the game**

Network Congestion

A regulator **observes the outcome** of the interaction but **does not know** some utilities/actions

A regulator wants to **intervene in the game**

Decision-making is rarely an individual task

Self-driven interactions with other decision-makers

Decision-making is rarely an individual task

Self-driven interactions with other decision-makers

We consider the perspective of an external regulator

Self-driven behavior often **conflicts with societal goals**

External regulators should **learn the agents' preferences** and **intervene**

Decision-making is rarely an individual task

External regulators should **learn the agents' preferences** and **intervene**

Learn the Parameters

Some information regarding the players' optimization problems **is missing**

Intervene in "complex" settings

Select Nash equilibria when players solve **mixed-integer optimization problems**

Learning Rationality in Potential Games

Stefan Clarke, Bartolomeo Stellato, and Jaime Fernandez Fisac
(Princeton University, USA)



Learn the Parameters

Problem setup

Simultaneous and non-cooperative game where $i = 1, \ldots, n$ solves

$$\begin{aligned} \min_{x^i} \quad & u_i(x_i; x_{-i}, \theta, \mu) \\ \text{s.t.} \quad & x_i \in \mathcal{X}_i = \{B_i(\theta, \mu) x_i + D_i(\theta, \mu) x_{-i} \leq b_i(\theta, \mu)\} \end{aligned}$$

A set of unknown **rationality parameters**

Known and observable **context parameters**

The utility u_i is convex-quadratic in x, and

There exists a convex-quadratic potential function $\Phi(x; \theta, \mu)$

Our approach

Simultaneous and non-cooperative game where $i = 1, \ldots, n$ solves

$$\begin{aligned} \min_{x^i} \quad & u_i(x_i; x_{-i}, \theta, \mu) \\ \text{s.t.} \quad & x_i \in \mathcal{X}_i = \{B_i(\theta, \mu) x_i + D_i(\theta, \mu) x_{-i} \leq b_i(\theta, \mu)\} \end{aligned}$$

We observe data $\mathcal{D} = \{(\bar{x}^k, \bar{\mu}^k)\}_{k=1}^K$ with equilibria and context

Inverse equilibrium task

Estimate θ so that it predicts the Nash equilibria $\bar{x}^k$

Data → Our Algorithm → Parameters

The three ingredients

1 **Potentiality**

Nash equilibria: $\min_{x}\{\Phi(x;\theta,\mu) : x_i \in \mathcal{X}_i, \quad i = 1,\dots,n\}$

2 **Learning Problem**

$\min_{x^k,\lambda^k,\theta}$ $\mathcal{L}(\theta;\mathcal{D})$ L2 norm between target and prediction

subject to Target is a Nash equilibrium,

θ belongs to a set of feasible parameters

$x^k \in \mathbb{R}^{mn}, \lambda^k \in \mathbb{R}^{ln}_+ \quad k = 1,\dots,K,$

$\theta \in \Theta.$

The three ingredients

1 **Potentiality**

Nash equilibria: $\min_x \{\Phi(x;\theta,\mu) : x_i \in \mathcal{X}_i, \quad i = 1,\dots,n\}$

2 **Learning Problem**

$$
\begin{aligned}
\min_{x^k,\lambda^k,\theta} \quad & (1/K)\textstyle\sum_{k=1}^{K} \|x^k - \bar{x}^k\|_2^2 \\
\text{subject to} \quad & 0 = R(\theta,\bar{\mu}^k)x^k + c(\theta,\bar{\mu}^k) + A(\theta,\bar{\mu}^k)^T\lambda^k, \\
& 0 \le b(\theta,\bar{\mu}^k) - A(\theta,\bar{\mu}^k)x^k \perp \lambda^k \ge 0 \\
& x^k \in \mathbb{R}^{mn}, \lambda^k \in \mathbb{R}_+^{ln} \quad k = 1,\dots,K, \\
& \theta \in \Theta.
\end{aligned}
$$

The three ingredients

1 **Potentiality**

Nash equilibria: $\min_{x} \{\Phi(x;\theta,\mu) : x_i \in \mathcal{X}_i, \quad i = 1,\dots,n\}$

2 **Learning Problem**

$$
\begin{aligned}
\min_{x^k,\lambda^k,\theta} \quad & (1/K)\textstyle\sum_{k=1}^{K} \|x^k - \bar{x}^k\|_2^2 \\
\text{subject to} \quad & 0 = R(\theta,\bar{\mu}^k)x^k + c(\theta,\bar{\mu}^k) + A(\theta,\bar{\mu}^k)^T\lambda^k, \\
& 0 \le b(\theta,\bar{\mu}^k) - A(\theta,\bar{\mu}^k)x^k \perp \lambda^k \ge 0 \\
& x^k \in \mathbb{R}^{mn}, \lambda^k \in \mathbb{R}_+^{ln} \quad k = 1,\dots,K, \\
& \theta \in \Theta.
\end{aligned}
$$

We would like to find a (local) minimum of the learning problem with a **first-order method**

The three ingredients

3 **Differentiation**

Since the previous formulation is non-convex:

- We differentiate $\mathcal{L}(\theta; \mathcal{D})$ with respect to the **parameters θ**
- **How?** We fix the "tight" complementarity constraints to get a **convex inner approximation of the learning problem**

Active set, i.e., *the set of indices of **tight complementarity constraints***

$$Z \leftarrow \{z : b(\theta, \bar{\mu}^k)_z - A(\theta, \bar{\mu}^k)_z x^k = 0\}$$

We employ $\nabla_\theta \mathcal{L}(\theta; \mathcal{D})$ to update our estimates of θ

The Algorithm

INPUT Max iterations T, step sizes $\{\eta\}_{t=1}^{T}$, and data $\mathcal{D} = \{(\bar{x}^k, \bar{\mu}^k)\}_{k=1}^{K}$

1. **Initialization** — Initial parameters $\theta^{(0)}$

Loop T Times

2.1 **Select** — Sample a data point $(\bar{x}^k, \bar{\mu}^k)$

2.2 **Play** — $(x^t, \lambda^t) \leftarrow \min_x \{\Phi(x; \theta^{(t)}, \bar{\mu}^k) : x_i \in \mathcal{X}_i(\theta^{(t)}, \bar{\mu}^k) \ \ \forall i\}$

2.3 **Differentiate** — Compute $\nabla_\theta \mathcal{L}(\theta; \mathcal{D})$ on the current active set

$$\theta^{(t+1)} \leftarrow \theta^{(t)} - \eta_t \nabla_\theta \mathcal{L}_t(x^{(t)})$$

OUTPUT $\theta^{(T)}$

Convergence

Convergence

The sampled gradient of the loss with respect to $\theta^{(t)}$ *converges to* <u>*zero*</u>

With a **careful choice of the active set** at each iteration, the algorithm mimics a **stochastic gradient descent** with well-behaved derivates

$$\lim_{T\to\infty} \mathbb{E}[||\nabla g(\theta^{(T)})||_2] = 0$$

Where g is a **smoothened version of the loss**

Network Congestion

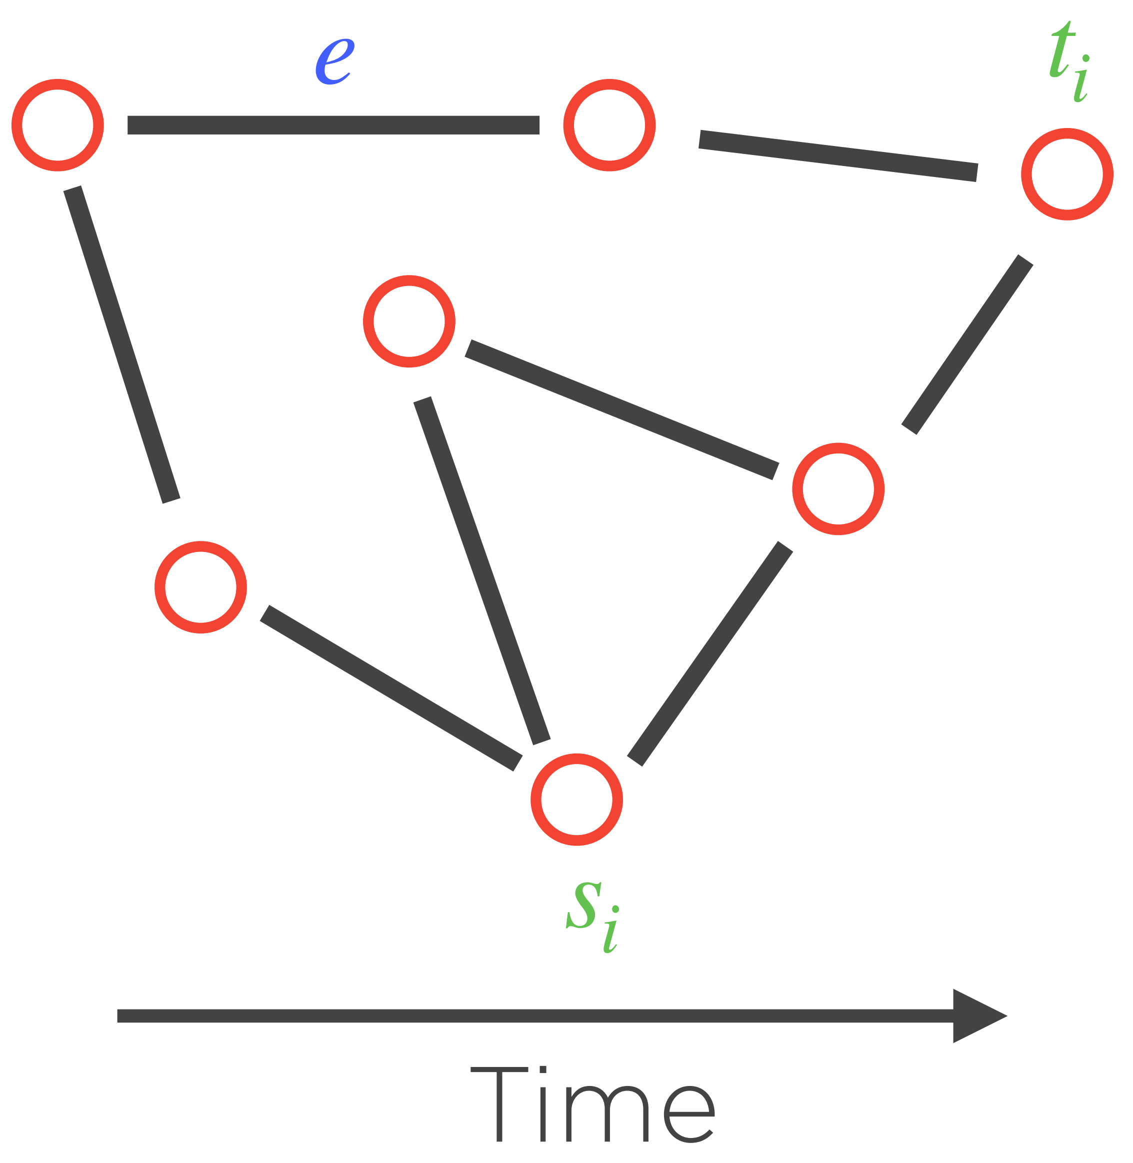

$$u_i(x_i; x_{-i}, \theta, \mu) = \sum_{e \in E} \theta_{ie}^{\top} l_e x_{ie}(x_{1e} + \cdots + x_{ne})$$

A set of unknown **rationality parameters**

Personal preferences

Known and observable **context parameters**

Traffic, weather, road conditions

Network Congestion

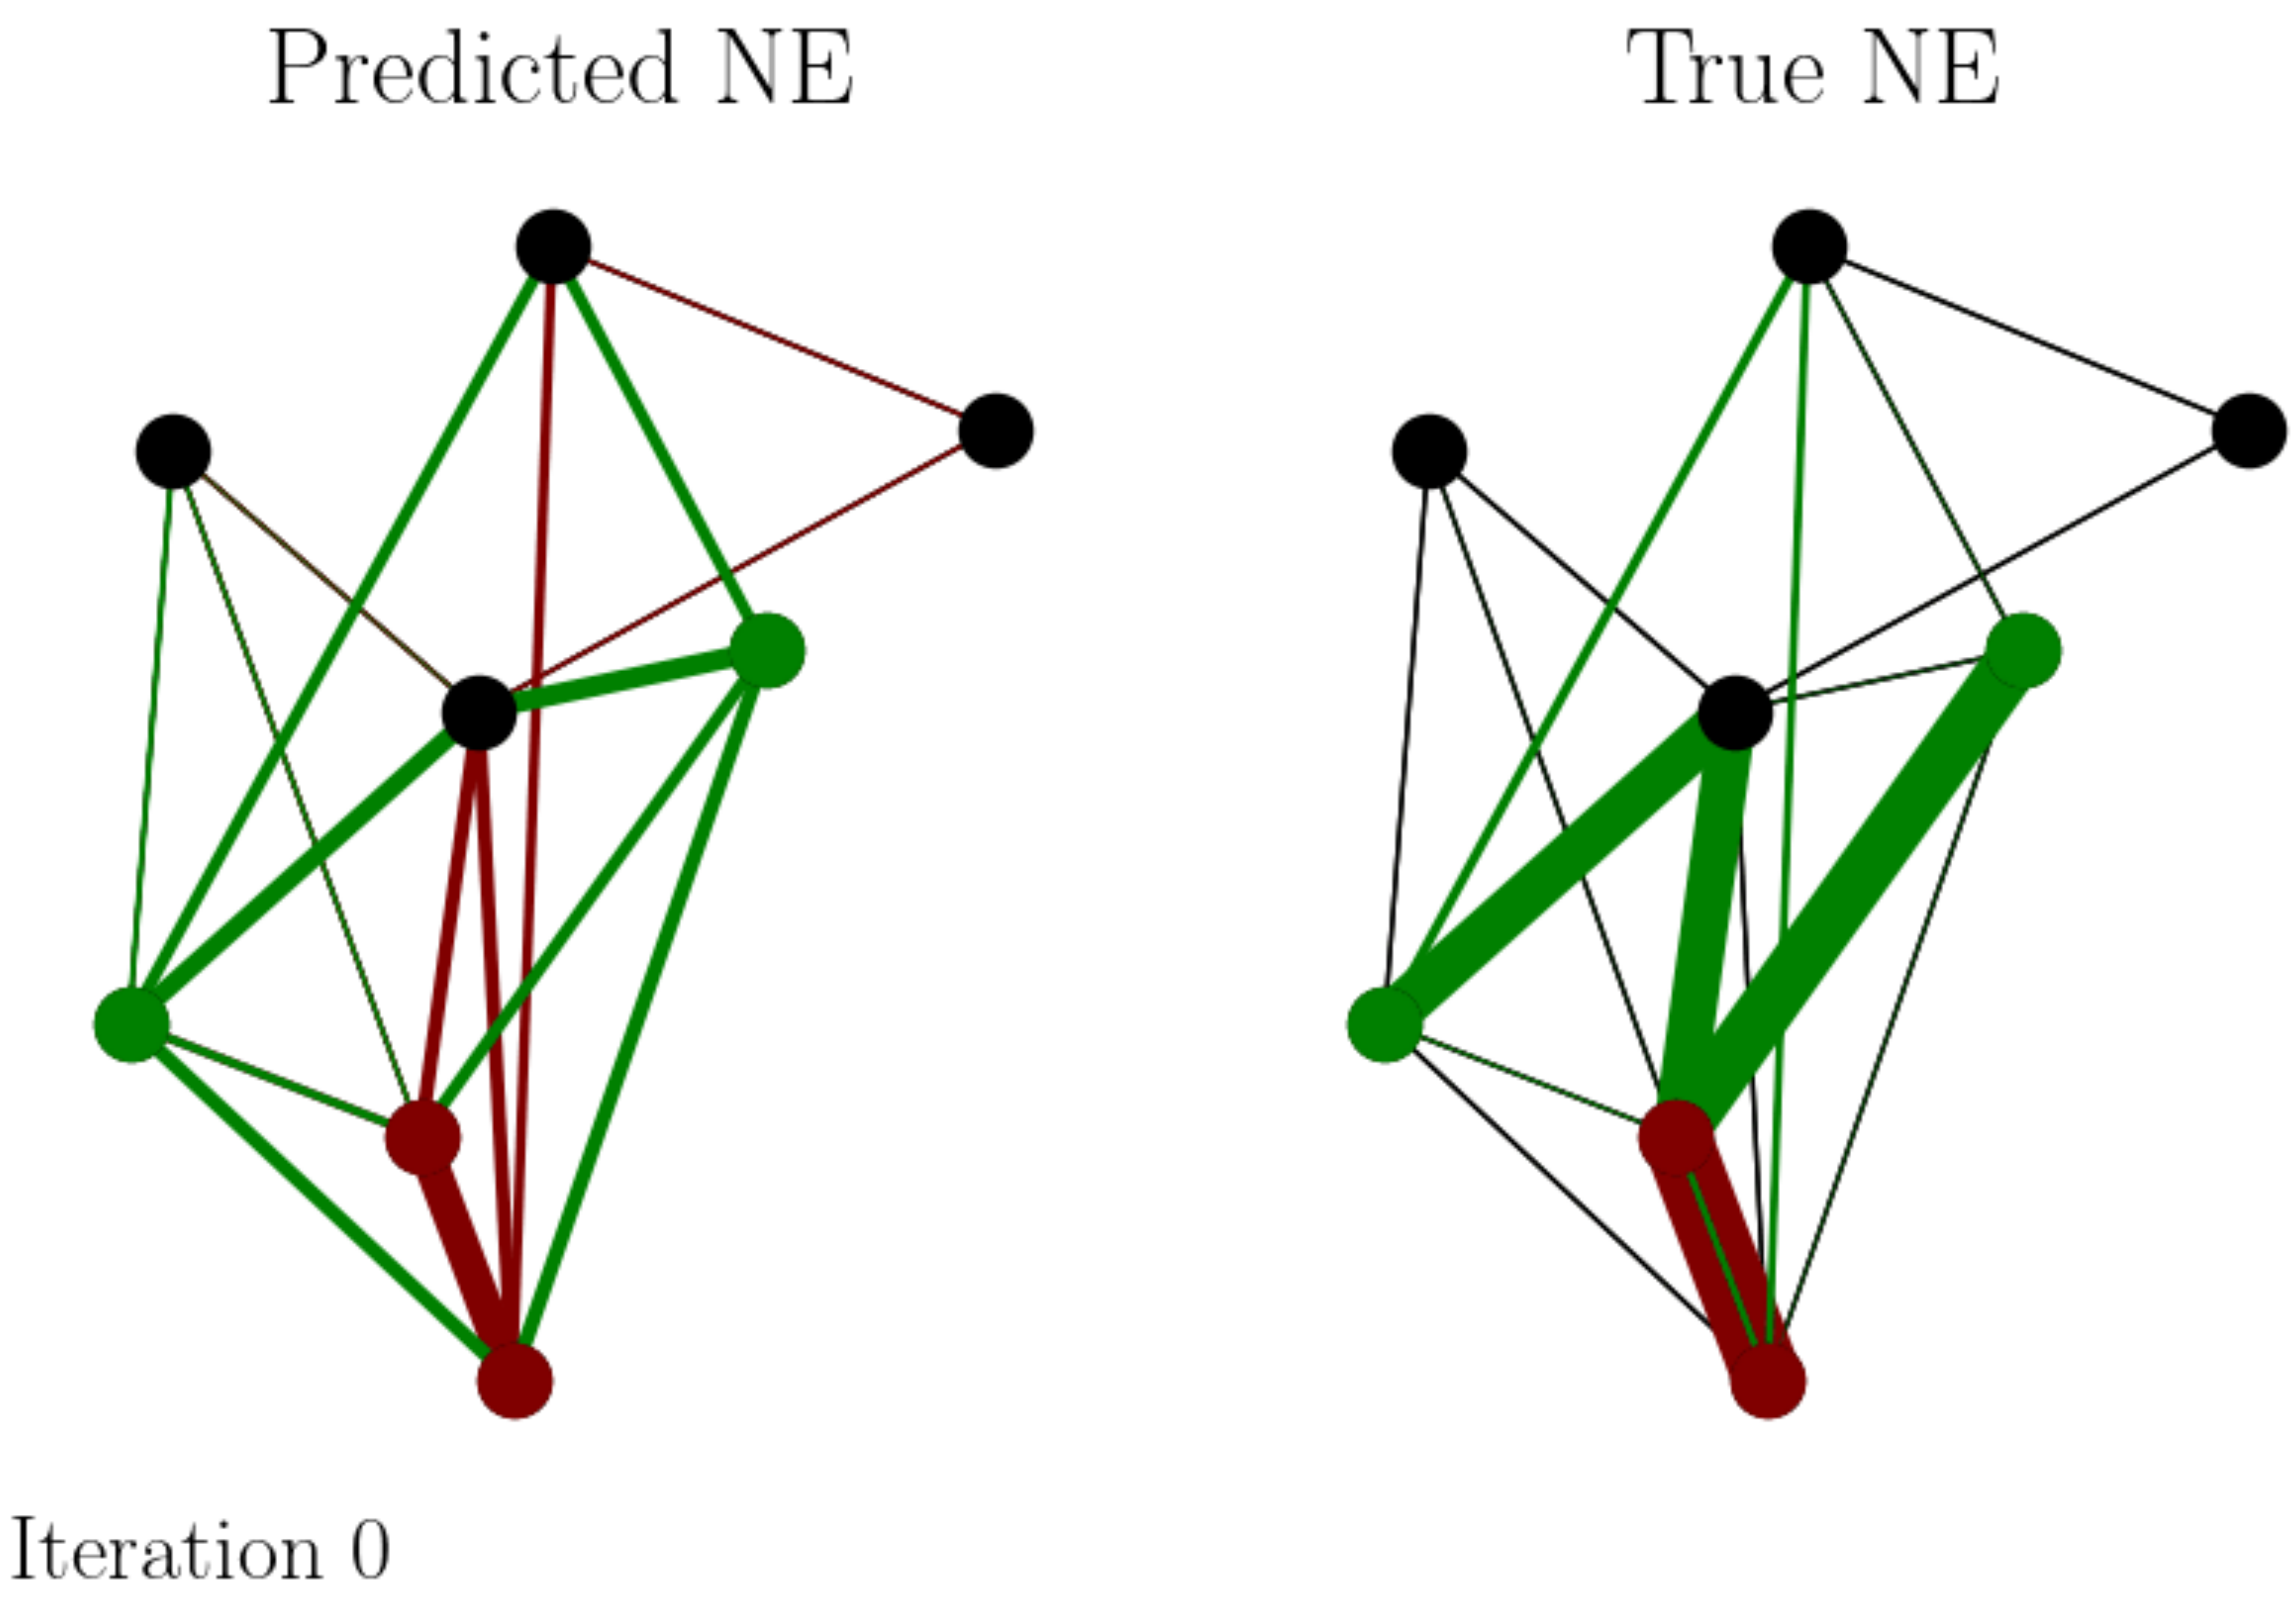

Network Congestion

Test error

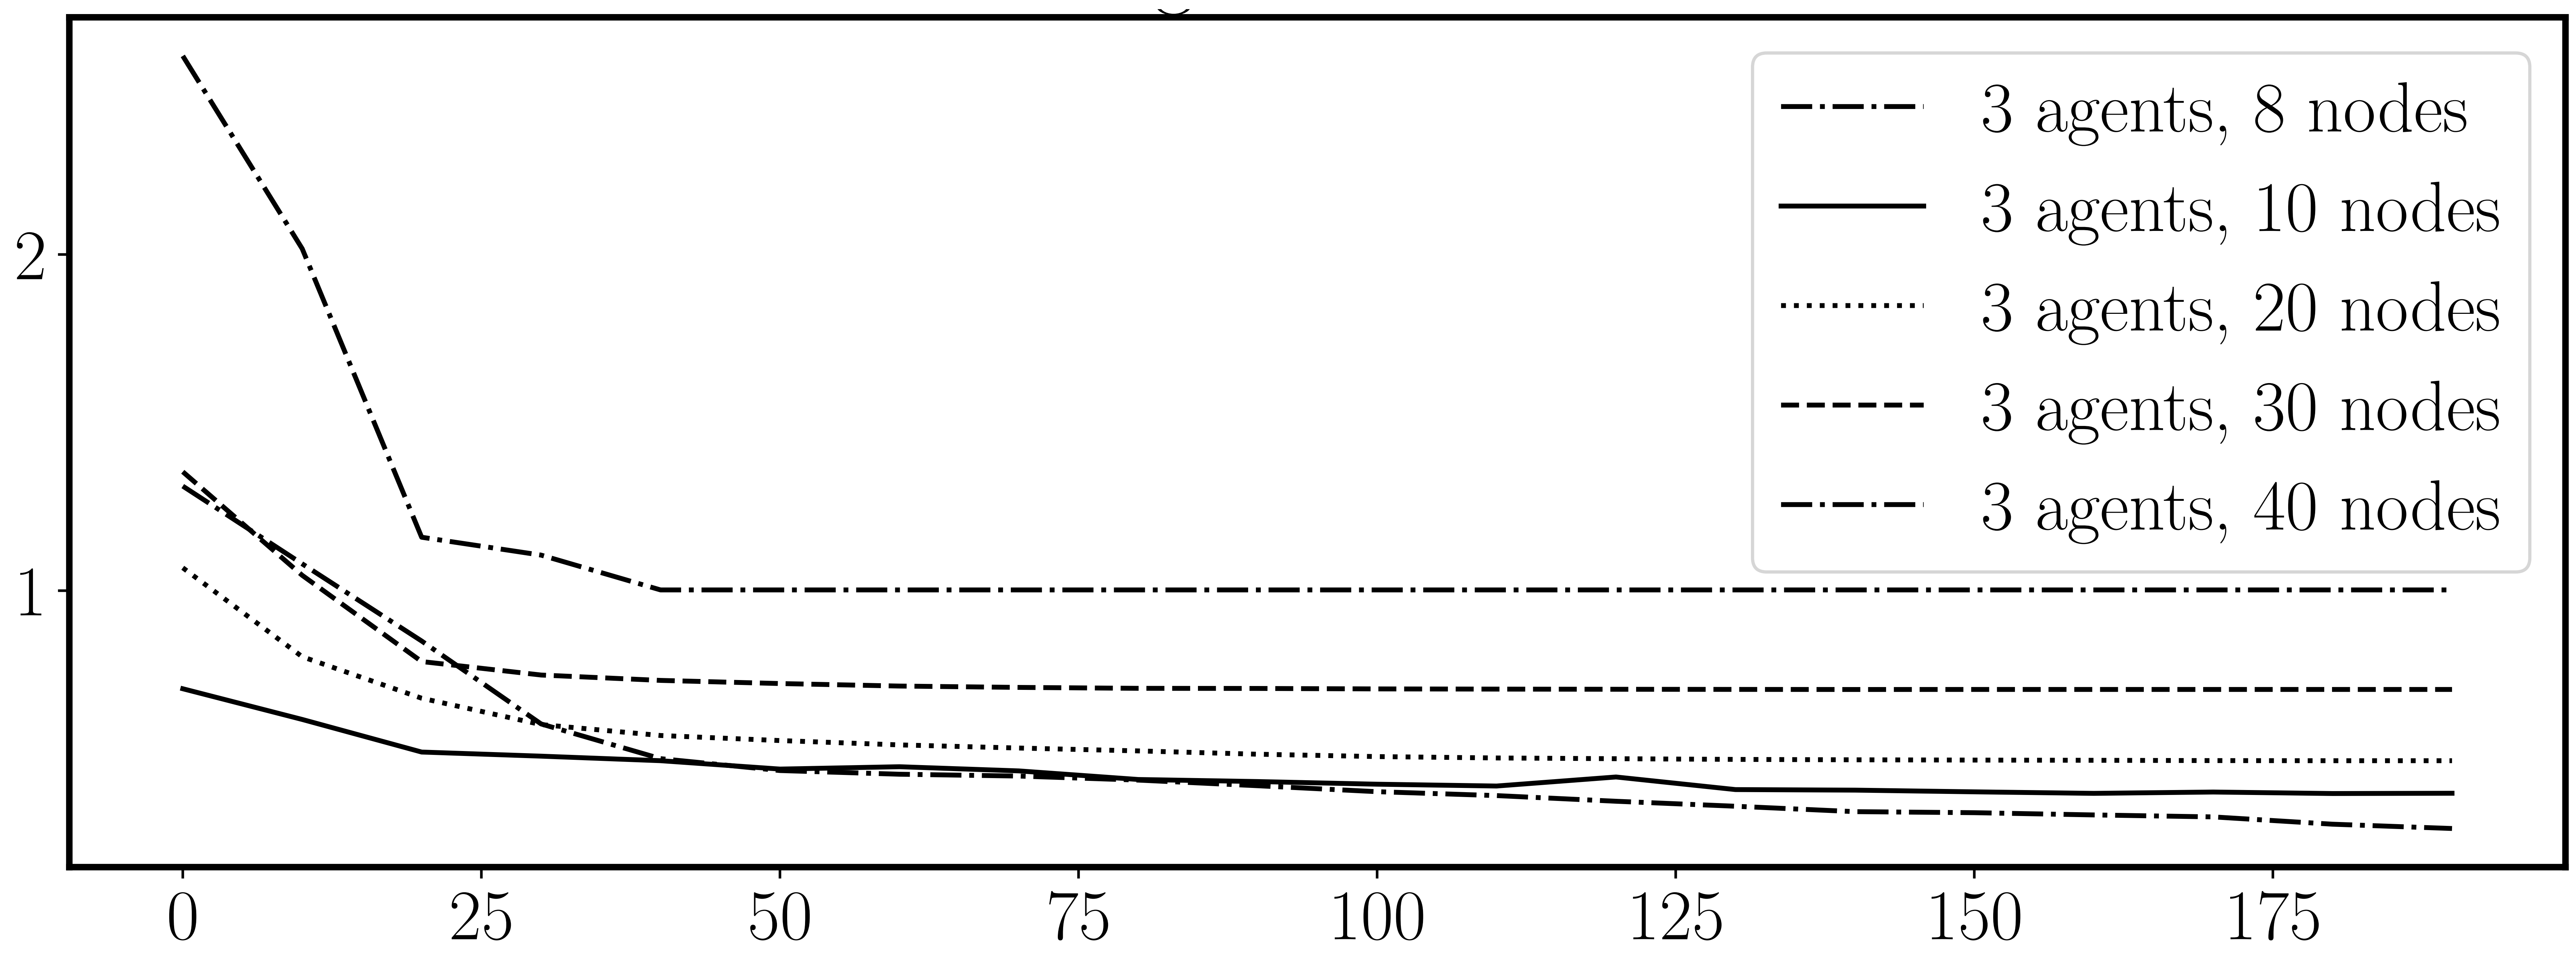

Iterations

Dataset of *90 equilibria*

We learn **good estimates** of the rationality parameters

Cournot Games

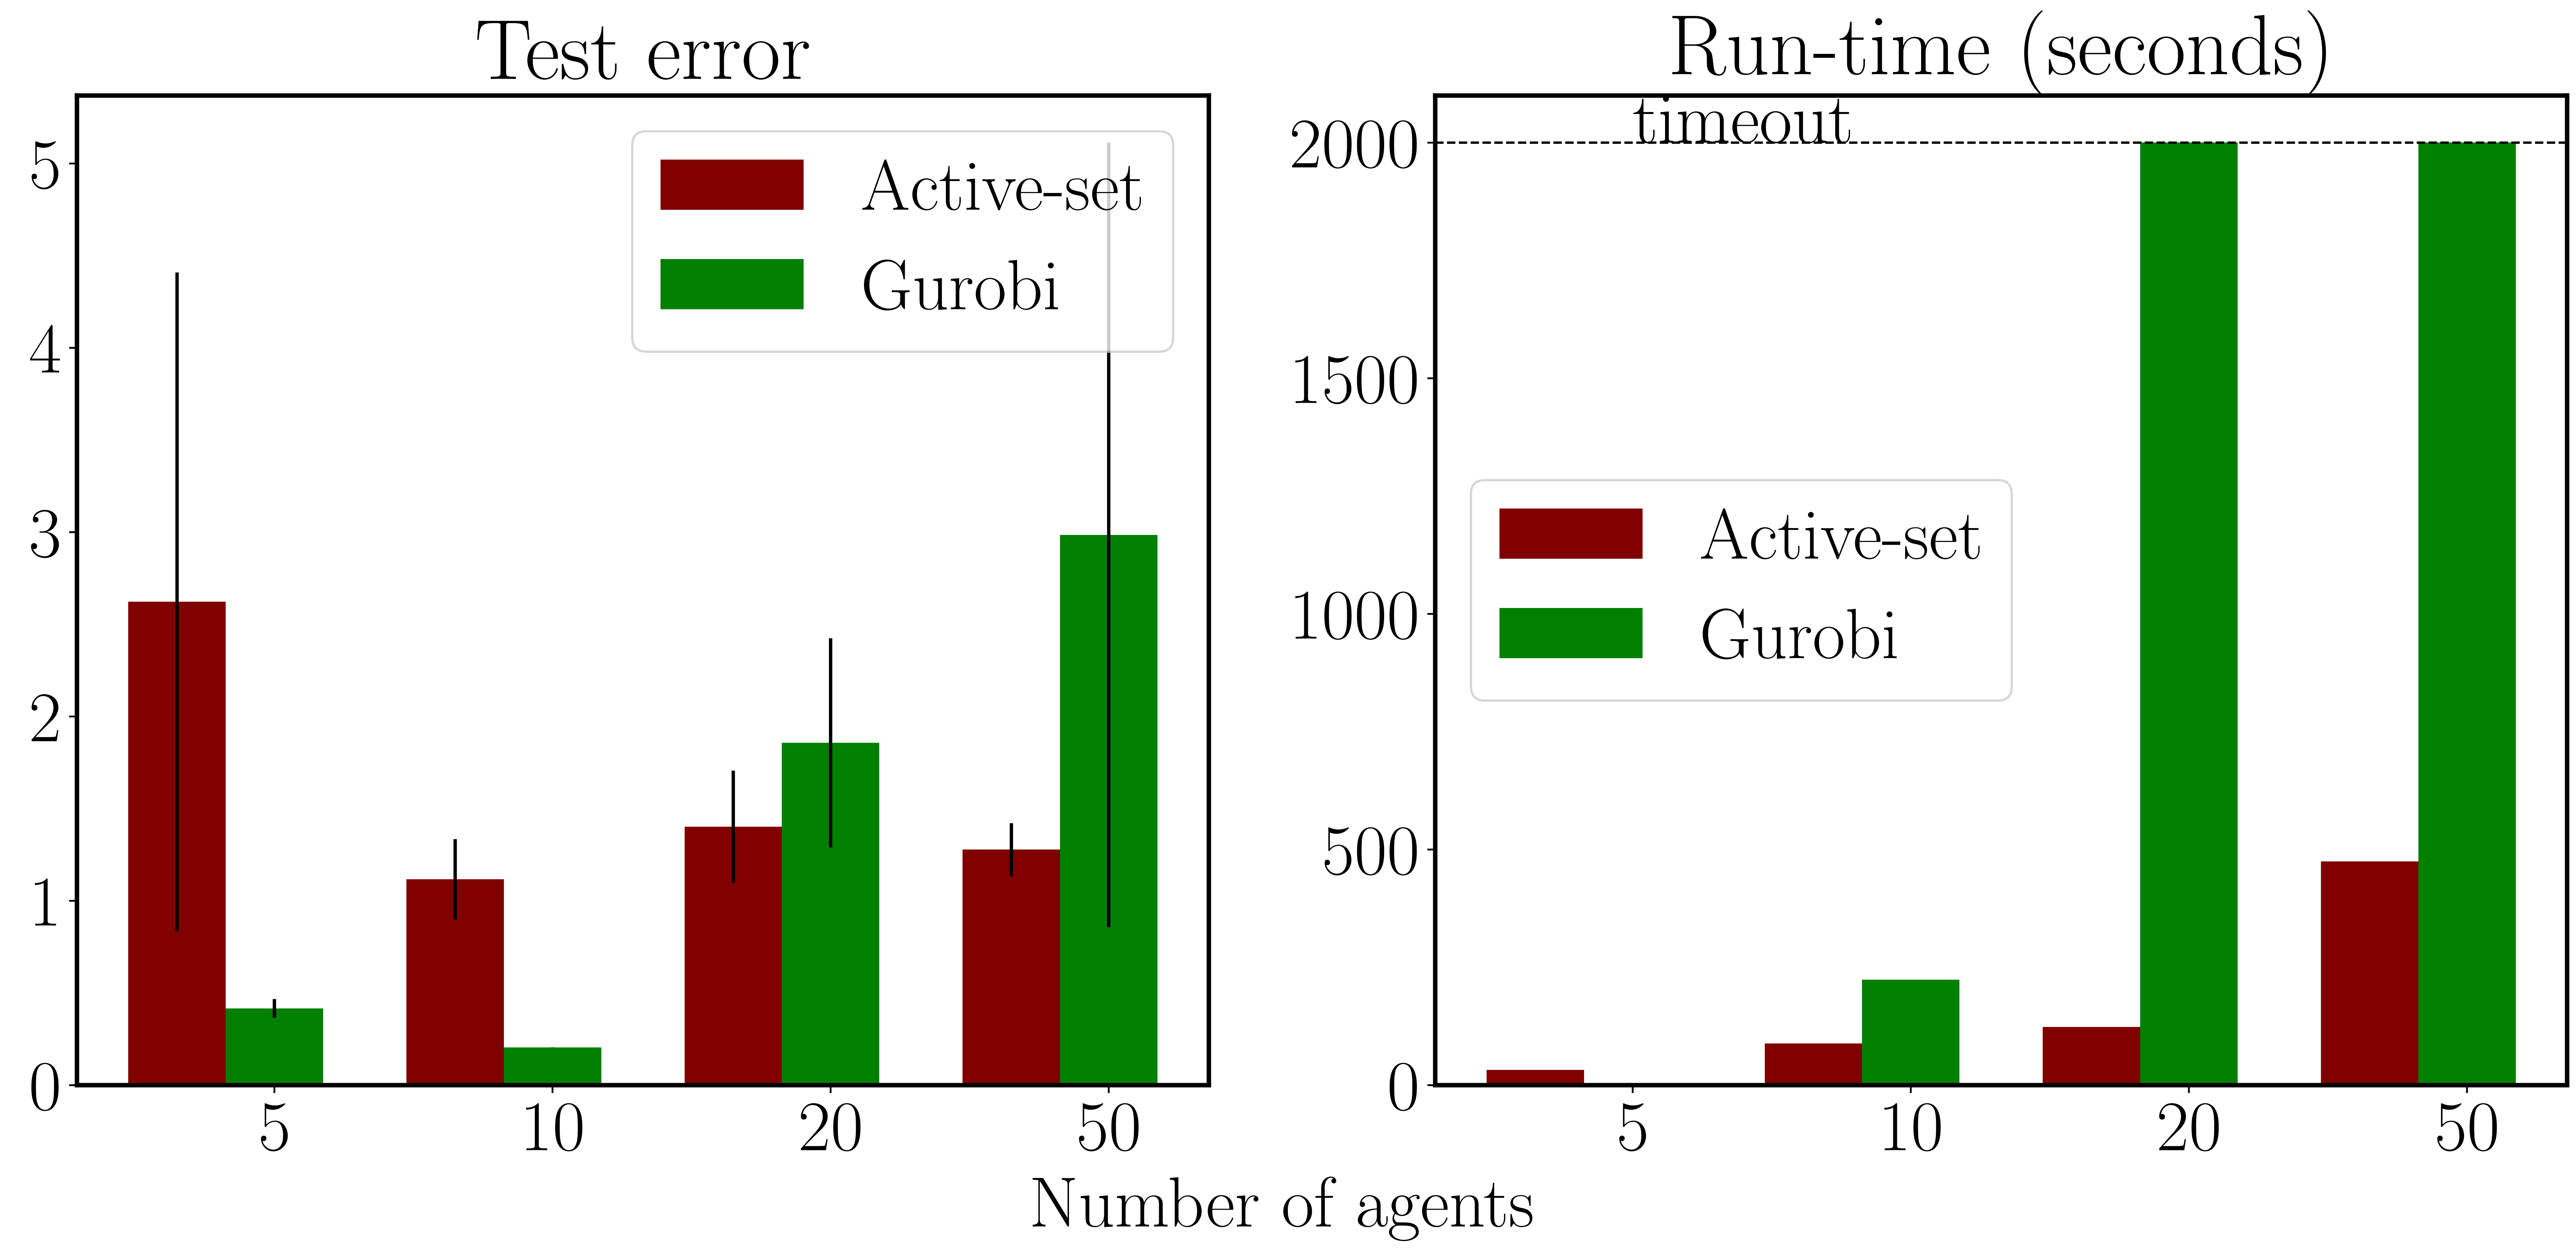

Our algorithm **scales to large datasets**

Once we know the parameters...

We would like to prescribe strategies

Decision-making is rarely an individual task

External regulators should **learn the agents' preferences** and **intervene**

Learn the Parameters

Intervene in "complex" settings

Integer Programming Games

Rosario Scatamacchia (Politecnico di Torino, Italy) and
Margarida Carvalho (Universitè de Montréal, Canada),
Andrea Lodi (Cornell Tech, USA), and
Sriram Sankaranarayanan (IIM Ahmedabad, India)




Intervene in “complex” settings

$$\begin{aligned} \max_{x_1} \quad & 6x_{11} + x_{12} \\ \text{s.t.} \quad & 3x_{11} + 2x_{12} \leq 4 \\ & x_1 \in \{0,1\}^2 \end{aligned}$$

Their "profits" **interact**

$$\begin{aligned} \max_{x_1} \quad & 6x_{11} + x_{12} - 4x_{11}x_{21} + 6x_{12}x_{22} \\ \text{s.t.} \quad & 3x_{11} + 2x_{12} \leq 4 \\ & x_1 \in \{0, 1\}^2 \end{aligned}$$

$$\begin{aligned} \max_{x_2} \quad & 4x_{21} + 2x_{22} - x_{21}x_{11} - x_{22}x_{12} \\ \text{s.t.} \quad & 2x_{21} + 3x_{22} \leq 4 \\ & x_2 \in \{0, 1\}^2 \end{aligned}$$

Knapsack Games (Carvalho et al., 2022; ***D.*** *and Scatamacchia, 2022), Critical Node Game (**D.** et al., 2023)*

And it can get more complex...

Facility Location and Design Game

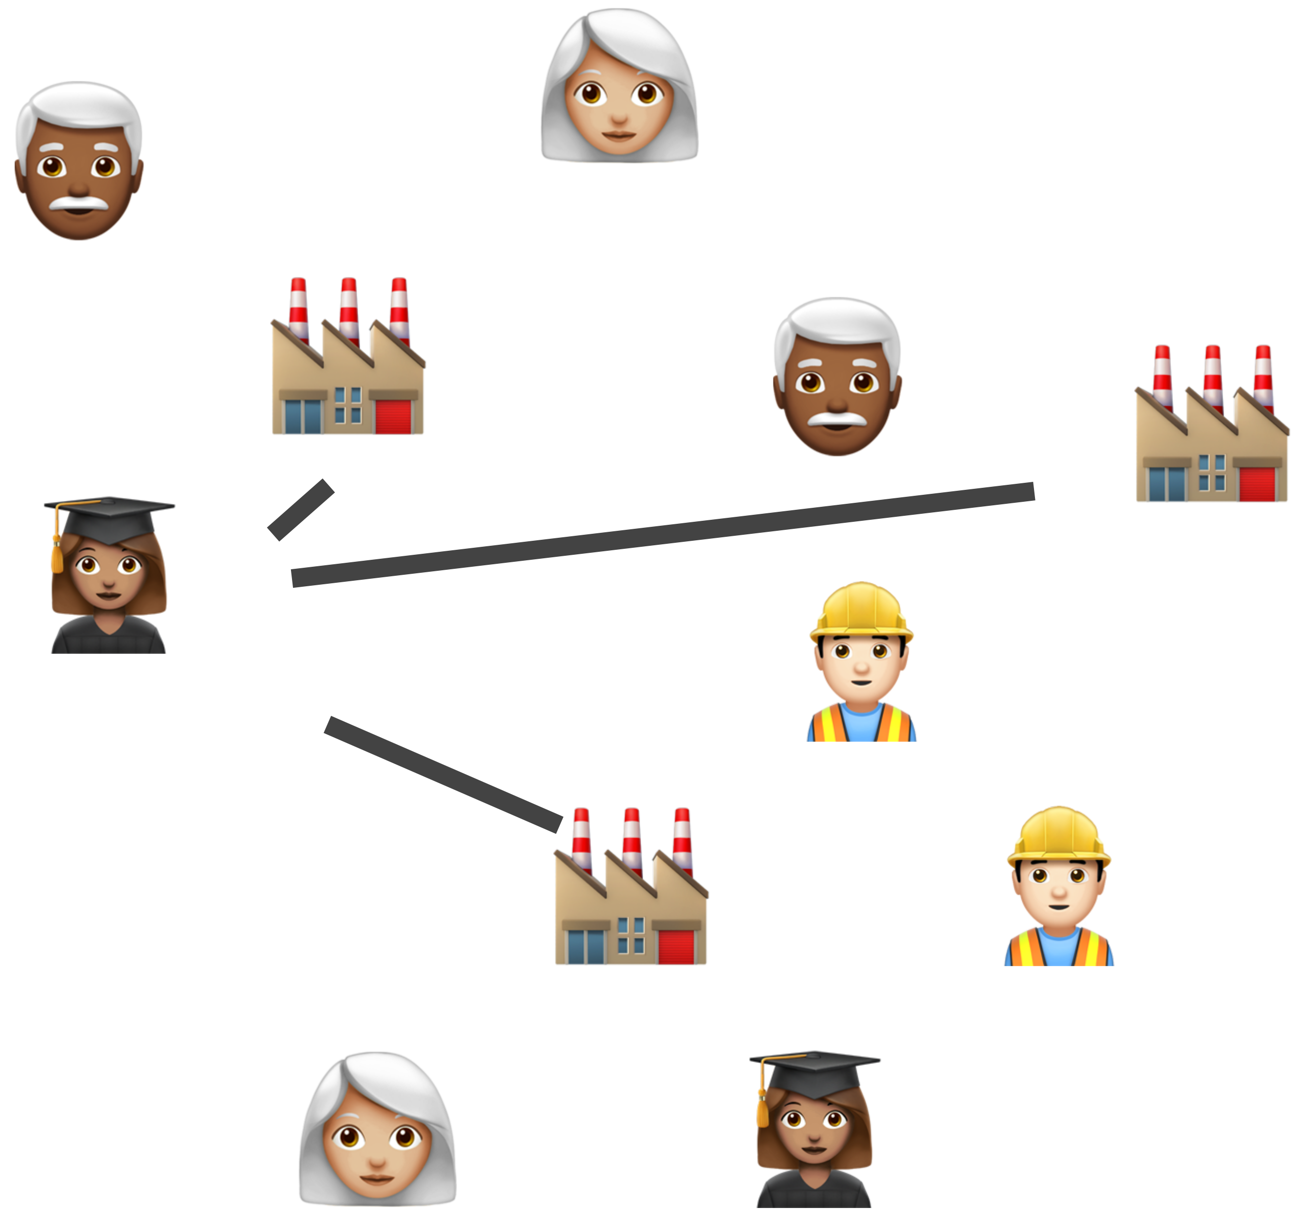

Aboolian et al. (2007),
Cronert and Minner (2020),

Sellers (players) compete for the demand of customers located in a given geographical area. Each player decides:

- **Where** to open its selling facilities
- **What** assortment to sell (i.e., what design)

$$\max_{x_i} \sum_{j \in J} w_j \frac{\sum_{l \in L} \sum_{r \in R_l} u_{iljr} x_{ilr}}{\sum_{k=1}^{n} \sum_{l \in L} \sum_{r \in R_l} u_{kljr} x_{klr}}$$ Share of customers' demand

$$\text{s.t.} \sum_{l \in L} \sum_{r \in R_l} f_{ilr} x_{ilr} \leq B_i$$ Budget

$$\sum_{r \in R_l} x_{ilr} \leq 1 \quad \forall l \in L,$$ One facility per location

$$x_{ilr} \in \{0,1\} \quad \forall l \in L, \forall r \in R_l$$ Integrality

What are these games?

An Integer Programming Game (IPG) is a *simultaneous one-shot (static)* game among n players where each player $i = 1,\ldots,n$ solves

$$\min_{x_i}\{u_i(x_i; x_{-i}) : x_i \in \mathcal{X}_i\}$$

$$\mathcal{X}_i := \{A_i x_i \leq b_i, \;\; x_i \in \mathbb{Z}^{\alpha_i} \times \mathbb{R}^{\beta_i}\}$$

There is **common knowledge of rationality**, i.e., each player is **rational** and there is **complete information**

Köppe et al., (2011), Sagratella (2015), ***D.*** *et al (2021),* **D.** *(2022)*

Why?

Flexible Modeling

They extend traditional **resource-allocation tasks and combinatorial optimization** problems to a multi-agent setting

Indivisible quantities, **fixed production costs** and **logical disjunctions** often require discrete variables

In general, they allow to model **complex operational requirements** in the players' optimization problems

Energy — Gabriel et al. (2013), David Fuller and Çelebi (2017)
Supply Chain — Anderson et al. (2017)
Assortment-Price competitions — Federgruen and Hu (2015)
Kidney Exchange Problems — Carvalho et al. (2017)
Cybersecurity — Dragotto et al. (2023)

Nash equilibria

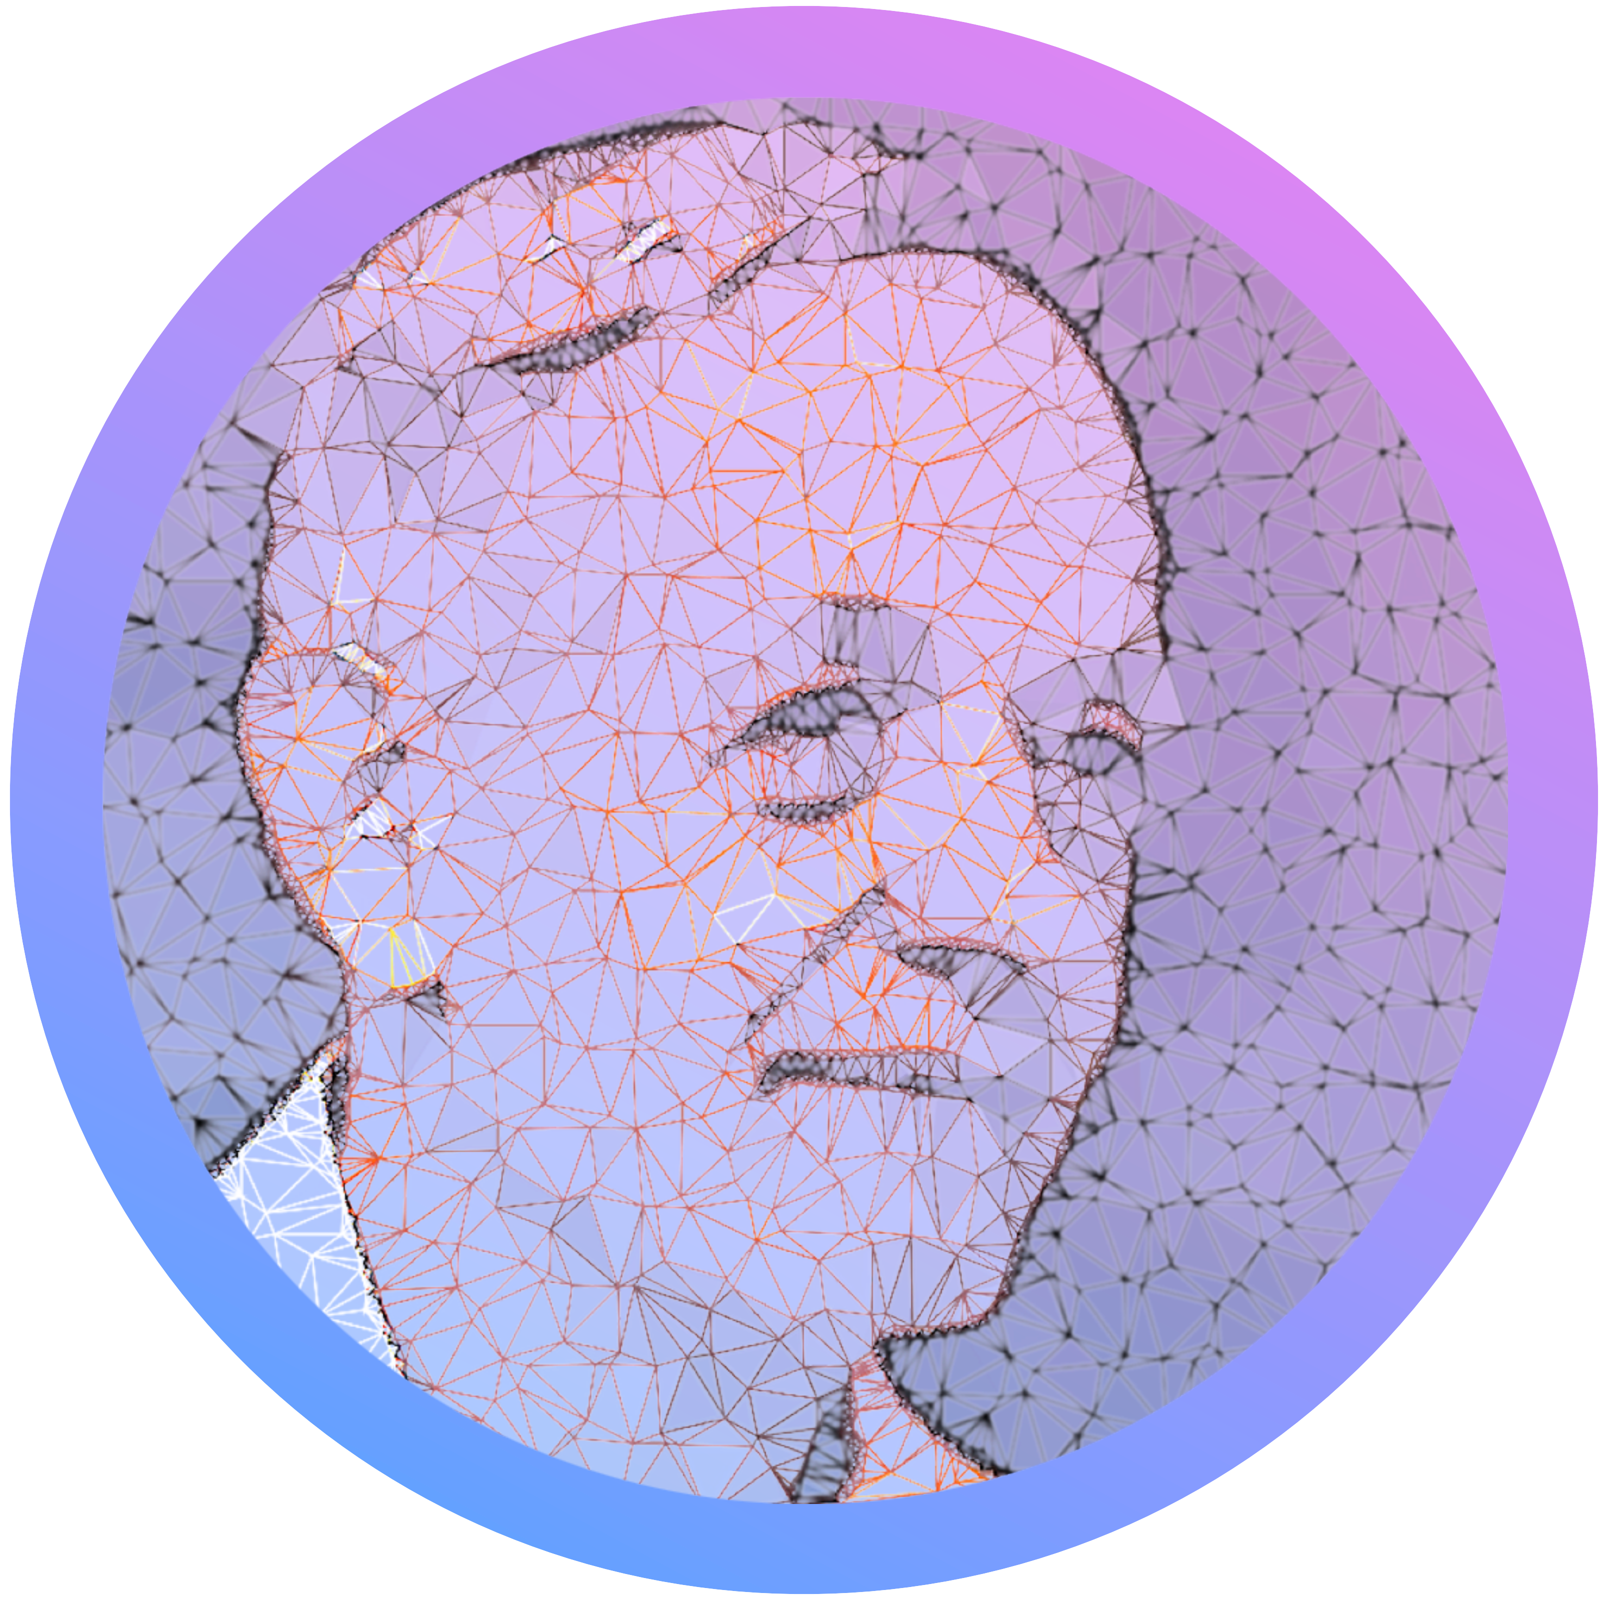

$\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)$ *is a Pure Nash Equilibrium (***PNE***) if, for any player* $\boldsymbol{i}$*,*

$$u_i(\bar{x}_i, \bar{x}_{-i}) \leq u_i(\hat{x}_i, \bar{x}_{-i}) \quad \forall \hat{x}_i \in \mathcal{X}_i$$

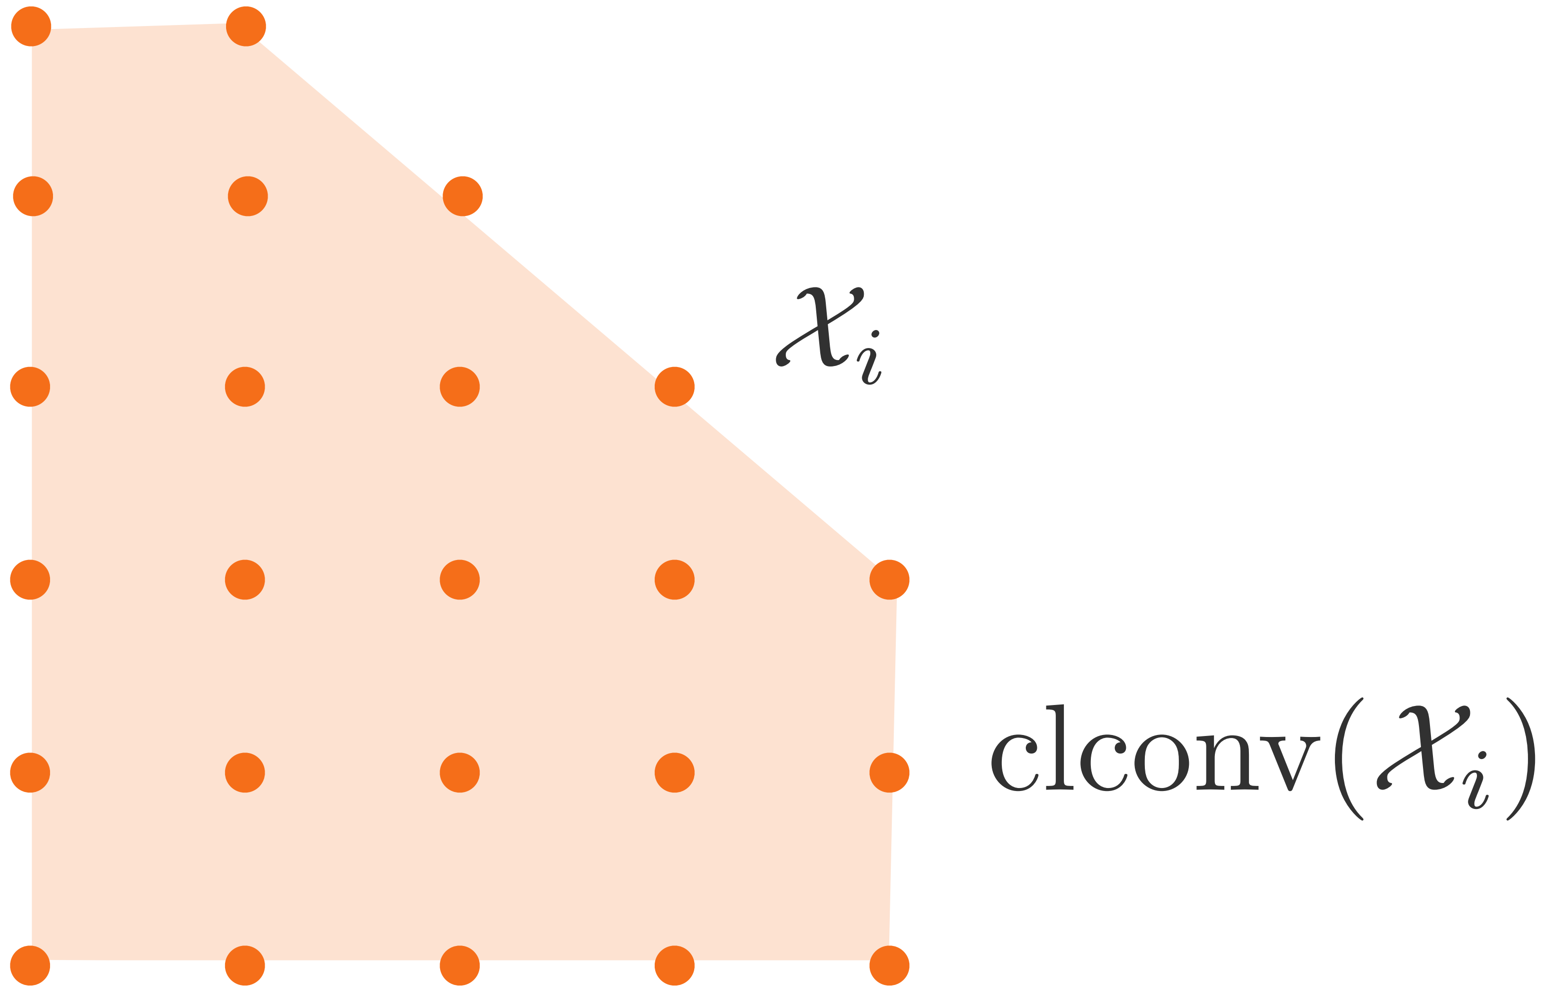

Mixed strategies = *randomizing over the convex-hull*

The goal? Zero Regrets

Flexible Modeling

Equilibria Computation

Equilibria Enumeration

Equilibria Selection

Without assuming any specific structure of the game

- Compute PoA/PoS?
- Select (optimize over) a **pure equilibrium**?
- Determine if one exists?

The goal? Zero Regrets

Flexible Modeling

Equilibria Computation

Equilibria Enumeration

Equilibria Selection

	Type of Equilibrium						
	General	Enumer.	Select	Pure	Mixed	Approx	Notes
Zero Regrets	✅	✅	✅	✅	✅	✅	**Most efficient, selection, existence, enumeration**
Koeppe et al. (2011)	✅	✅	❌	✅	❌	❌	No (practical) algorithm
Sagratella (2016)	✅	✅	❌	✅	❌	❌	Convex payoffs
Del Pia et al. (2017)	❌	❌	❌	✅	❌	❌	Problem-specific (unimodular)
Carvalho, D., Lodi, Sankaranarayanan (2020)	✅	❌	❌	❌	✅	❌	Bilinear payoffs
Cronert and Minner (2021)	✅	✅	❌	❌	✅	❌	No selection, expensive, existence?
Carvalho et al. (2022)	✅	❌	❌	❌	✅	✅	No selection/enumeration, existence?
Schwarze and Stein (2022)	✅	✅	❌	✅	❌	❌	Expensive Branch-and-Prune

Our Algorithm

Given an instance, compute *a* Nash equilibrium minimizing a function $f(x_1, \ldots, x_n)$

Our Algorithm

Given an instance, compute *a* Nash equilibrium minimizing a function $f(x_1, \ldots, x_n)$

Practical assumptions

We can tractably optimize f over $\prod_i \mathcal{X}_i$

We can express u_i as a linear function in x_i

High-level idea

1 **Initialization**

$$\mathcal{K} = \{(x,z) : x \in \prod_i \mathcal{X}_i, (x,z) \in \mathcal{L}\} \qquad \Phi := \{0 \leq 1\}$$

2 **Optimization**

$$\bar{x} = \arg \min_{x_1,\ldots,x_n,z} \{f(x,z) : (x,z) \in \mathcal{K},\ (x,z) \in \Phi\}$$

3 **Separation**

$$\tilde{x}_i = \arg\min_{x_i} \{u_i(x_i, \bar{x}_{-i}) : A_i x_i \leq b_i, x^i \in \mathbb{R}^{\alpha_i} \times \mathbb{Z}^{\beta_i}\}$$

If there is a player i so that $u_i(\tilde{x}_i, \bar{x}_{-i}) \leq u_i(\bar{x}_i, \bar{x}_{-i})$

$$\Phi = \Phi \cup \{\ u_i(\tilde{x}_i, x_{-i}) \geq u_i(x_i, x_{-i})\ \}$$ and goto 2

Else: $\bar{x}$ **is the PNE maximizing** f

Why does it work?

An inequality is an **equilibrium inequality** *if it is valid for the set of Nash equilibria*

$$u_i(\tilde{x}_i, x_{-i}) \geq u_i(x_i, x_{-i}) \quad \forall \tilde{x}_i \in \mathcal{X}_i$$

Theorem *(D. and Scatamacchia, 2022)*

$$P^e := \operatorname{conv}\left(\left\{(x,z) \in \mathcal{K} : \begin{array}{l} u_i(\tilde{x}_i, x_{-i}) \geq u_i(x_i, x_{-i}) \\ \forall \tilde{x} : \tilde{x}_i \in \mathcal{BR}(i, \tilde{x}_{-i}),\ i = 1, \dots, n \end{array}\right\}\right)$$

(1) P^e is a polyhedron

(2) P^e does not contain feasible “profiles” in its interior

(3) The extreme points of P^e are pure Nash equilibria

Applications

	Applications	Baselines	Select	Enumer.	Improvement
Knapsack Game	Packing, Assortment Optimization	Carvalho, **D.,** et al. (2021, 2022)	❌	❌	N.A.
Network Formation Games	Network design, the Internet, cloud infrastructure	Chen and Roughgarden (2006), Anshelevich, et al. (2008), Nisan et al. (2008)	✅	❌	N.A.
Facility Location Games	Retail, cloud service provisioning	Cronert and Minner (2021)	✅	❌	>50x
Cybersecurity	Ericsson Cloud Security	**D.** et al. (2023)	✅	❌	N.A.
Simultaneous-Bilevel Games	Energy, Insurance,	Carvalho, **D.,** et al. (2022)	✅	❌	N.A.

Network Formation

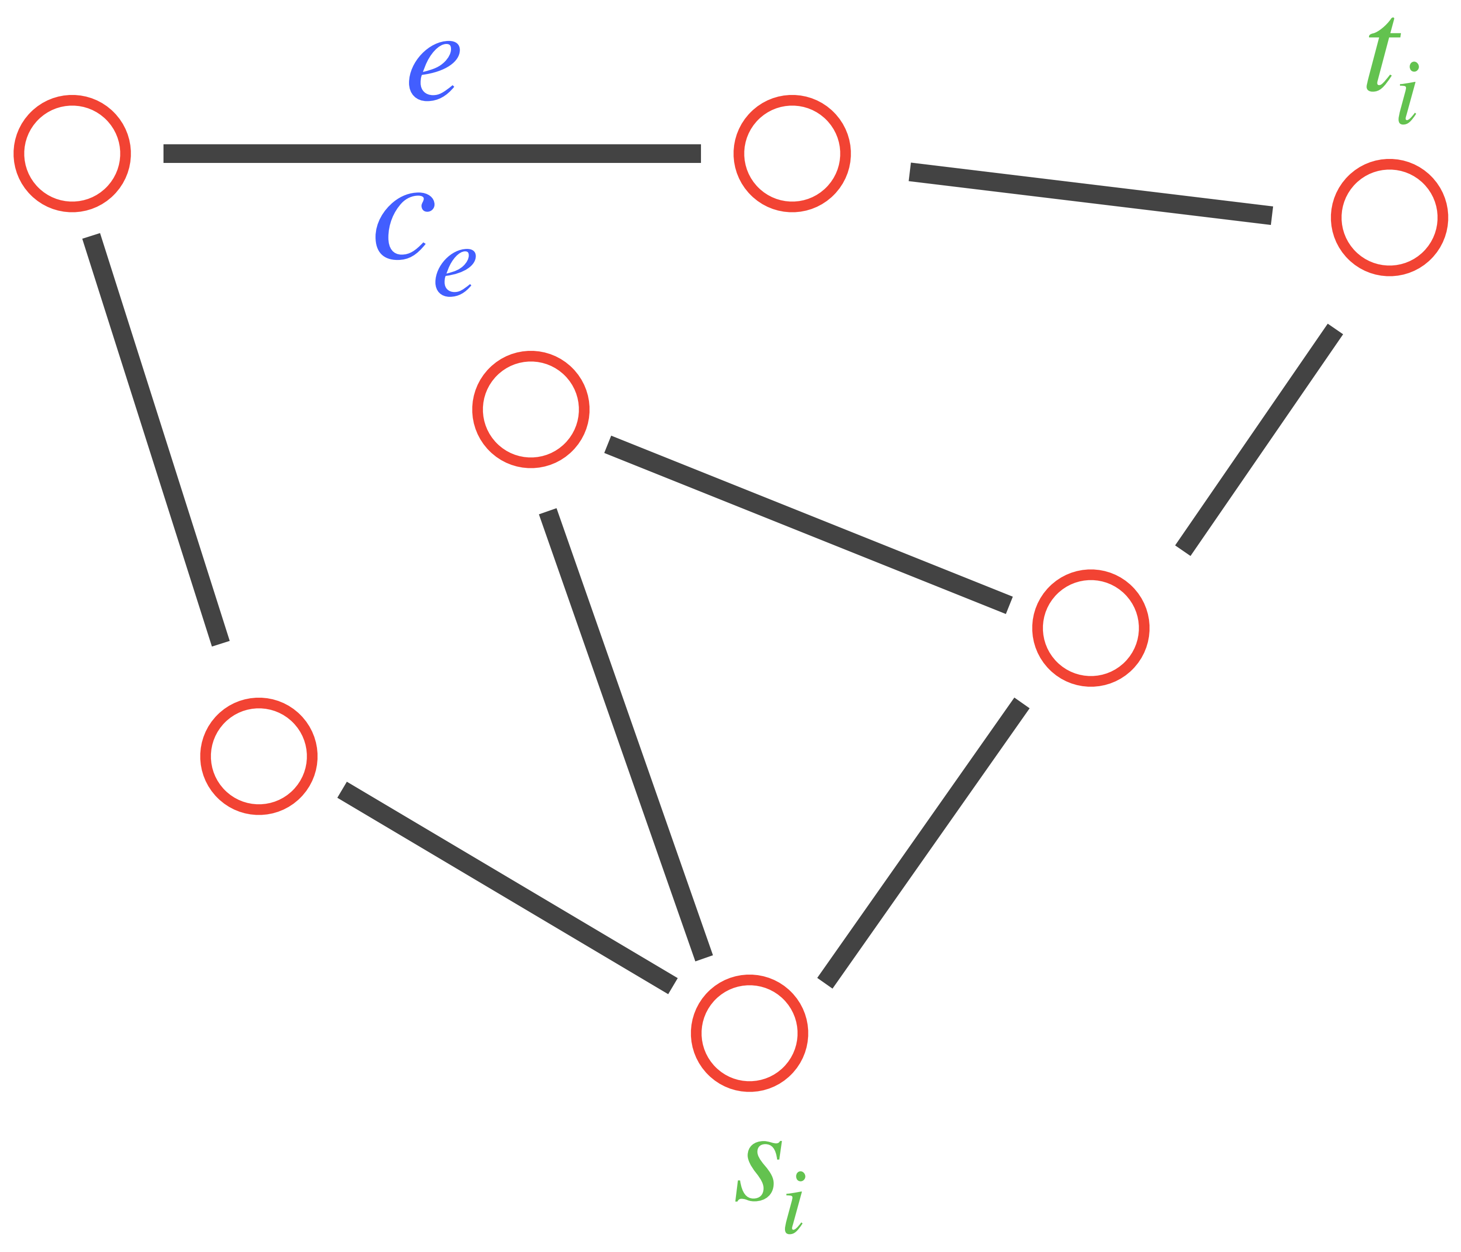

There are n players optimizing simultaneously the shortest path on a graph $G = (V, E)$ so that:

- The player i needs to go **from s_i to t_i**
- $x_{ie} = 1$ if player i selects the edge $e \in E$
- $\mathcal{X}_i$ are linear flow constraints for the path $s_i \to t_i$
- **The player i has a weight w_i**
- **Players share the cost c_e of building e**

Network **Formation**

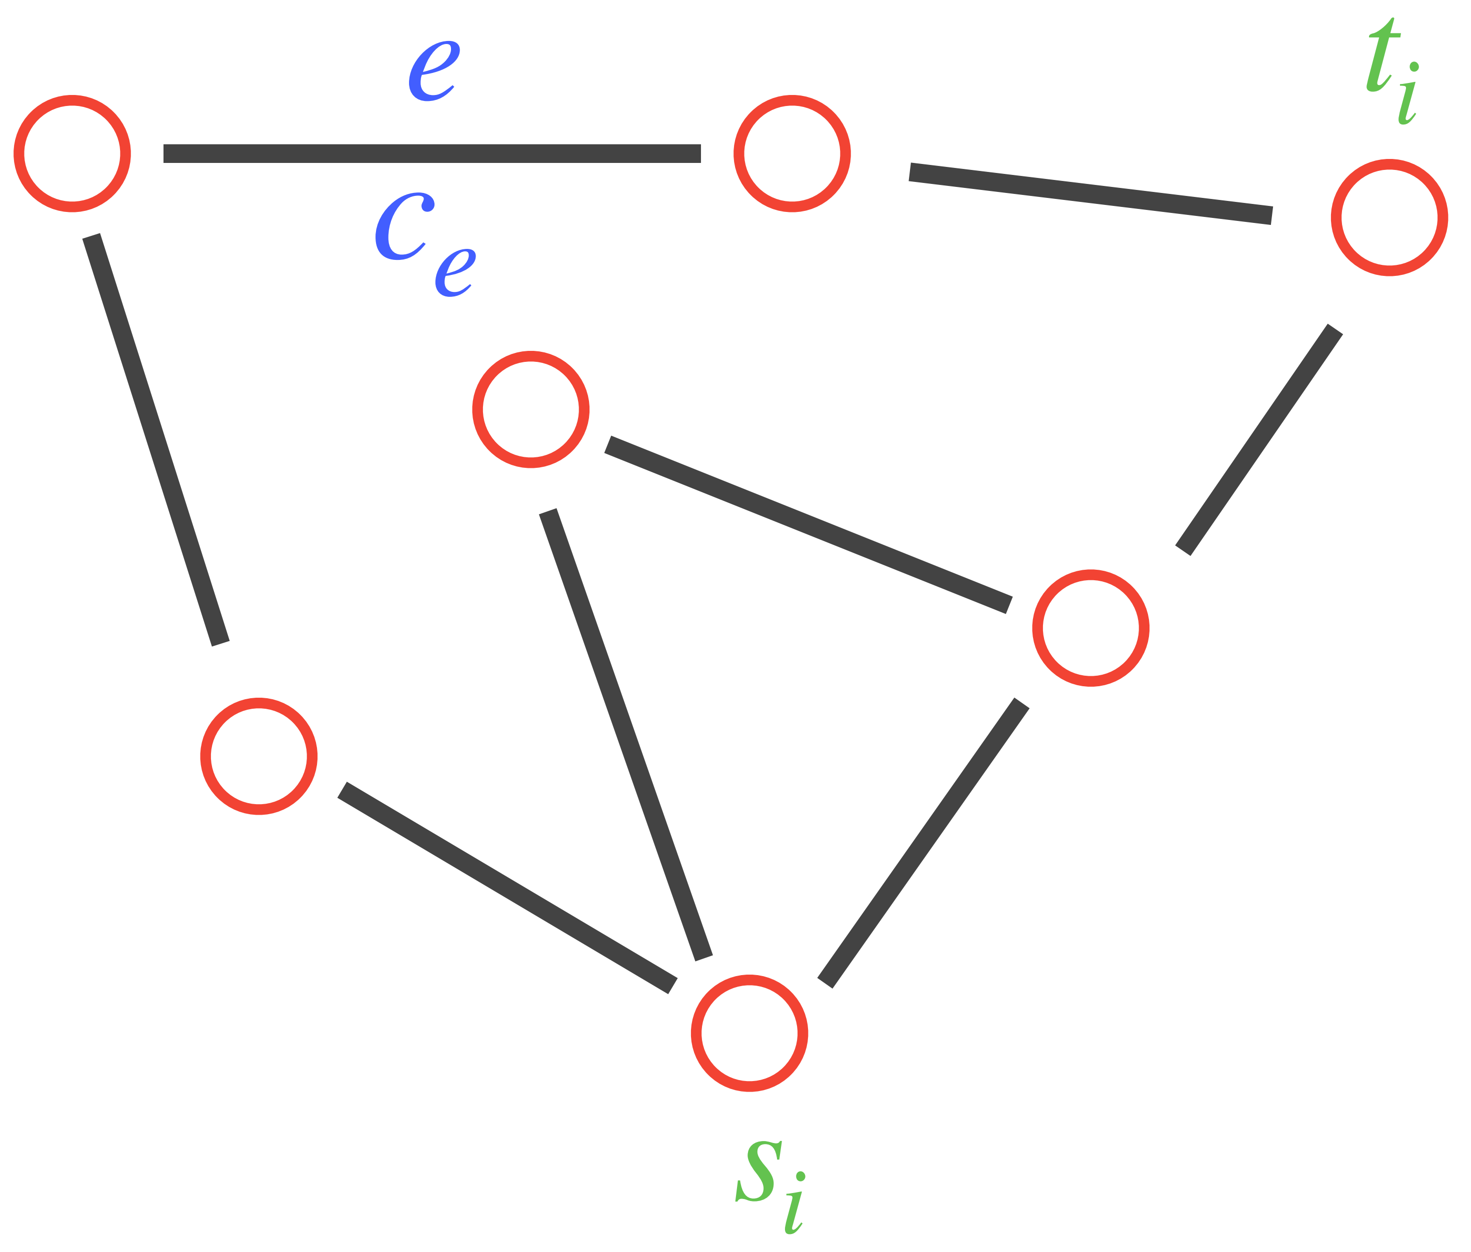

$$\min_{x_i}\{\sum_{e\in E}\frac{w_i c_e x_{ie}}{\sum_{k=1}^{n} w_k x_{ke}} : x_i \in \mathcal{X}_i\}.$$

A few remarks

- No algorithms to **select** equilibria in arbitrary NFGs
- Several bounds on PoS/PoA in some specific instances
- We consider the **weighted version** with $n = 3$

Network Formation

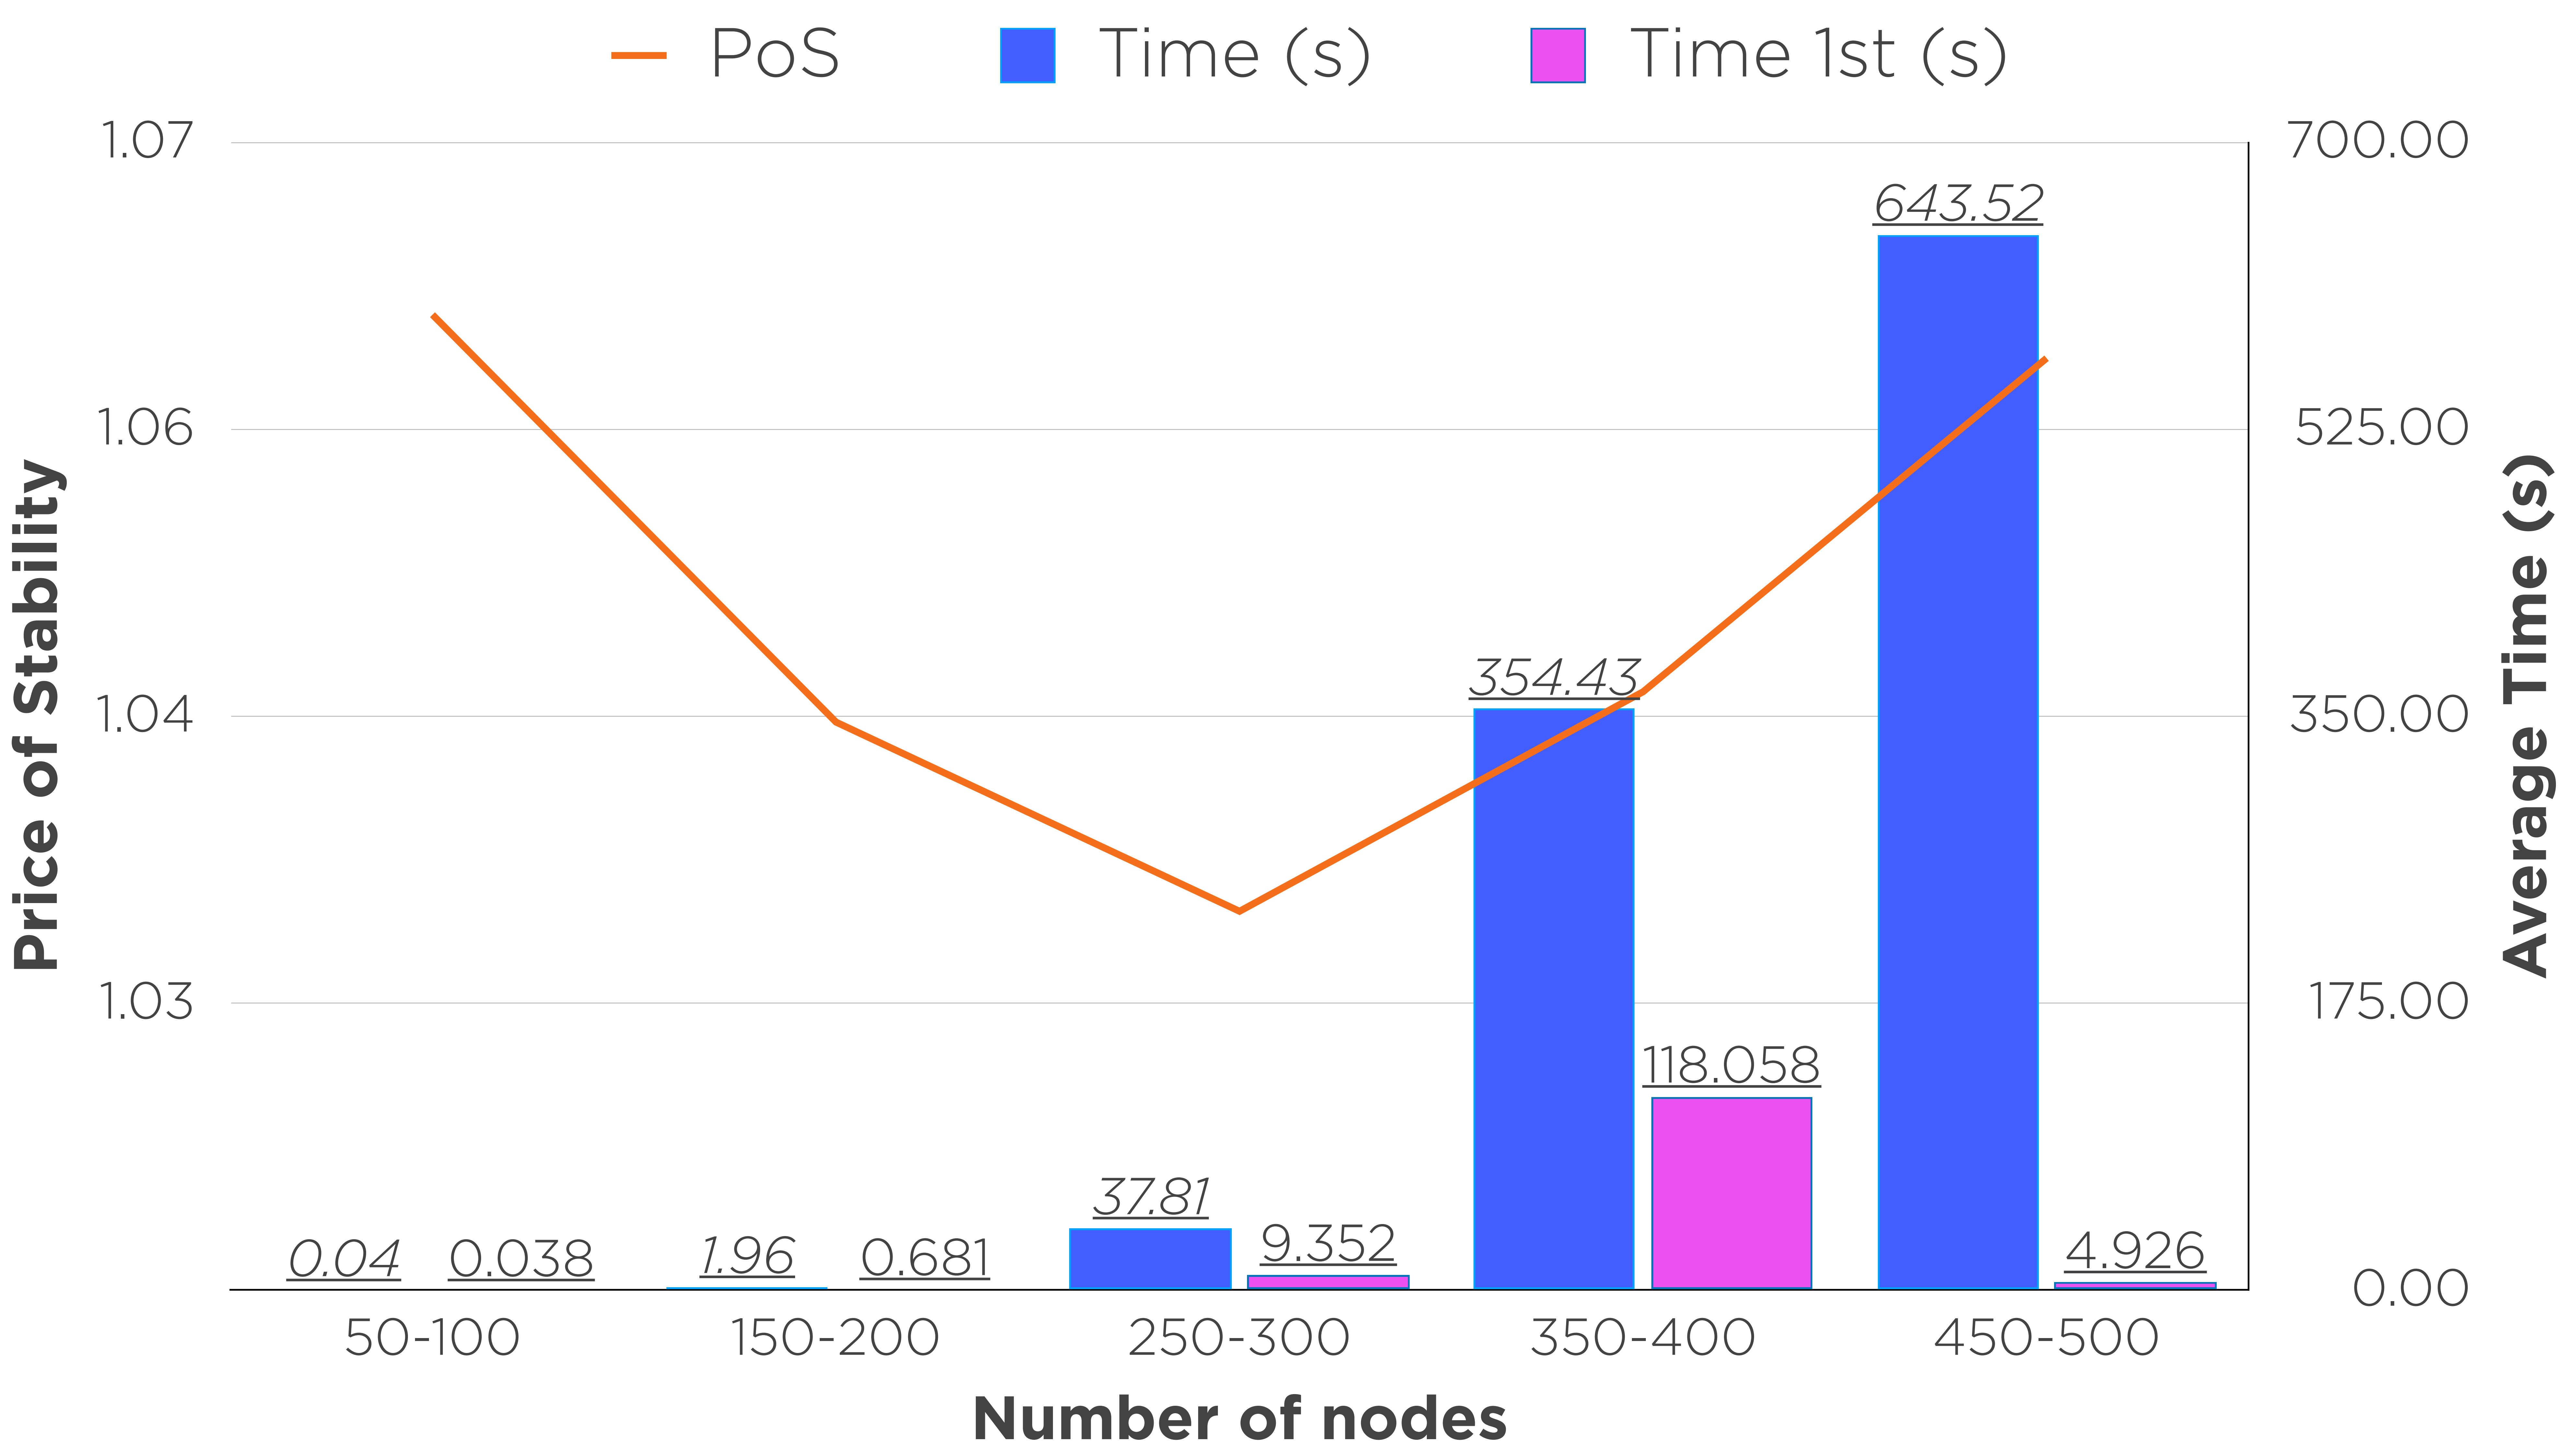

Summing up

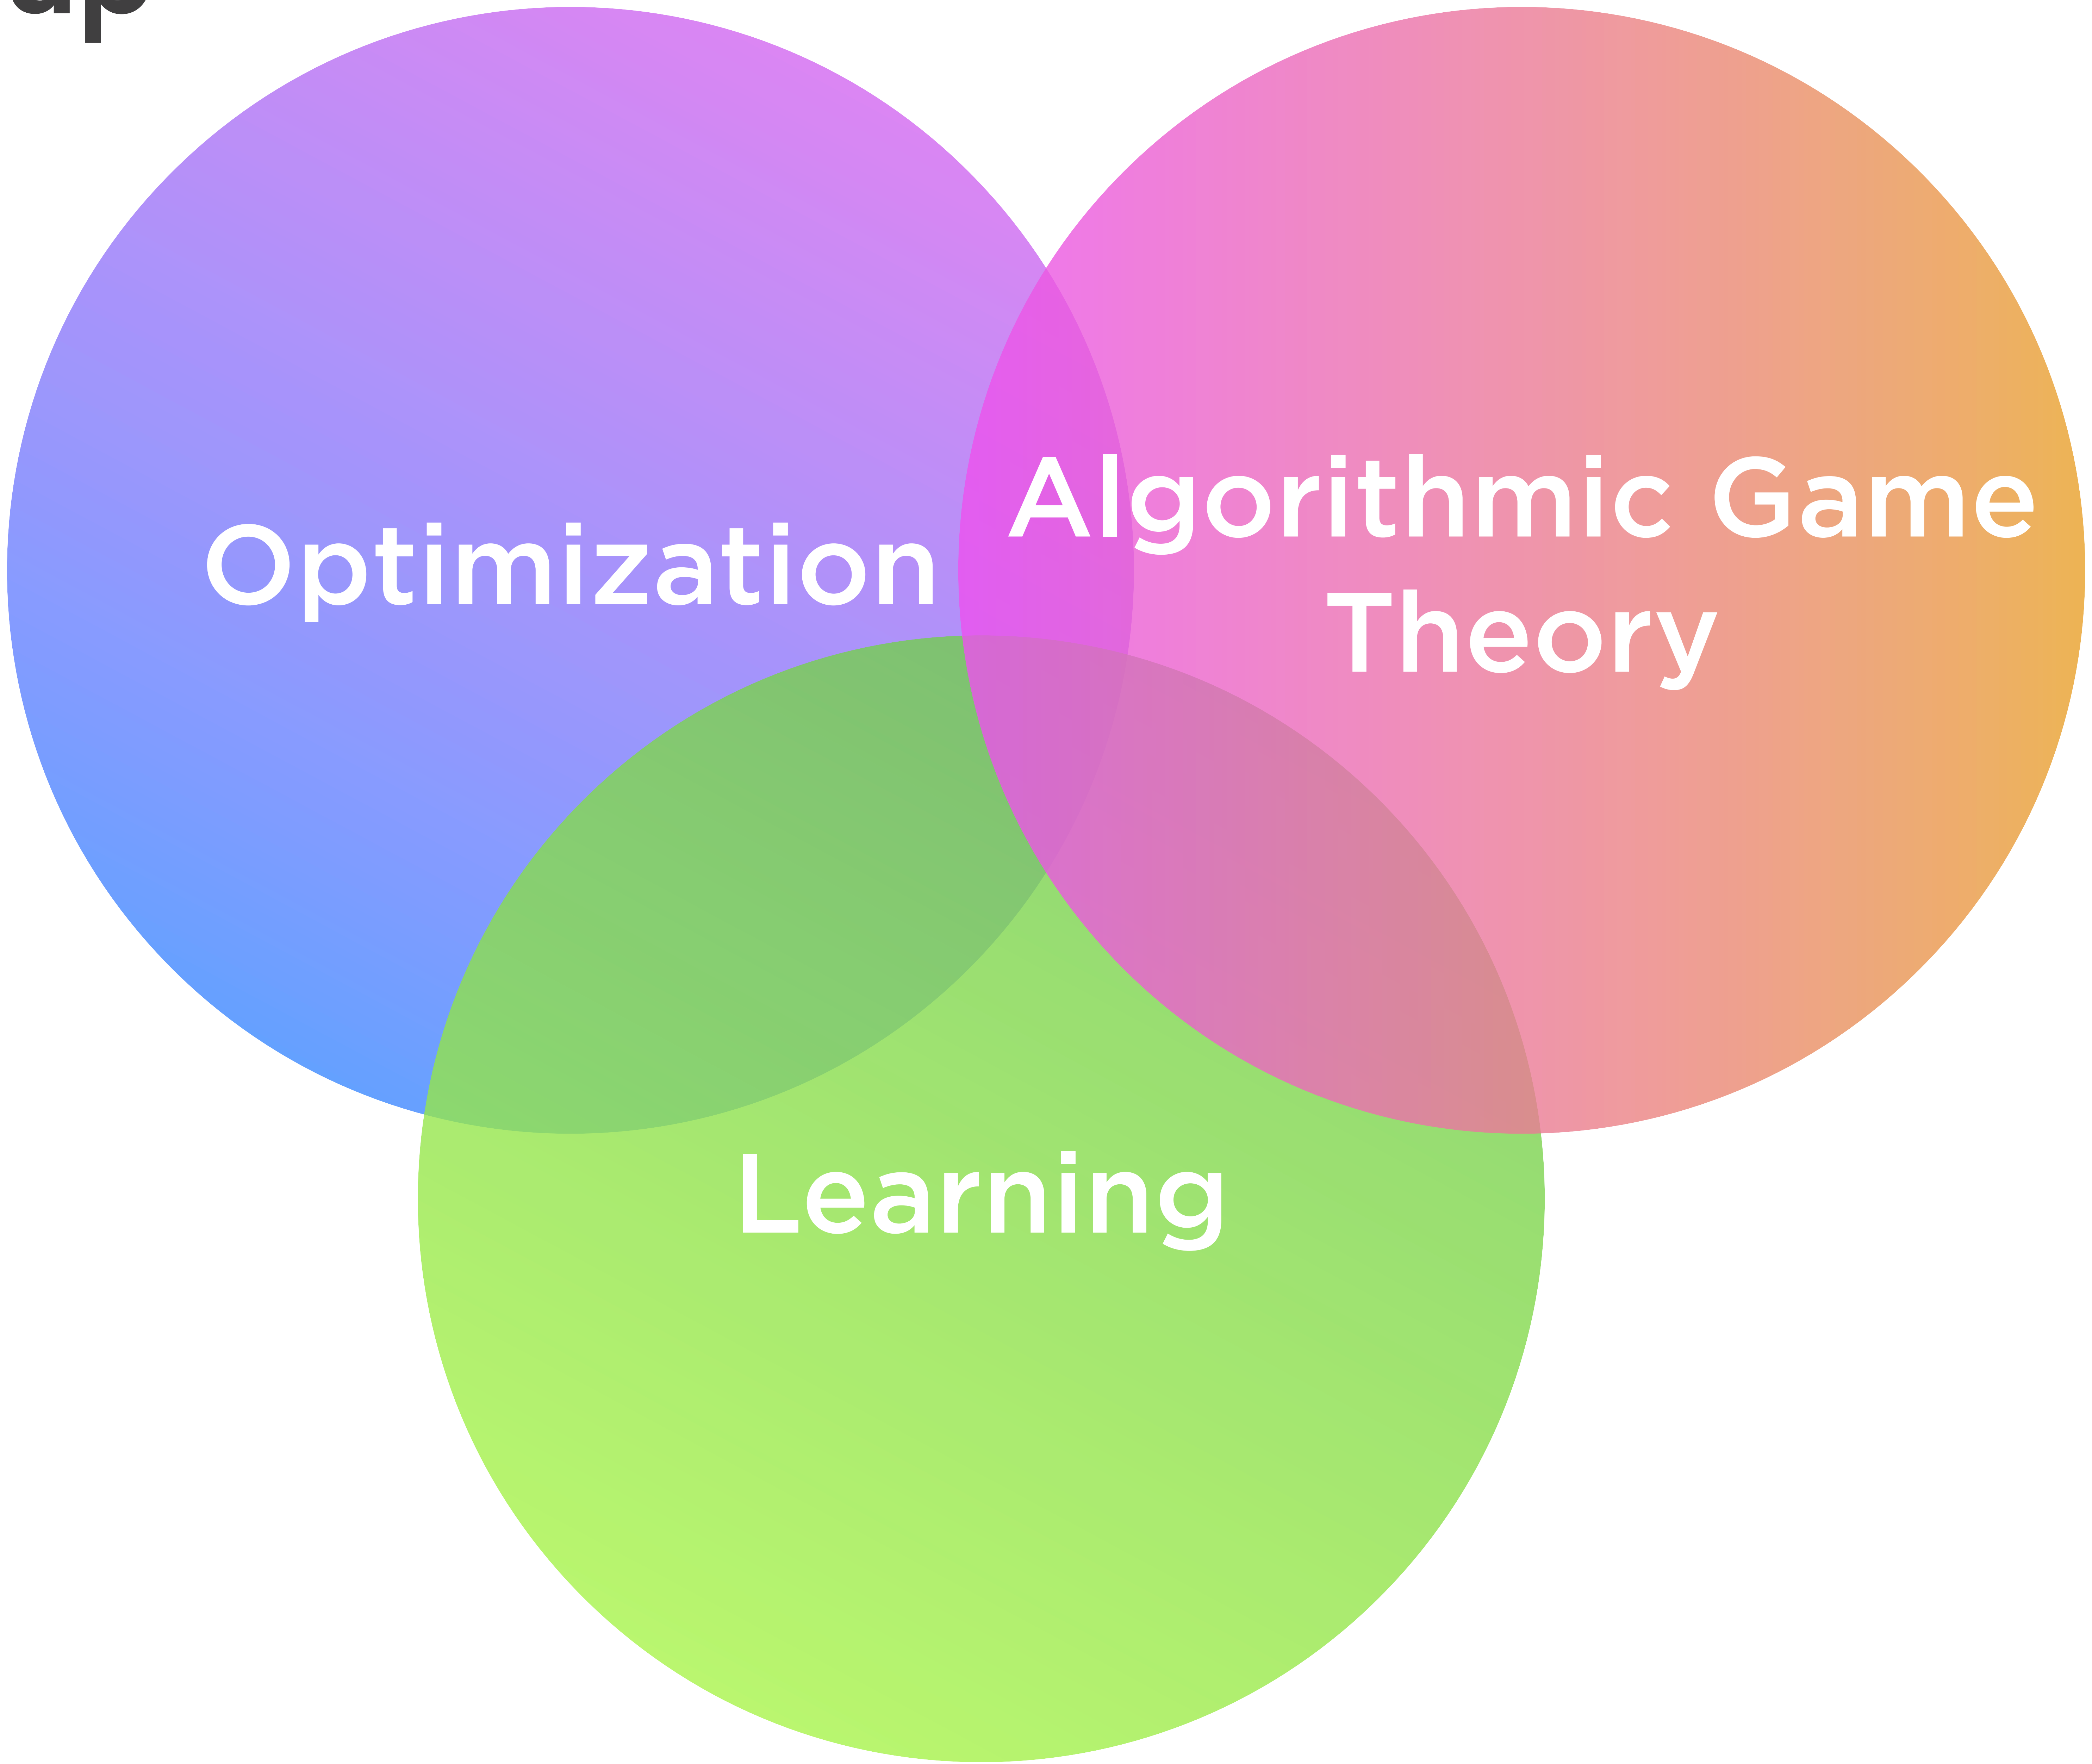

Summing up

Algorithmic Game Theory

Model complex and hierarchical structure of **interactions** among agents

Learning

Learn games' parameters from data

Optimization

Prescribe effective regulatory **interventions**

Learning Rationality in Potential Games

arXiv 2303:11188

The Zero Regrets Algorithm

arXiv 2111.06382

Integer Programming Games: A Gentle Computational Introduction

INFORMS 2023 TutORial - October 2023

DataX
Accelerating Scientific Discovery at Princeton

NUS
National University of Singapore

PRINCETON UNIVERSITY

Or even hierarchical

When Nash Meets Stackelberg (Carvalho, **D.***, et al., 2022)*

Knapsack Game