Learn-and-Play

Data, Uncertainty and Interventions

Gabriele Dragotto

ALOP Colloquium - Trier
July 17th 2023

DataX
Accelerating Scientific Discovery at Princeton
UNIVERSITÄT TRIER
PRINCETON UNIVERSITY

Commuting to work

Network Congestion

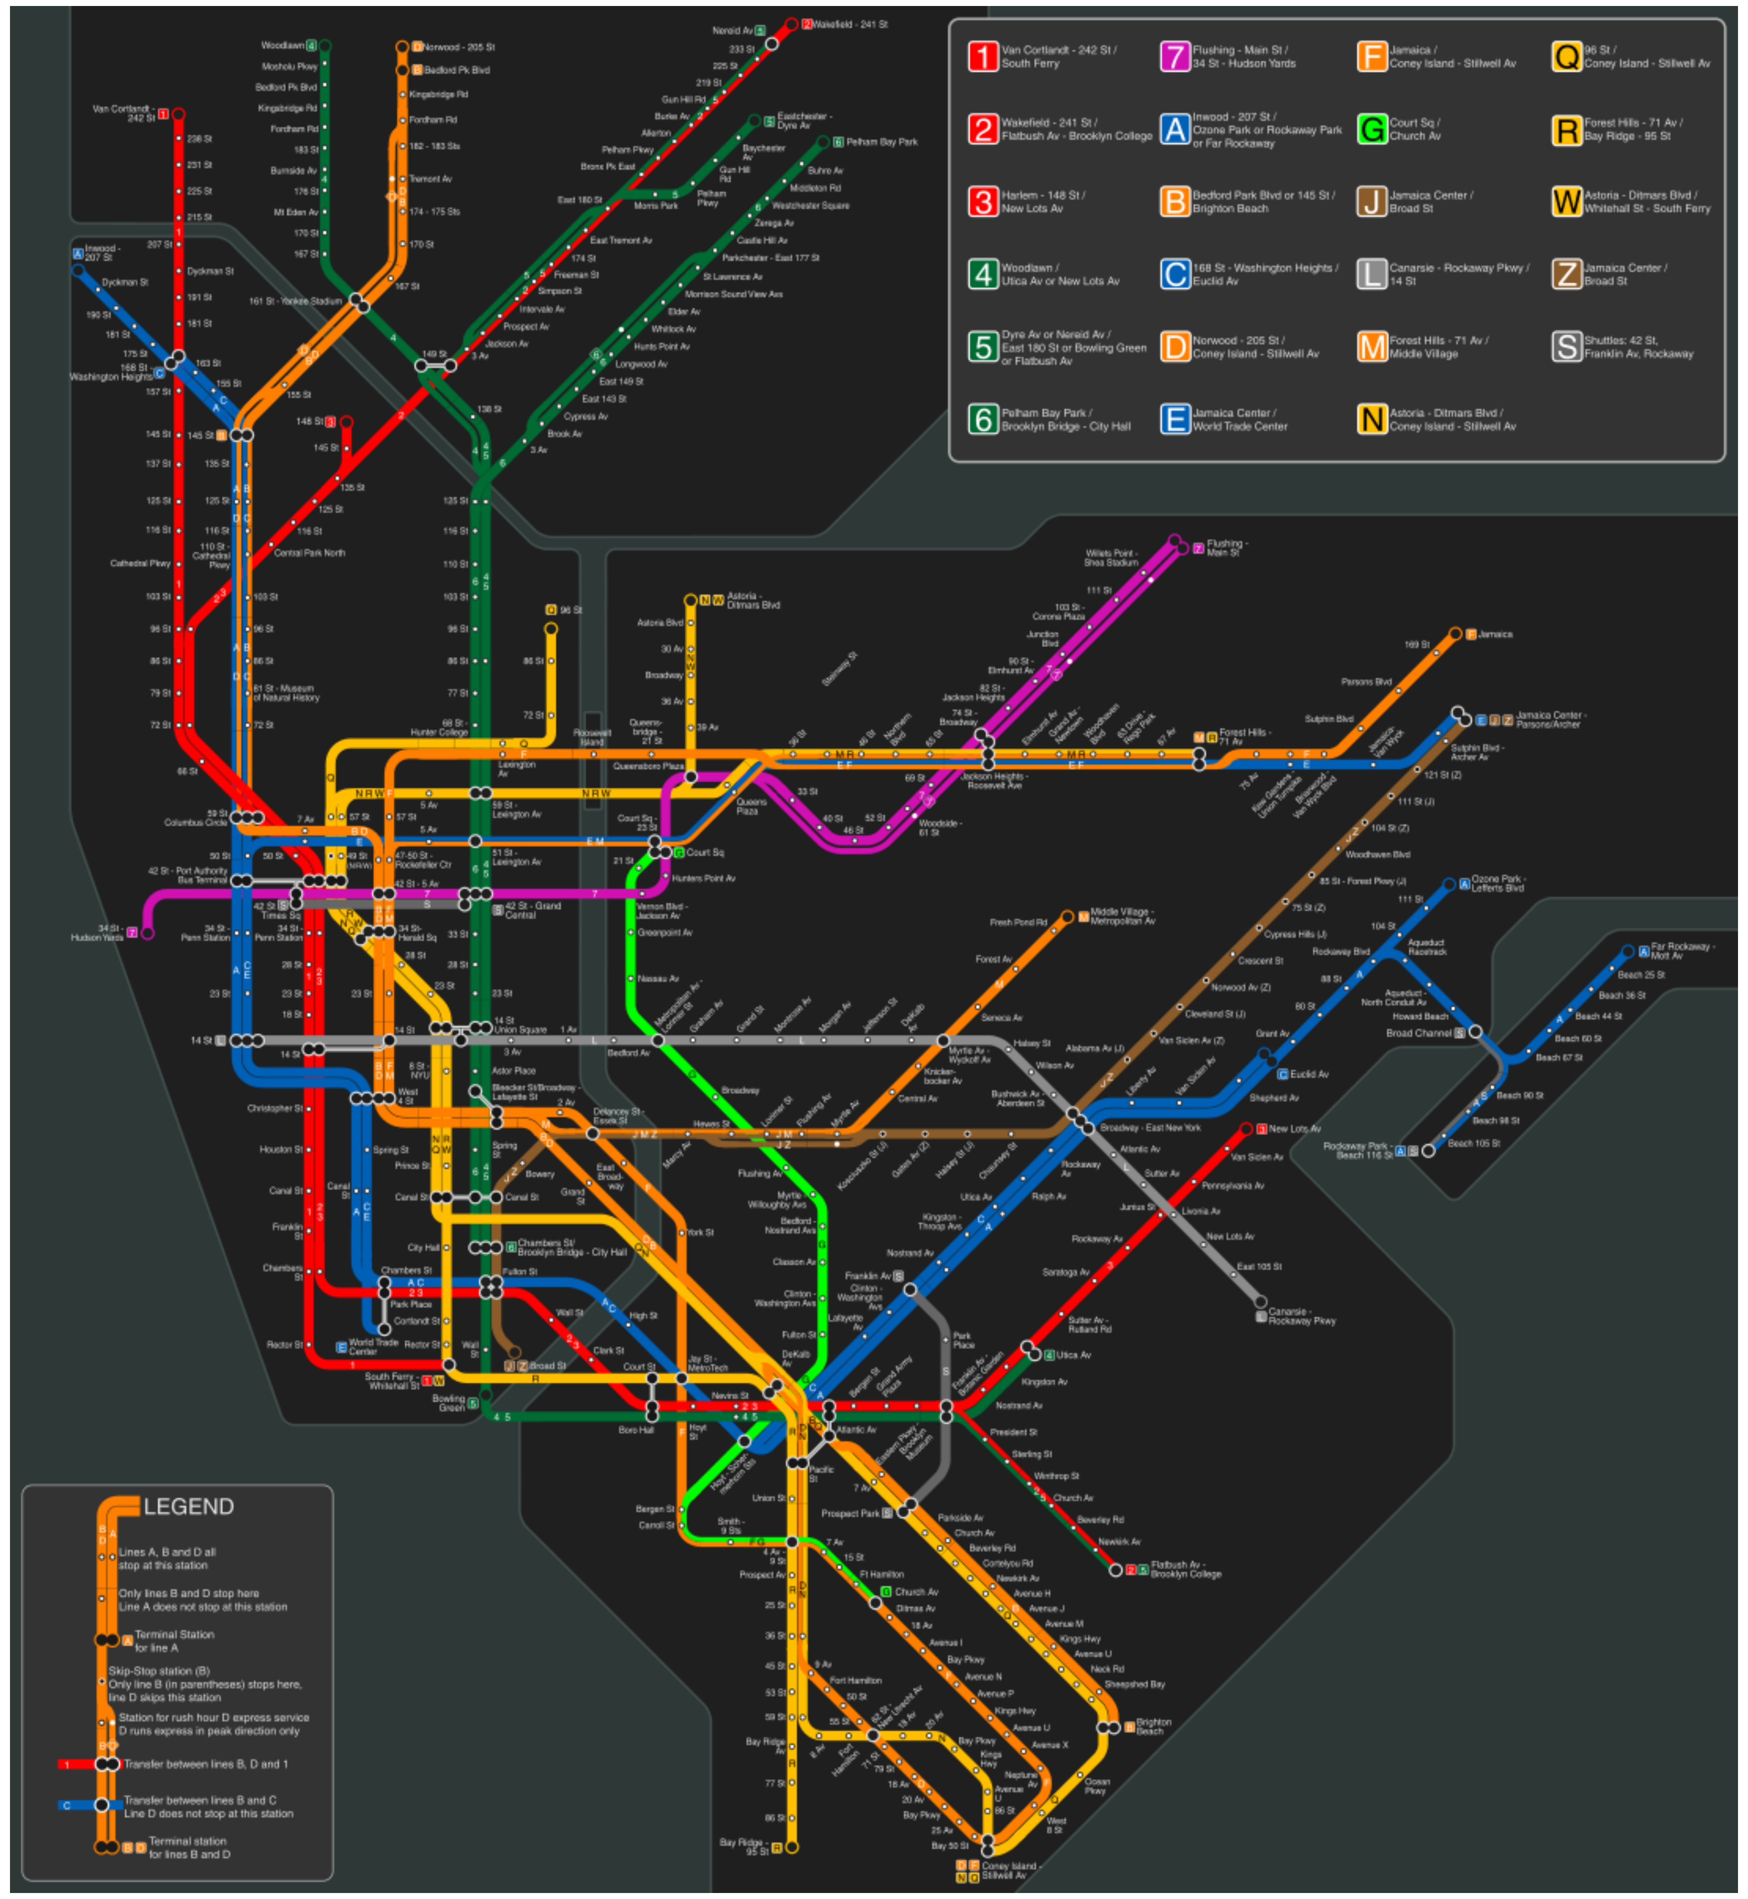

There are n players optimizing simultaneously the **shortest path** on a graph

Choices of **other players**

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

Choices of **player i**

Network Congestion

A regulator **observes the outcome** of the interaction but **is uncertain** of the agents' utilities and actions

It wants to **intervene in the game**

Decision-making is rarely an individual task

A regulator **observes the outcome** of the interaction but **is uncertain** of the agents' utilities and actions

Learn the Parameters

Intervene

Decision-making is rarely an individual task

A regulator **observes the outcome** of the interaction but **is uncertain** of the agents' utilities and actions

Learn the Parameters

Intervene

Learning Rationality in Potential Games

Stefan Clarke, Bartolomeo Stellato, and Jaime Fernandez Fisac



Learn the Parameters

Problem setup

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

Choices of **other players**

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

A set of unknown **rationality parameters**

Known and observable **context parameters**

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

Minimizing this function yields a **Nash equilibrium**

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**

$\min_{x^k,\lambda^k,\theta}$

subject to

$\mathcal{L}(\theta;\mathcal{D})$ L2 norm between target and prediction

Prediction is a Nash equilibrium,

θ belongs to a set of feasible parameters

A and **b**, and **R** and **c** are just "compact" way to represent the players constraints and objectives

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**

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

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**

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

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

The three ingredients

3 **Differentiation**

The learning problem 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***

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

Our algorithm eventually finds either a **local minimum of the smoothed loss** or a saddle point

The algorithm mimics a **stochastic gradient descent**

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

smoothened version of the loss

Network Congestion

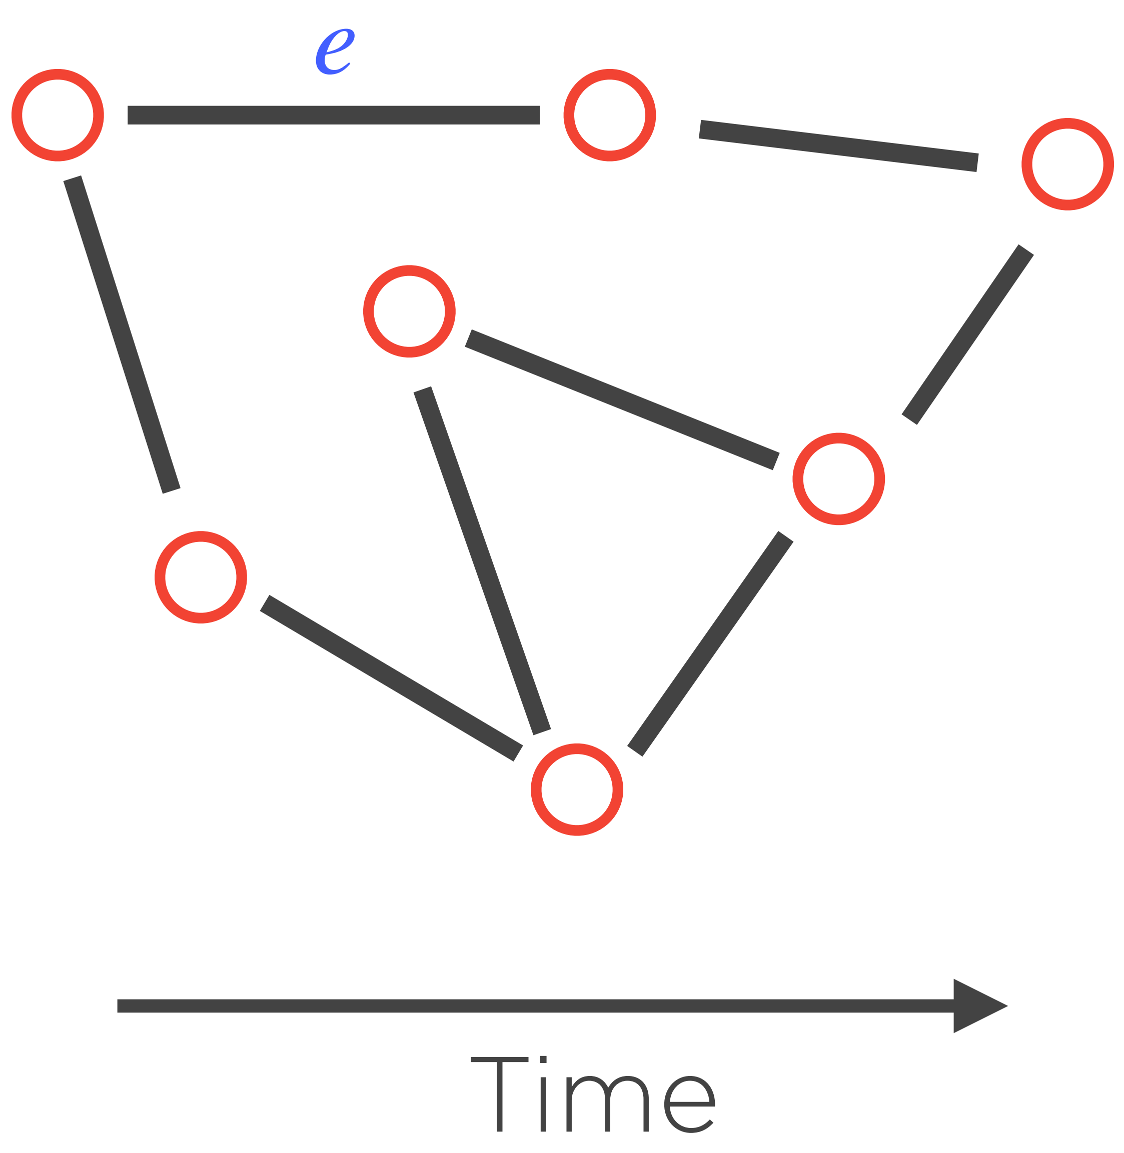

$$u_i(x_i; x_{-i}, \theta, \mu) = \sum_{e \in E} \theta_{ie}^\top l_e x_{ie}(x_{1e} + \cdots + x_{ne})$$

A set of unknown **rationality parameters**

Known and observable **context parameters**

Personal preferences

Traffic, weather, road conditions

Network Congestion

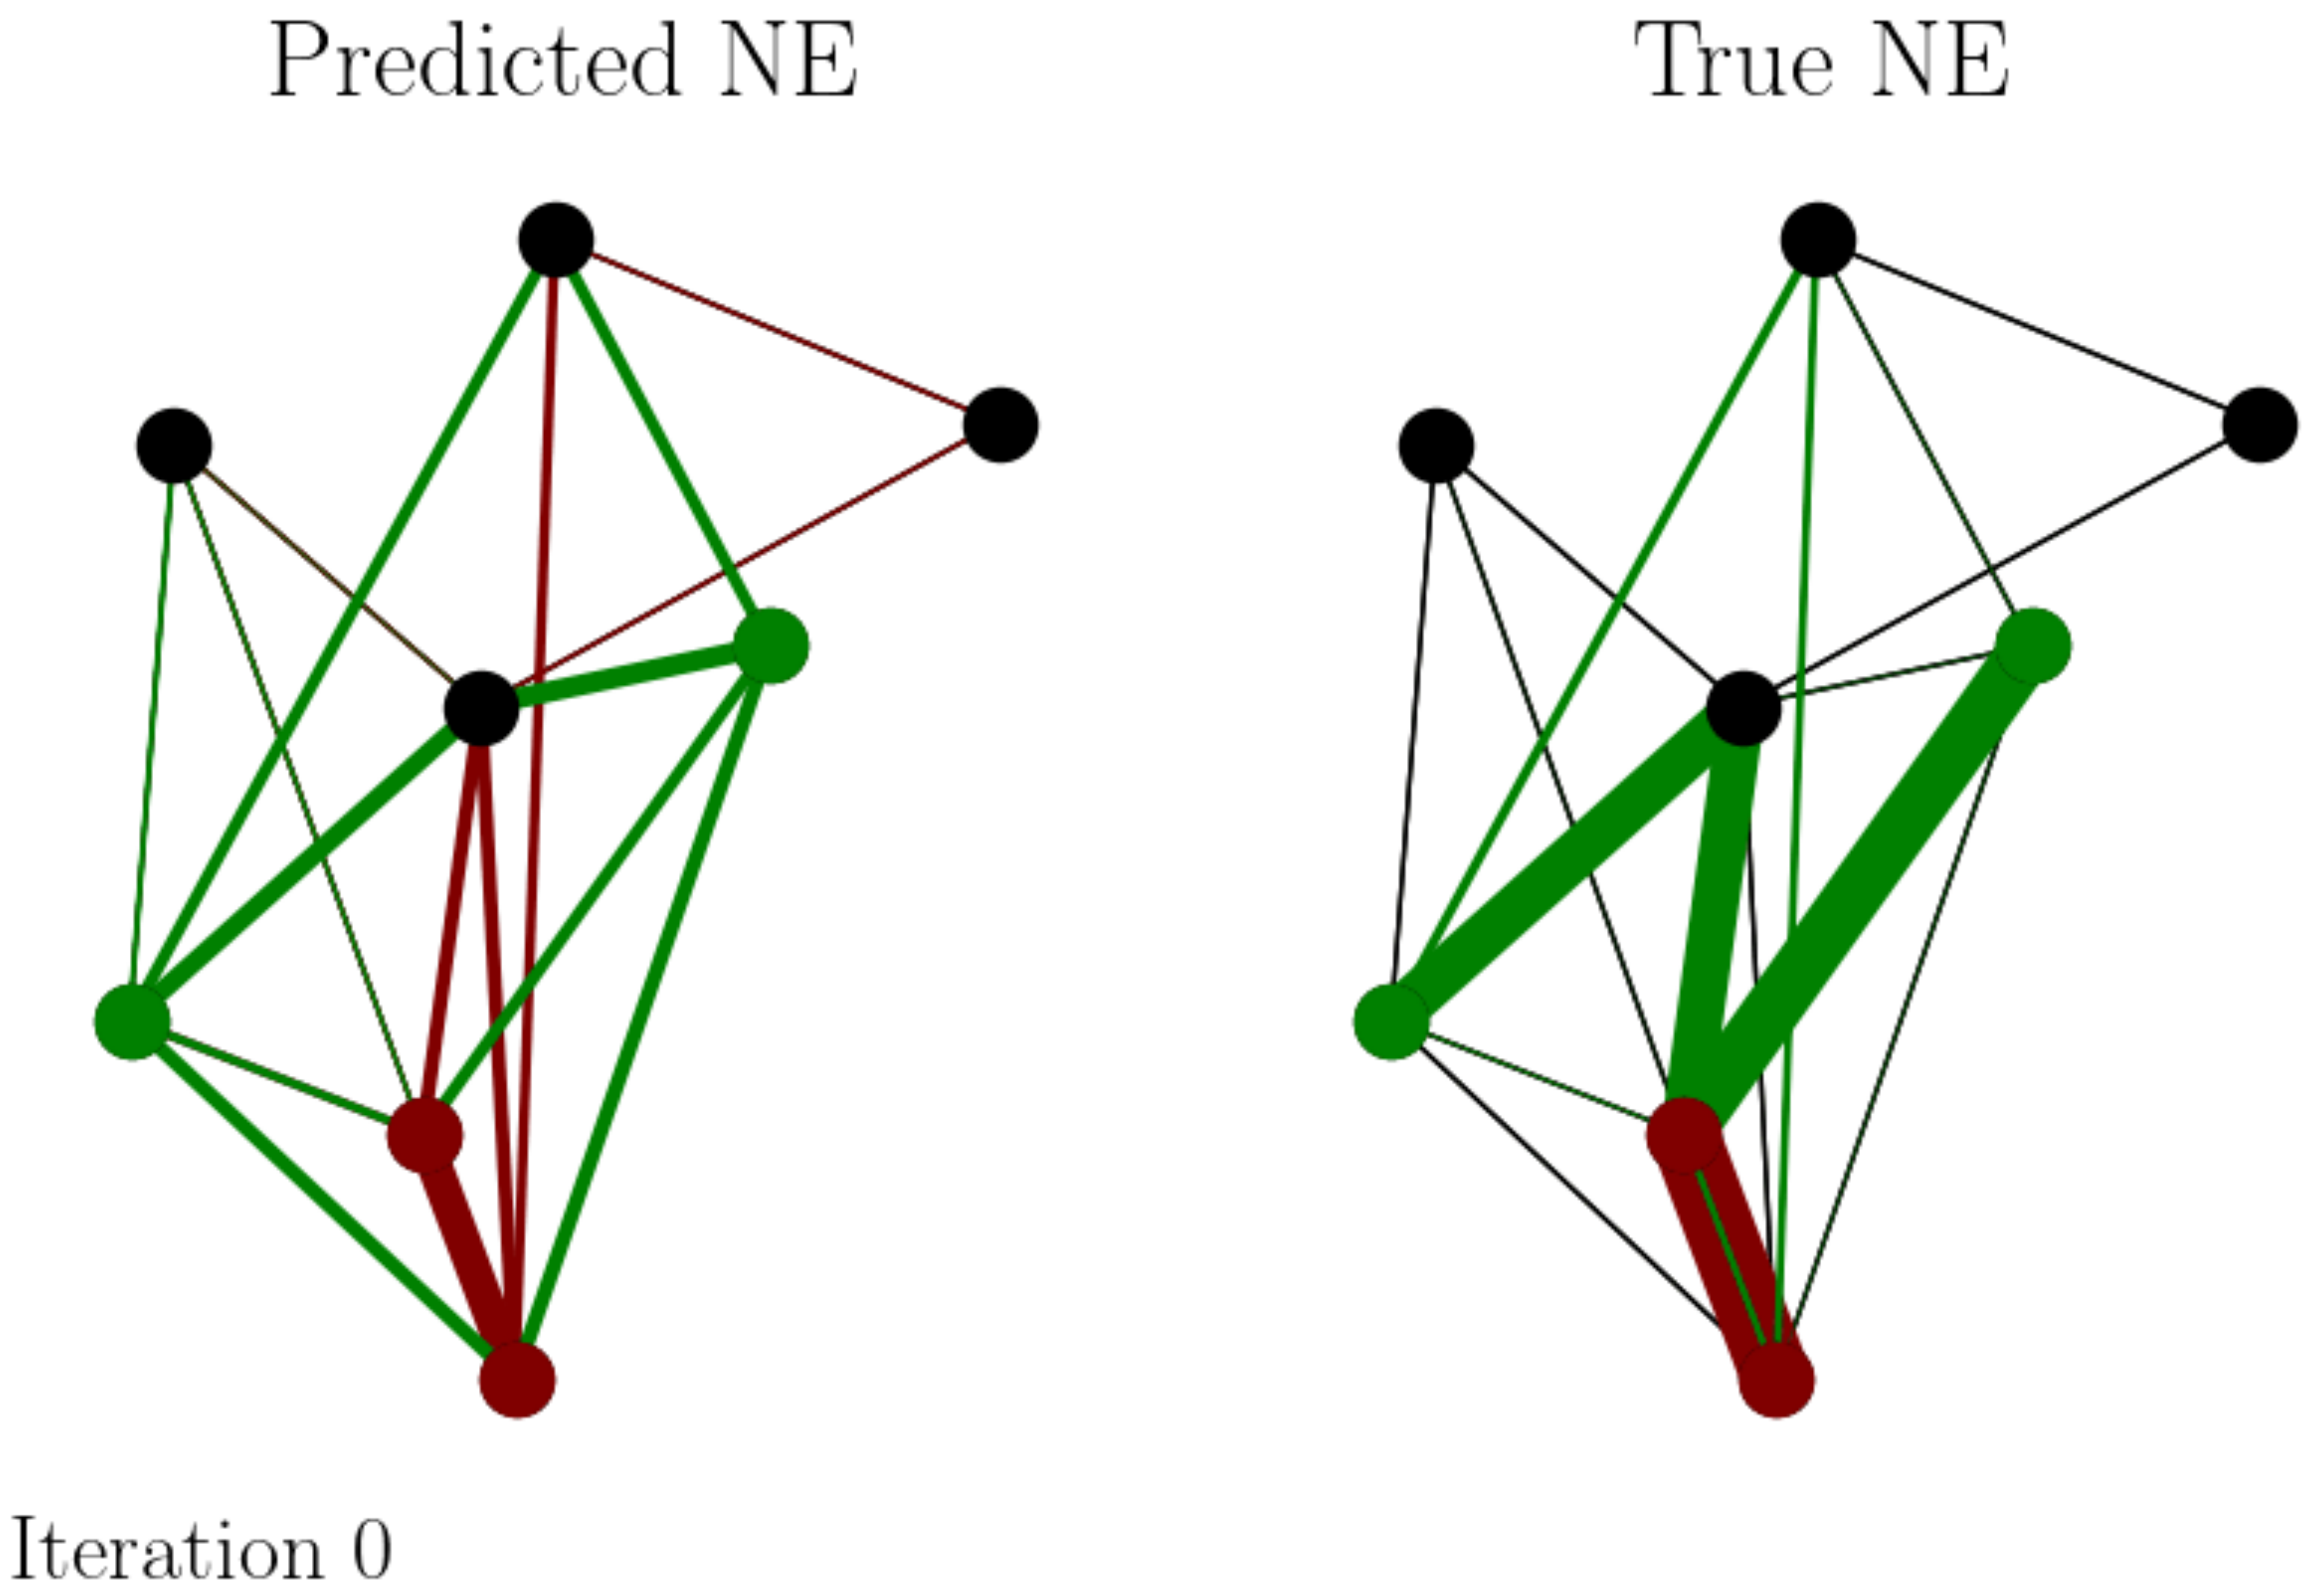

Network Congestion

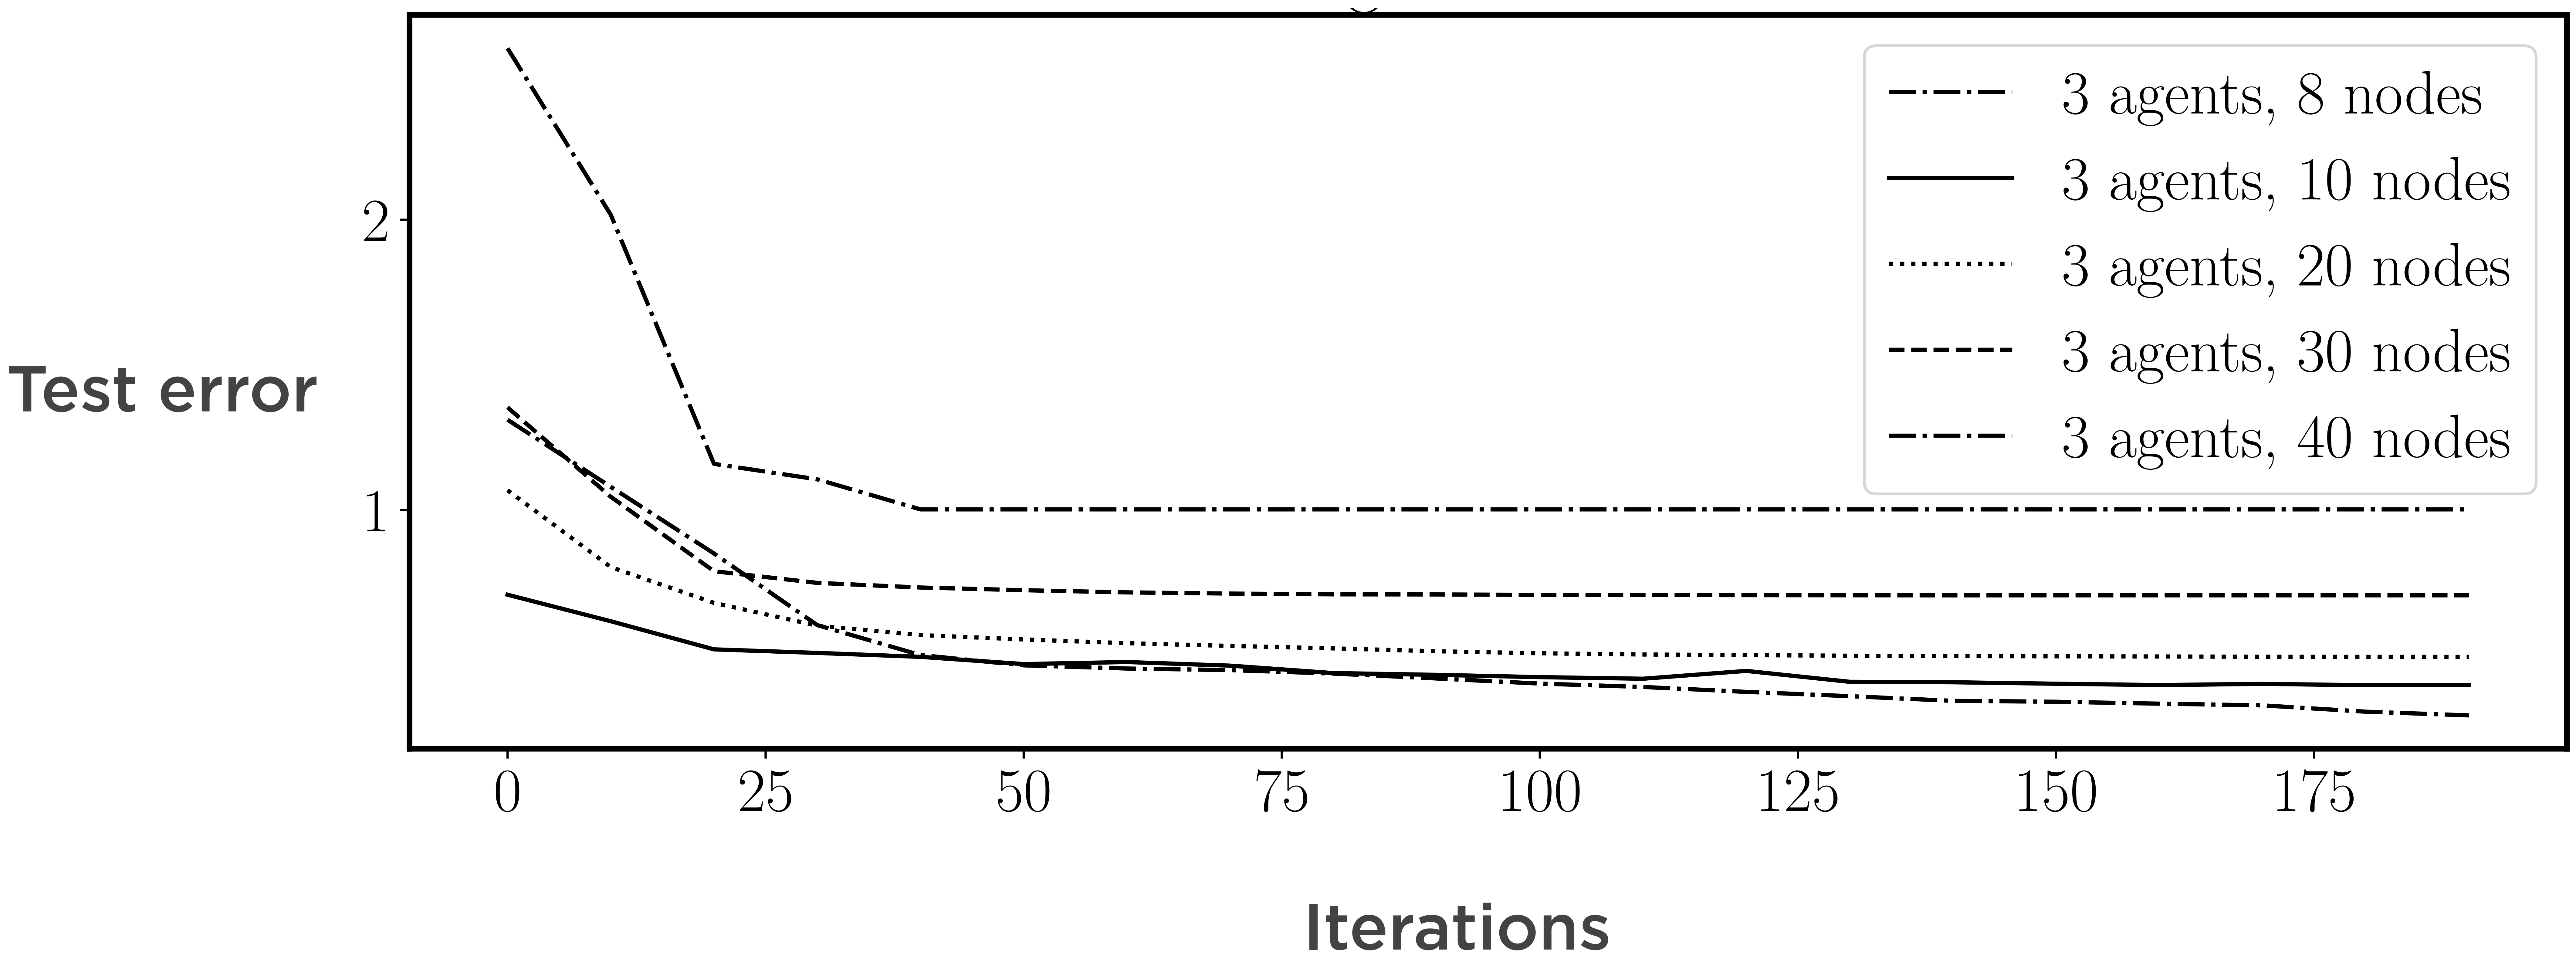

Dataset of *90 equilibria*

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

Cournot Games

There are n players producing a homogeneous good 🍎

cost × quantity

$$u_i(x_i; x_{-i}, \theta, \mu) = -F(x)x_i + c_i x_i$$

price × quantity

$$F(x) = a - b\sum_{j=1}^{n} x_j$$

Unknown **rationality parameters**

Cournot Games

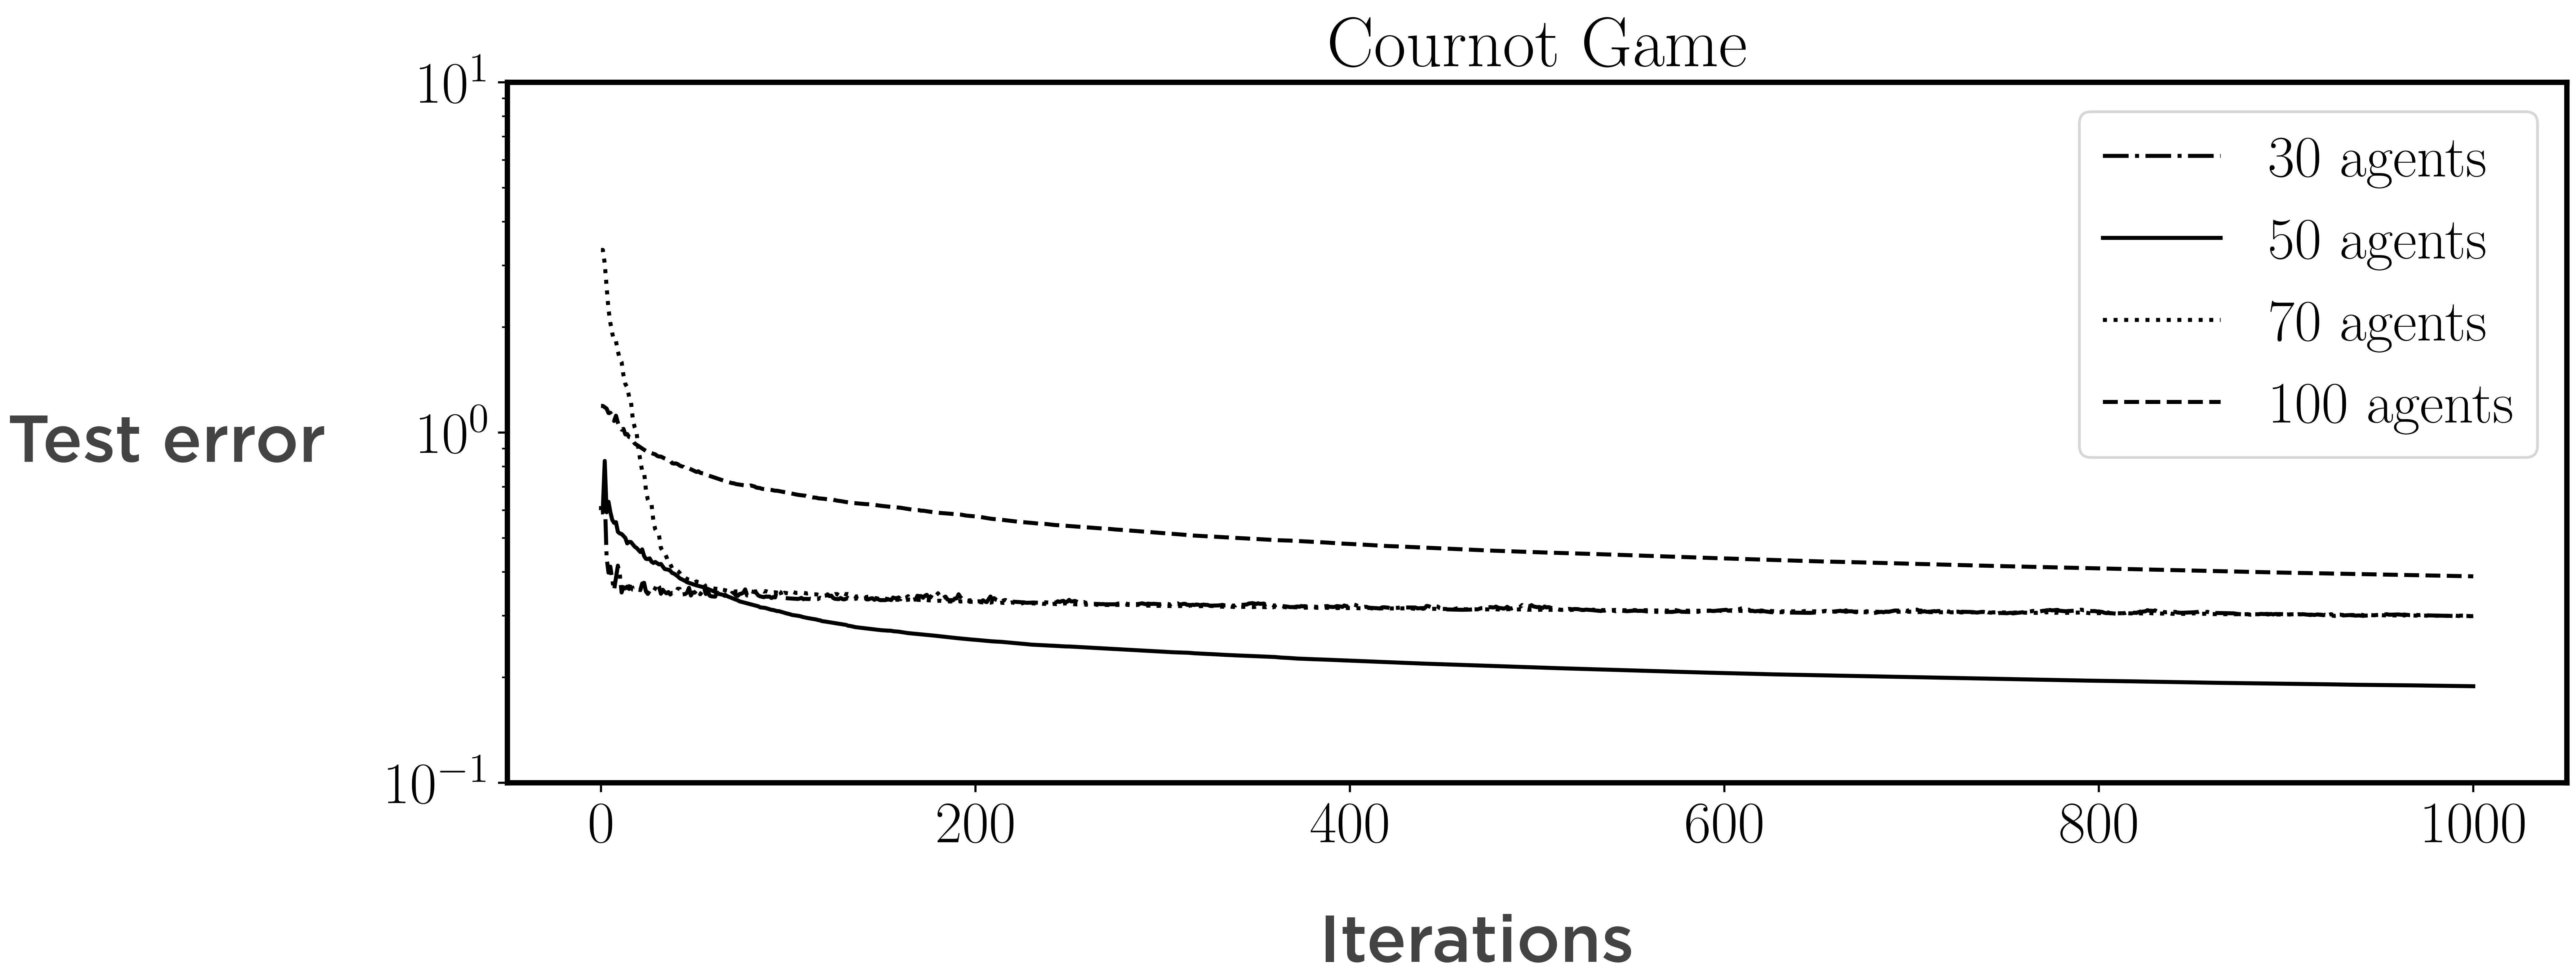

Our algorithm **scales to large datasets**

Cournot Games

Test error

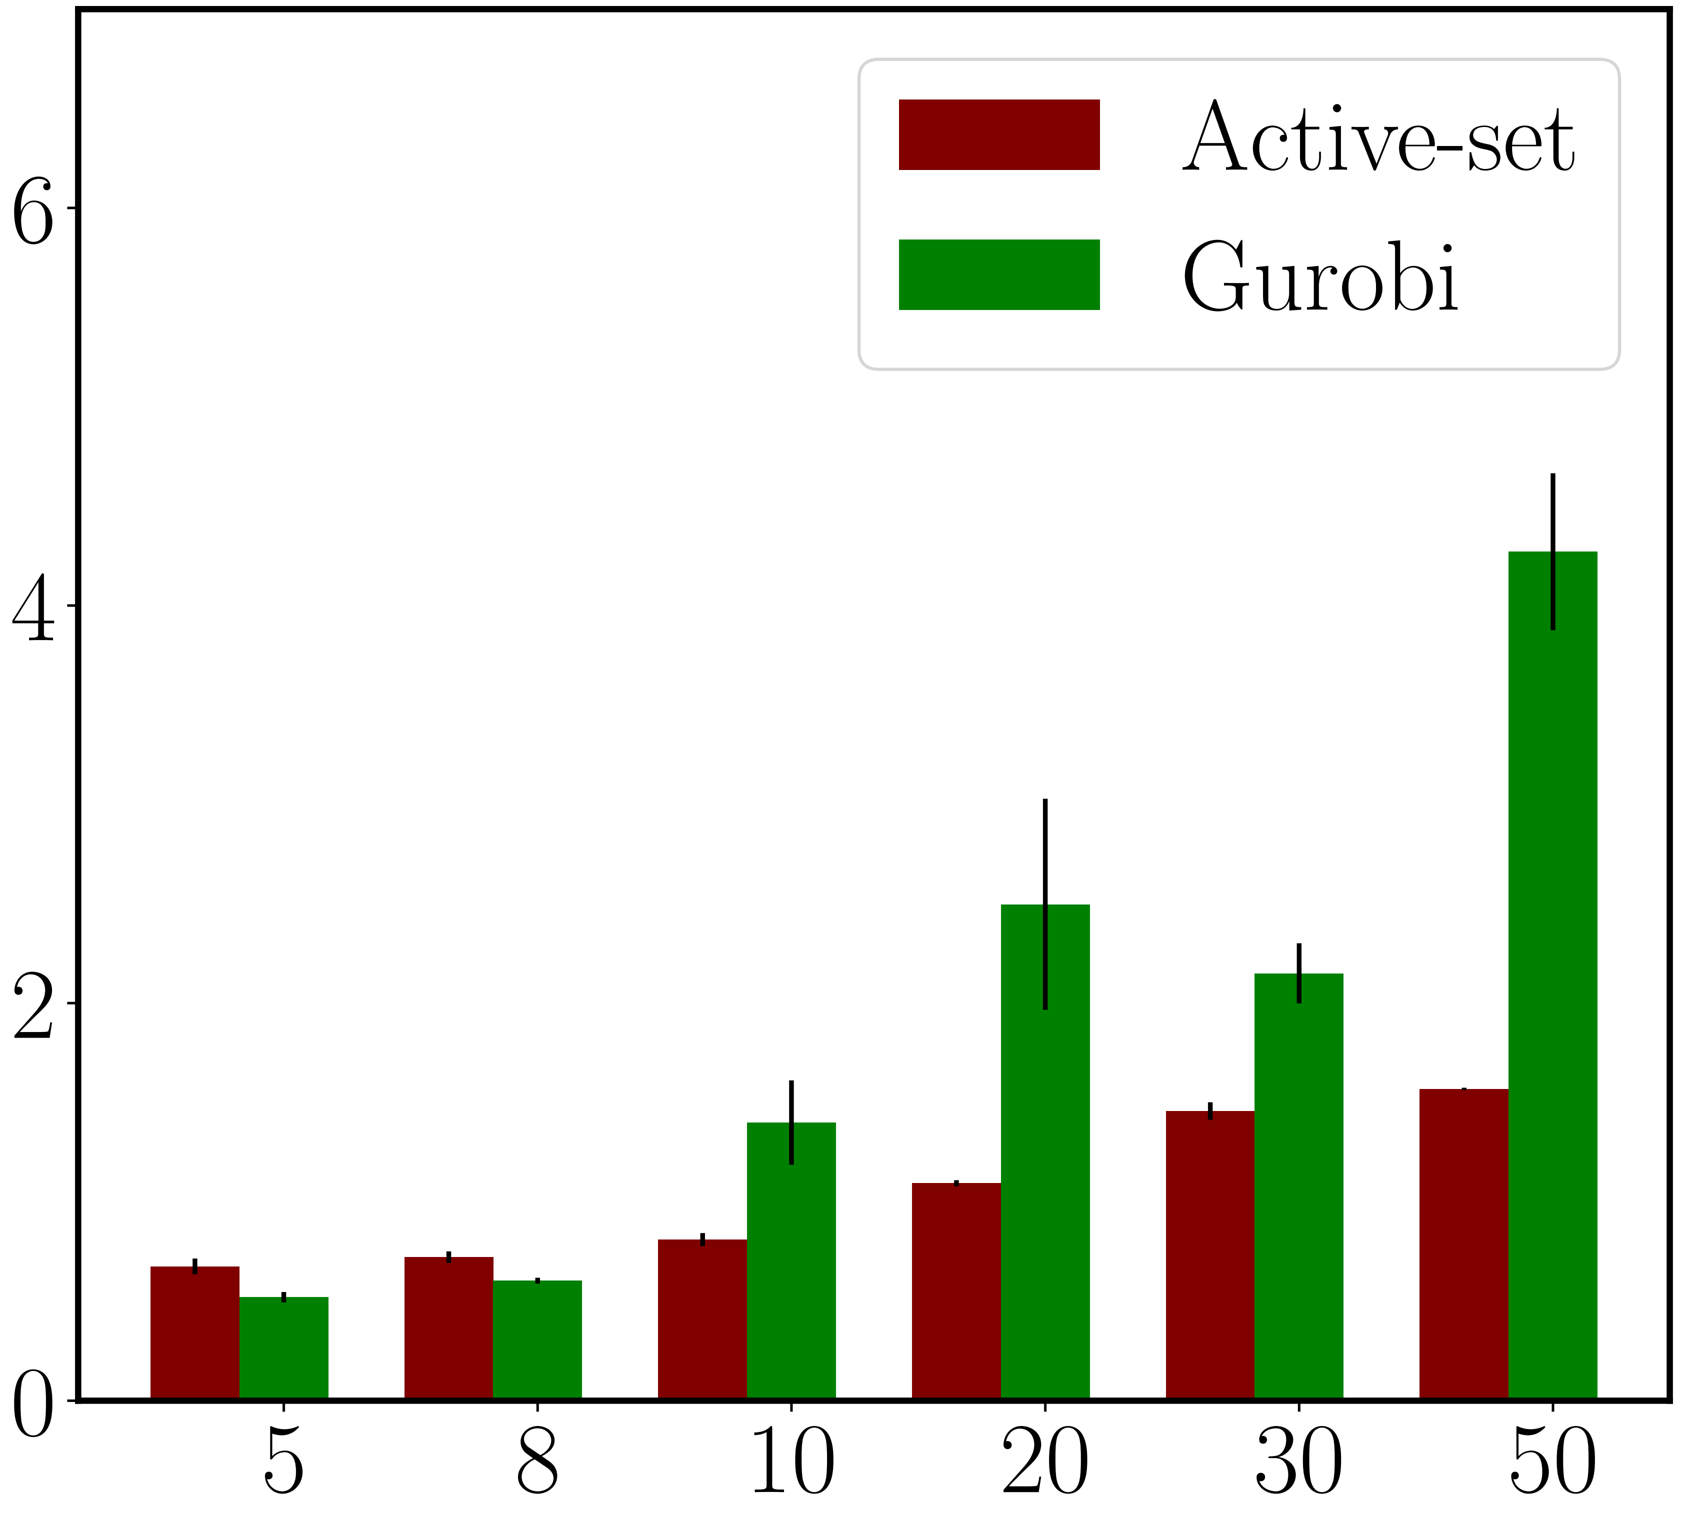

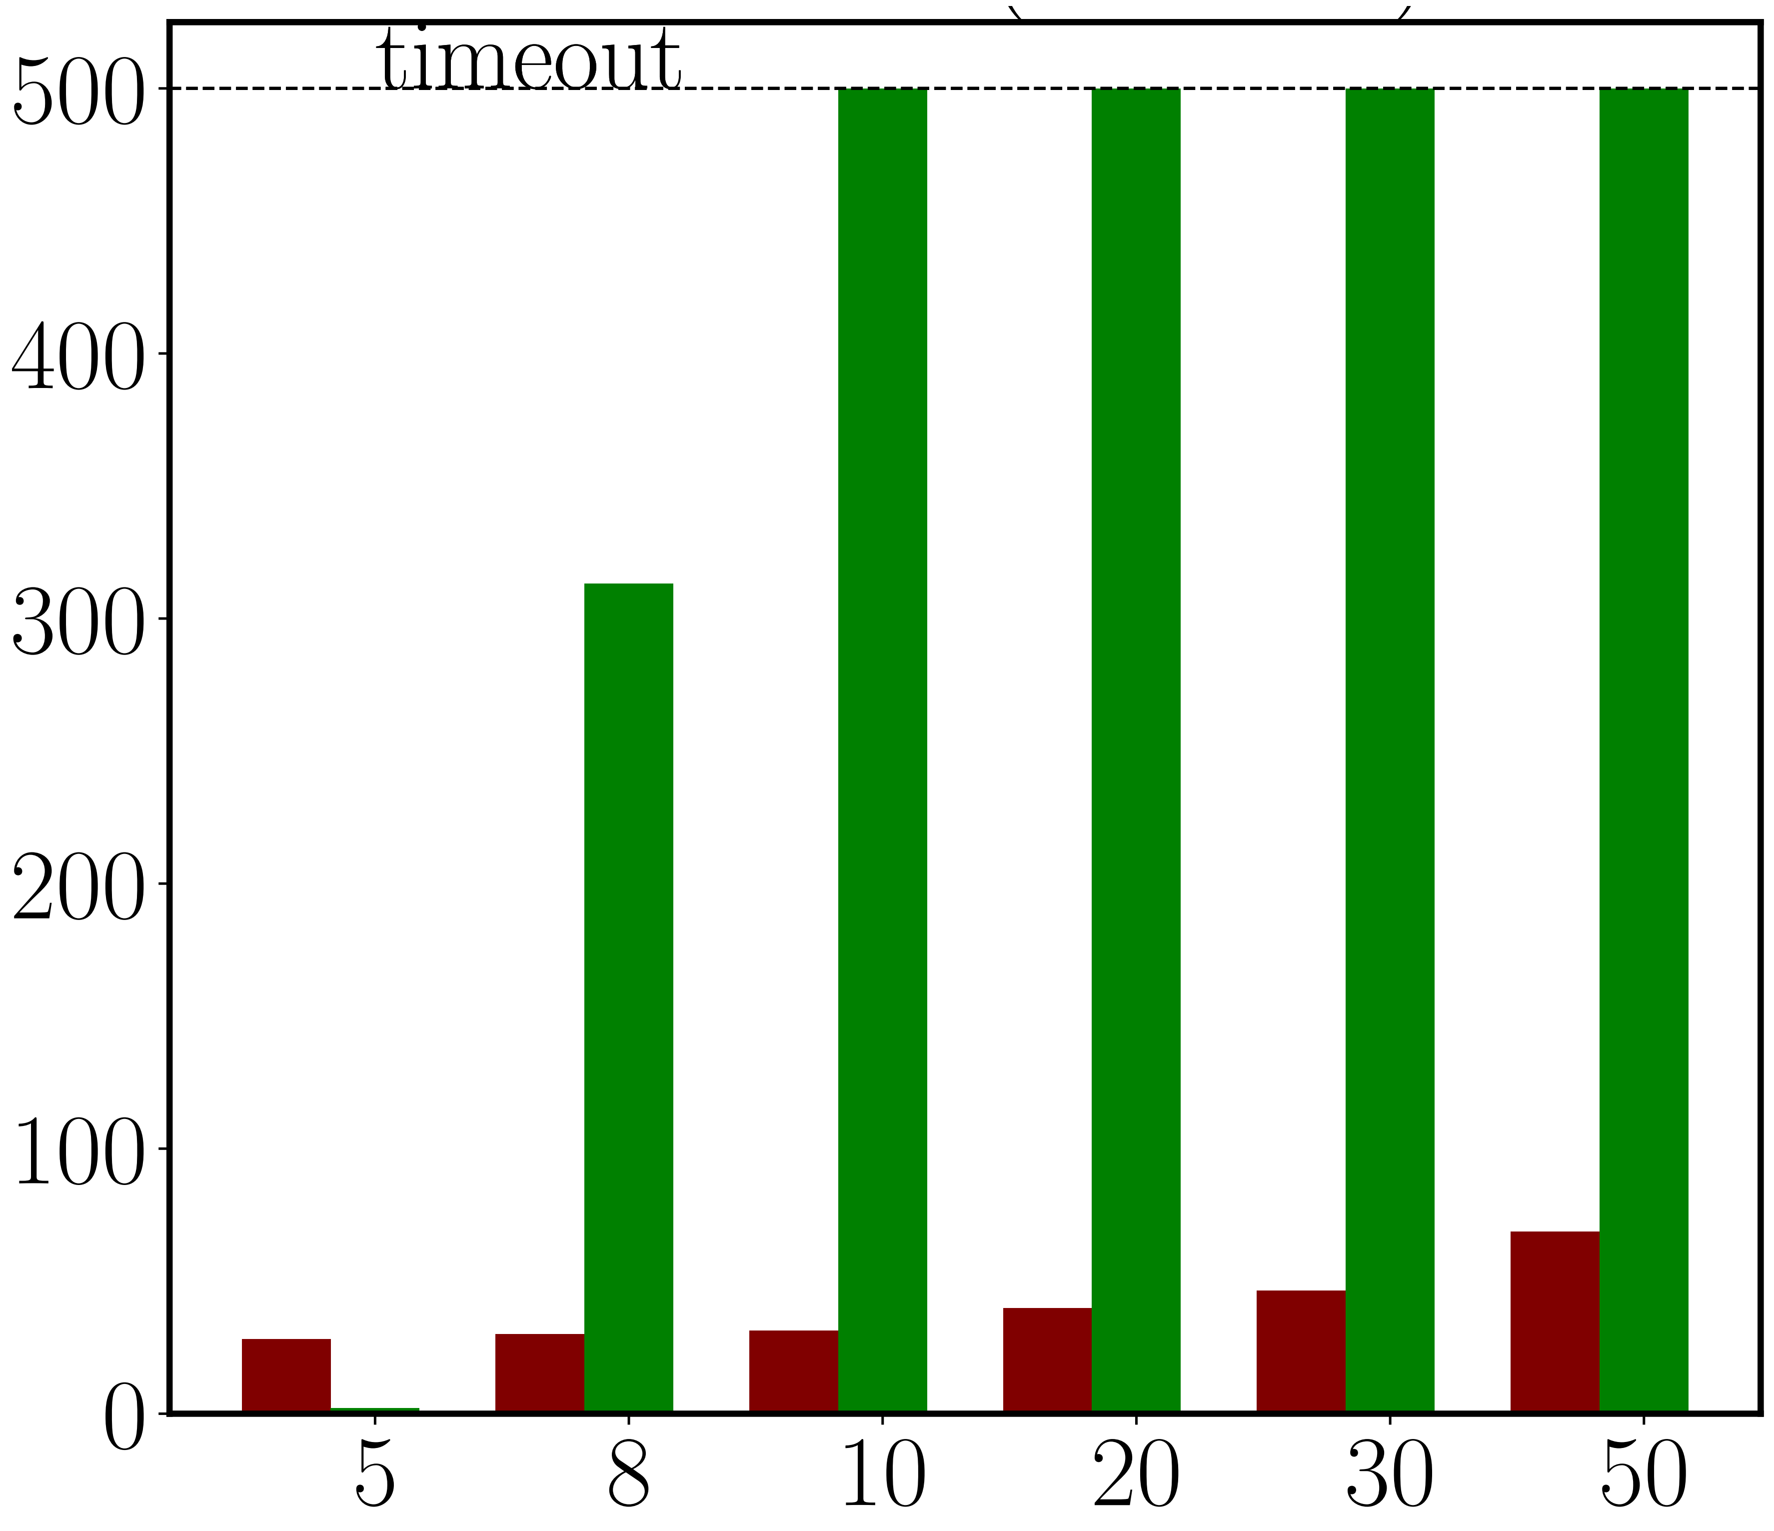

Run time (s)

Number of agents

Our algorithm **scales to large datasets**

Once we know the parameters...

We design market interventions

Decision-making is rarely an individual task

A regulator **observes the outcome** of the interaction but **is uncertain** of the agents' utilities and actions

Learn the Parameters

Intervene

Econometrica, Vol. 70, No. 4 (July, 2002), 1341–1378

THE ECONOMIST AS ENGINEER: GAME THEORY, EXPERIMENTATION, AND COMPUTATION AS TOOLS FOR DESIGN ECONOMICS[1]

BY ALVIN E. ROTH[2]

"Designers therefore **cannot work** only **with the simple conceptual models** used for theoretical insights into the general working of markets. Instead, **market design calls for an engineering approach**.

Experimental and computational economics are natural complements to game theory in the work of design."

The Toolkit: Integer Programming Games

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

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

$$\mathcal{X}_i := \{g_i(x_i) \leq b_i, \quad x_i \in \mathbb{Z}^{\alpha_i} \times \mathbb{R}^{\beta_i}\}$$

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

Integer Programming Games: A Gentle Computational Introduction

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

Decision-making is rarely an individual task

Self-driven interactions with other decision-makers

deciding by solving **complex (e.g., non-convex) optimization problems**

Modeling Capabilities

Informative Solutions

Practically Useful

A few examples

Supply Chain and Transportation

Cronert and Minner, 2021 (OR, TR-B)

Sagratella et al 2020 (EJOR)

Carvalho et al. 2018 (IJ Production Economics)

Simultaneous game among “bilevel” players

*Carvalho, **D.** et al, 2023*

(Management Science)

Cybersecurity

***D.** et al, 2023*

(Ericsson Inc, - Patent pending)

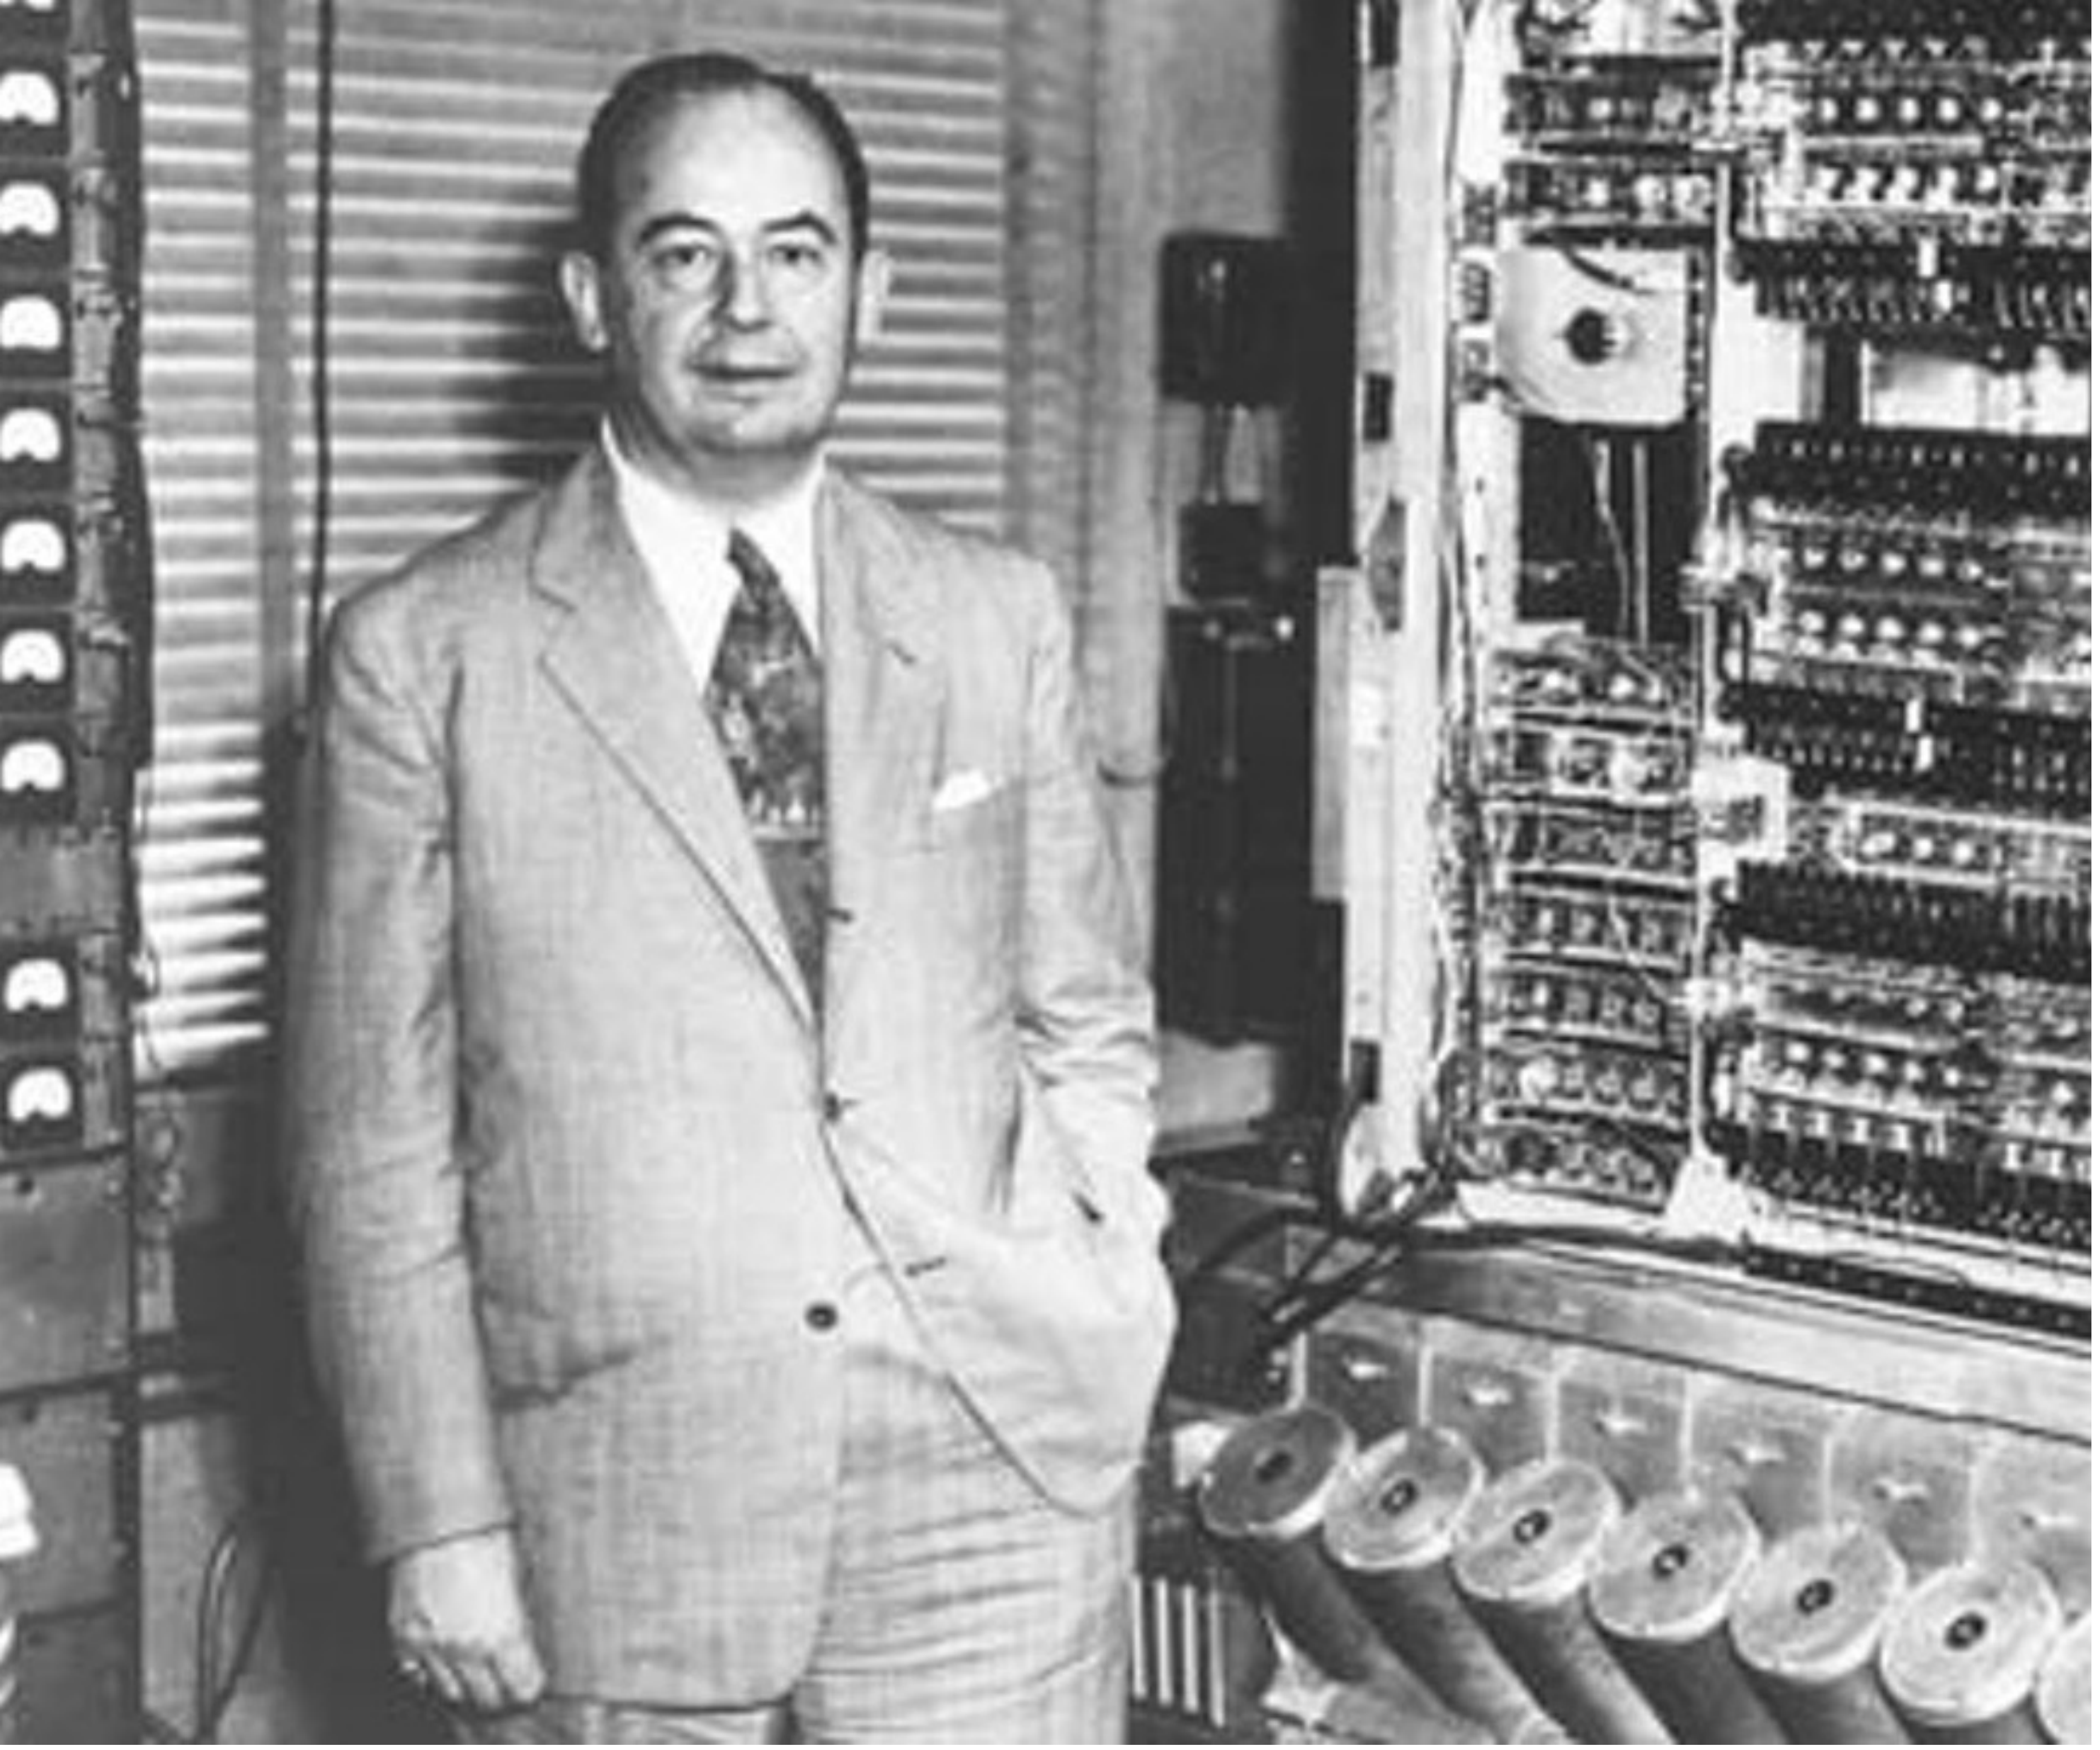

The bad news: non-convexity

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

Historically, **convexity** played a central role in shedding light on the existence and computation of Nash equilibria

State-of-the-art

		Payoff Types	Constraints
Sagratella (2016)	**Branching Method**	Convex payoffs	Bounded Convex Integer
Carvalho et al. (2022)	**Sample Generation Method**	Separable Payoffs	Bounded Mixed-Integer Linear
Schwarze and Stein (2022)	**Branch-and-Prune**	Quadratic Payoffs	Bounded Convex-Integer
Carvalho, D., Lodi, Sankaranarayanan (2021)	**Cut-And-Play**	Separable Payoffs	Polyhedral convex-hull
Cronert and Minner (2021)	**Exhaustive Sample Generation Method**	Separable Payoffs	Bounded Pure-Integer
D. and Scatamacchia (2023)	**Zero Regrets**	Linearizable Payoffs	Bounded Mixed-Integer Linearizable

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** www.dragotto.net

Learning Rationality in Potential Games

arXiv 2303:11188

Integer Programming Games: A Gentle Computational Introduction

INFORMS 2023 TutORial - October 2023

DataX
Accelerating Scientific Discovery at Princeton

UNIVERSITÄT TRIER

PRINCETON UNIVERSITY