Decentralized Noncooperative Games With Coupled Decision-Dependent Distributions

Wenjing Yan   Xuanyu Cao

Nov. 2024

香港科技大學
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Motivations

■ Paradigm of stochastic optimization and machine learning

Make predictions or decisions on new, unseen data

The environment is static

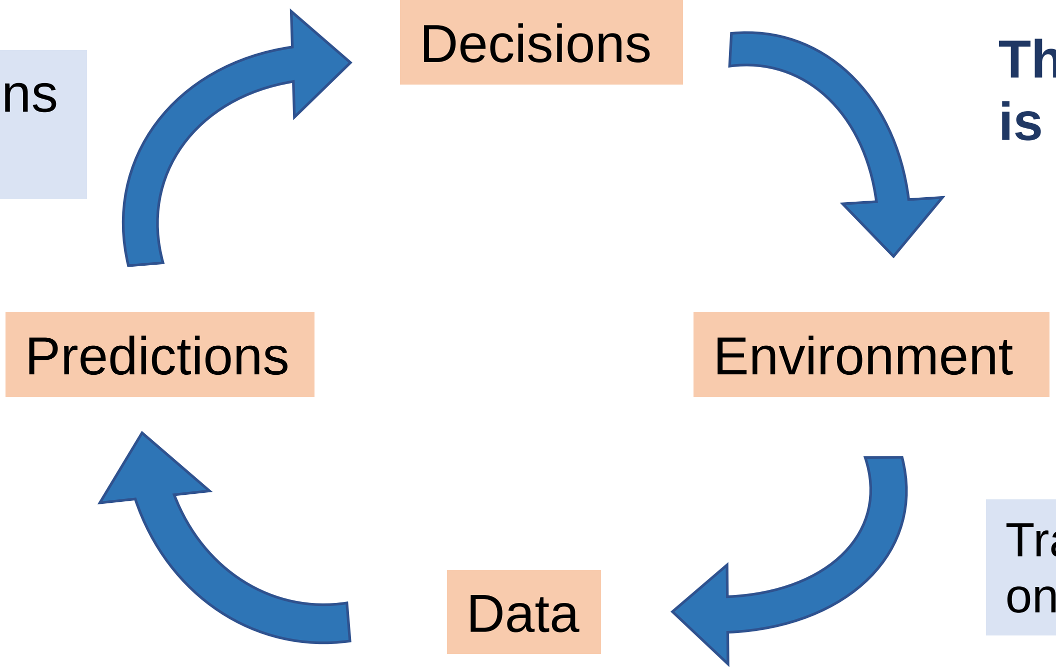

Train models based on given datasets

The training data and the testing data are i.i.d.

Motivations

■ Paradigm of machine learning and stochastic optimization

Make predictions or decisions on new, unseen data

Inaccurate

The training data and the testing data have different distributions

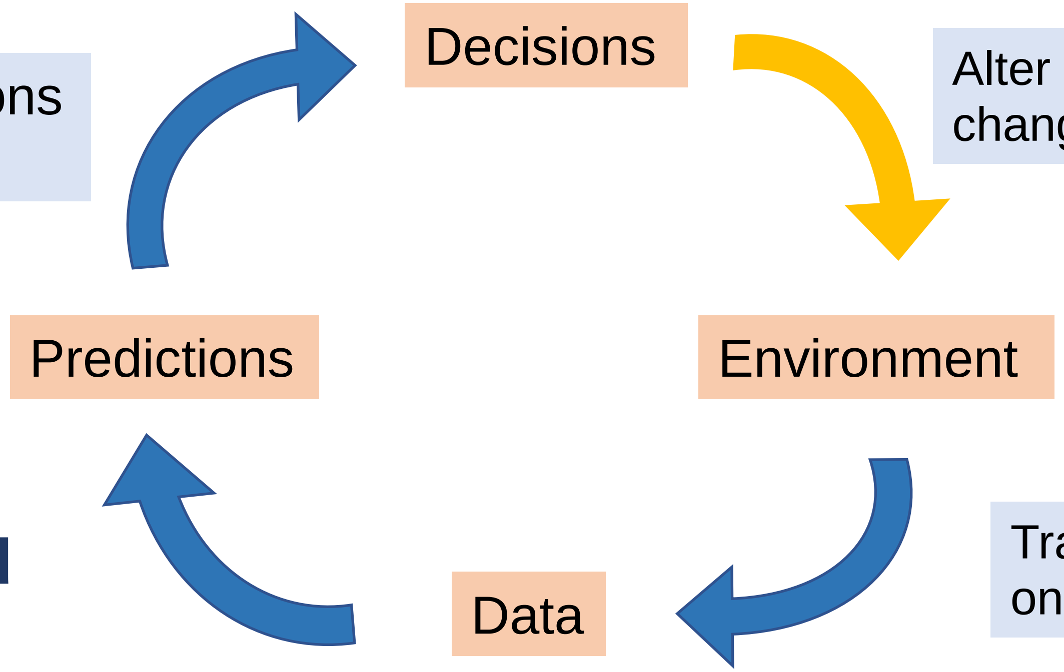

Alter the environment and further change the data distribution.

Train models based on old datasets

Performative Prediction

- **Conventional Learning:**
  - Data $Z \sim \mathcal{D}$
  - Static distribution
  - Goal: minimize risk

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{Z\sim\mathcal{D}}[\ell(\boldsymbol{\theta}; Z)]$$

- **Performative Prediction:**
  - Data $Z \sim \mathcal{D}(\theta)$
  - Decision-dependent distribution
  - Goal: minimize *performative risk*

$$\min_{\boldsymbol{\theta}} \mathrm{PR}(\boldsymbol{\theta}) := \mathbb{E}_{Z\sim\mathcal{D}(\boldsymbol{\theta})}[\ell(\boldsymbol{\theta}; Z)]$$

- Predictions guide decision-making and hence influence future data distributions.
- Initially formalized as *performative prediction* by [Perdomo et al., 2020]
- Represent the strategic responses of data distributions to the taken decisions by a decision-dependent distribution mapping $Z \sim \mathcal{D}(\theta)$.

Problem Formulation

- n-player decentralized noncooperative games with coupled decision-dependent distributions:

$$\begin{aligned} \min_{\boldsymbol{\theta}_i \in \boldsymbol{\Omega}_i} \quad & \mathrm{PR}_i(\boldsymbol{\theta}) := \mathbb{E}_{\boldsymbol{\xi}_i \sim \mathcal{D}_i\left(\boldsymbol{\theta}_i, \boldsymbol{\theta}_{-i}\right)}\left[J_i\left(\boldsymbol{\xi}_i ; \boldsymbol{\theta}_i, \boldsymbol{\theta}_{-i}\right)\right] \\ \text{s.t.} \quad & \boldsymbol{g}_i\left(\boldsymbol{\theta}_i\right) + \sum_{j \neq i} \boldsymbol{g}_j\left(\boldsymbol{\theta}_j\right) \leq \mathbf{0} \qquad (1) \end{aligned}$$

where $\boldsymbol{\theta}_{-i} := \mathrm{col}(\boldsymbol{\theta}_1 \ldots \boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}_{i+1} \ldots \boldsymbol{\theta}_n)$, $\boldsymbol{\theta} := \mathrm{col}(\boldsymbol{\theta}_i , \boldsymbol{\theta}_{-i})$, $\mathcal{D}_i(\boldsymbol{\theta})$ is the data distribution of player i.

- Applications:
  - **Autonomous vehicular networks:** multiple vehicles compete to select routes under constraints such as road capacities, traffic congestion, and travel costs. The route choices of each vehicle influence traffic patterns and consequently affect the travel times experienced by other vehicles.
  - **Networked Cournot games:** traders compete to maximize profits under constraints like market capacities and inventory levels. The trading strategies of these participants impact market volatility and the distribution of asset prices, creating a dynamic pricing landscape.

Main Contributions

- **Problem Formulation:** We formulate the problem of decentralized noncooperative game (1) with data performativity by coupled decision-dependent distributions.
- **Nash Equilibrium:** We examine the Nash equilibrium (NE) for the game (1) and establish sufficient conditions for its existence and uniqueness (E&U).
- **Performative Stable Equilibrium:** We examine the performative stable equilibrium (PSE) for the game (1) and establish sufficient conditions for its E&U.
- **Distance Bound:** We provide the first distance bound between the PSE and NE, which is challenging due to the absence of strong convexity on the joint cost function.
- **Algorithm Design:** We develop a decentralized stochastic primal-dual algorithm for the efficient computing of the PSE and provide rigorous analysis to demonstrate its comparable convergence performance to the case without data performativity.

Nash Equilibrium

■ Theorem 1. (Existence and Uniqueness of NE, informal)

If it is satisfied that

$$\mu - \Sigma_{i=1}^{n} L_i \varepsilon_i \max_{j\in[n]} \sqrt{p_{ij}} - \sqrt{\Sigma_{i=1}^{n} L_i^2 \varepsilon_i^2 p_{ii}} > 0, \qquad (2)$$

then, the PP game (1) is strongly monotone and admits a unique NE.

- μ is the monotone parameter of the fixed distribution counterpart of the game (1).
- L_i is the smoothness parameter of $J_i(.)$.
- ε_i is the sensitivity parameter of $\mathcal{D}_i$ satisfying $\mathcal{W}_1(\mathcal{D}_i(\boldsymbol{\theta}) - \mathcal{D}_i(\boldsymbol{\theta}')) \leq \varepsilon_i \sqrt{\Sigma_{j=1}^{n} p_{ij} \left\|\boldsymbol{\theta}_j - \boldsymbol{\theta}_j'\right\|_2}$.
- p_{ij} is the normalized influence of play j`s decision on $\mathcal{D}_i(.)$.

➢ If L_1 =… = L_n = L, ε_1 =… = ε_n = ε, and p_{ij} =$1/n$ for all i, j. the above condition is simplified to μ - 2$\varepsilon L > 0$, i.e., $\mu >$ 2εL, which coincides with the condition for the E&U of the performative optimal (PO) point in the single agent PP problem (Miller et al., 2021).

➢ Due to the presence of data performativity, i.e., $\varepsilon_i > 0$, the PP game (1) requires a more stringent condition for the E&U of NE.

Performative Stable Equilibrium

■ Definition of Performatively Stable Equilibrium (PSE):

The strategy profile $\boldsymbol{\theta}^{\mathrm{pse}} := \mathrm{col}(\boldsymbol{\theta}_1^{\mathrm{pse}} \ldots \boldsymbol{\theta}_n^{\mathrm{pse}})$ is a PSE point of the game (1) if it holds for all $i \in [n]$ that

$$\boldsymbol{\theta}_i^{\mathrm{pse}} \in \underset{\boldsymbol{\theta}_i \in \boldsymbol{\Omega}_i}{\arg\min}\ \mathbb{E}_{\boldsymbol{\xi}_i \sim \mathcal{D}_i\left(\boldsymbol{\theta}^{\mathrm{pse}}\right)}\left[J_i(\boldsymbol{\xi}_i; \boldsymbol{\theta}_i, \boldsymbol{\theta}_{-i}^{\mathrm{pse}})\right] \ \text{s.t.}\ \boldsymbol{g}_i(\boldsymbol{\theta}_i) + \sum_{j \neq i} \boldsymbol{g}_j(\boldsymbol{\theta}_j^{\mathrm{pse}}) \leq \mathbf{0}.$$

- $\boldsymbol{\theta}^{\mathrm{pse}}$ achieves the NE of the game (1) under fixed data distribution $\{\mathcal{D}_i(\boldsymbol{\theta}_i^{\mathrm{pse}})\}_{i \in [n]}$.

■ Theorem 2. (Existence and Uniqueness of PSE, informal)

If it is satisfied that $\frac{1}{\mu}\sqrt{\sum_{i=1}^{n} L_i^2 \varepsilon_i^2 \max_{j \in [n]} p_{ij}} < 1, \quad (3)$ the PP game (1) admits a unique PSE, which can be found by repeatedly minimizing the game (1) under fixed data distribution induced by current decisions.

- If $L_1 = \ldots = L_n = L$, $\varepsilon_1 = \ldots = \varepsilon_n = \varepsilon$, and $p_{ij} = 1/n$ for all i, j. the above condition is simplified to $\mu > \varepsilon L$, which coincides with the condition for the E&U of performative stable (PS) point in the single agent PP problem (Perdomo et al., 2020).
- The E&U condition of the PSE is weaker than that of the NE ($\mu > \varepsilon L$ V.S. $\mu > 2\varepsilon L$).

Distance Between PSE and NE

■ **Theorem 3. (Distance Between PSE and NE, informal)**

Define $\tilde{\mu} := \mu - \Sigma_{i=1}^{n} L_i \varepsilon_i \max_{j\in[n]} \sqrt{p_{ij}}$ and $\alpha := \Sigma_{i=1}^{n} G_i \left(1+\varepsilon_i\right) \max_{j\in[n]} \sqrt{p_{ij}}$

Suppose that $\tilde{\mu} > 0$ holds. Then, for every PSE point and NE point, we have the following relations:

$$\left\|\theta^{\mathrm{pse}} - \theta^{\mathrm{ne}}\right\|_2 \leq \frac{1}{\tilde{\mu}} \sqrt{\Sigma_{i=1}^{n} G_i^2 \varepsilon_i^2} \quad \text{and} \quad \left\|\mathrm{PR}(\theta^{\mathrm{pse}}) - \mathrm{PR}(\theta^{\mathrm{ne}})\right\|_2 \leq \frac{\alpha}{\tilde{\mu}} \sqrt{\Sigma_{i=1}^{n} G_i^2 \varepsilon_i^2}$$

where G_i is the Lipschitz parameter of $J_i(.)$ for all $i\in[n]$.

- Larger performative strengths $\{\varepsilon_1\}_{i\in[n]}$ widen the gap, while a bigger monotonicity parameter μ reduces this gap.
- Comparable to the result in single agent PP case that $\|\boldsymbol{\theta}^{\mathrm{PO}} - \boldsymbol{\theta}^{\mathrm{PS}}\|_2 < 2L\varepsilon/\mu$.

Computation of PSE

Algorithm 1 Decentralized Stochastic Primal-Dual Algorithm: The Procedures at Player i, $\forall i \in [n]$:

Initialize $\boldsymbol{\theta}_i^1 \in \boldsymbol{\Xi}_i$ arbitrarily. Set $\boldsymbol{\lambda}_i^1 = \mathbf{0}$ and $\widehat{\boldsymbol{\theta}}_{ih}^1 = \mathbf{0}$ for all $h \neq i$.
for $t = 1$ to T **do**
    Exchange $\boldsymbol{\theta}_i^t$, $\widehat{\boldsymbol{\theta}}_i^t$, and $\boldsymbol{\lambda}_i^t$ with all neighbors;
    Update the estimate $\widehat{\boldsymbol{\theta}}_{ih}^t$ for all $h \neq i$ by: $\widehat{\boldsymbol{\theta}}_{ih}^{t+1} = \sum_{k \neq h} a_{ik} \widehat{\boldsymbol{\theta}}_{kh}^t + a_{ih} \boldsymbol{\theta}_h^t$;
    Deploy the model $\boldsymbol{\theta}_i^t$ and sample $\boldsymbol{\xi}_i^t \sim \mathcal{D}_i(\boldsymbol{\theta}_i^t, \boldsymbol{\theta}_{-i}^t)$;
    Update the primal variable by: $\boldsymbol{\theta}_i^{t+1} = P_{\boldsymbol{\Omega}_i}\left[\boldsymbol{\theta}_i^t - \gamma_t \left(\nabla_{\boldsymbol{\theta}_i} J_i\left(\boldsymbol{\xi}_i^t; \boldsymbol{\theta}_i^t, \widehat{\boldsymbol{\theta}}_i^t\right) + \gamma_t \nabla \boldsymbol{g}_i(\boldsymbol{\theta}_i^t)^\top \boldsymbol{\lambda}_i^t\right)\right]$;
    Update the dual variable by: $\boldsymbol{\lambda}_i^{t+1} = \left[\left(1 - \gamma_t^2\right) \sum_{j \in \mathcal{N}_i} a_{ij} \boldsymbol{\lambda}_j^t + \gamma_t \boldsymbol{g}_i\left(\boldsymbol{\theta}_i^t\right)\right]_+$.
end for

Convergence Analysis

■ Theorem 4. (Convergence of Algorithm 1, informal)

Define $\tilde{\mu} := \mu - \Sigma_{i=1}^{n} L_i \varepsilon_i \max_{j\in[n]} \sqrt{p_{ij}}$. Under standard assumptions, by running Algorithm 1 for T times of iterations, both the performative regret and the constraint violations of the game (1) is upper bounded by $\mathcal{O}(T^{3/4})$.

- The performative effect slows down the convergence speed through $\tilde{\mu}$.
- The performance of Algorithm 1 matches the convergence order of the case without data performativity (Lu et al., 2020).

Simulations on A Networked Cournot Game

- n **firms selling a single commodity across** m **markets.**

$$\min_{\boldsymbol{\theta}_i\in\boldsymbol{\Omega}_i} \ \mathbb{E}_{p_j\sim\mathcal{D}_i(\boldsymbol{\theta})}\left[\boldsymbol{d}_i^{\mathrm{T}}\boldsymbol{\theta}_i-\sum_{j=1}^{m} c_j\theta_{ij}\right]$$
$$\text{s.t.} \quad \theta_{ij}+\sum_{i'\neq i}\theta_{i'j}\leq B_j.$$

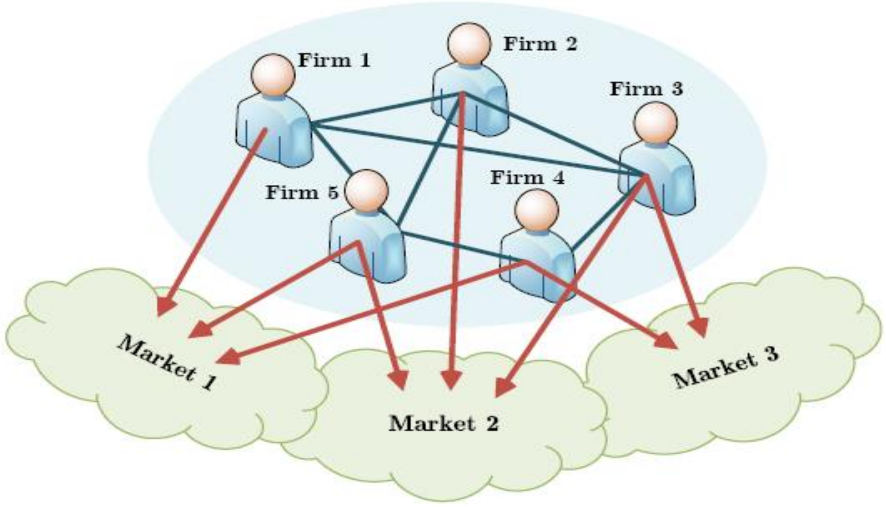

- θ_{ij} : the product's quantity that player i selling to the j-th market. $\boldsymbol{\Theta}_i=\mathrm{col}(\theta_{i1},\ldots,\theta_{im})$.
- c_j : unit demand price of the market j.

  $c_j:=\xi_j+\Lambda_j\,(\sum_{j=1}^{n}\theta_{ij})^{1/\tau_j}$, $\quad \xi_j:=\xi_j^0+\varepsilon\frac{\beta_j}{\sum_{j'=1}^{n}\beta_{j'}}(\sum_{j=1}^{n}\theta_{ij})$, ξ_j^0 : random base component,

  $\varepsilon>0$: performative strength, β_j : relative performative strength of market j.
- B_j : accommodating capacity of market j.

Simulations on A Networked Cournot Game

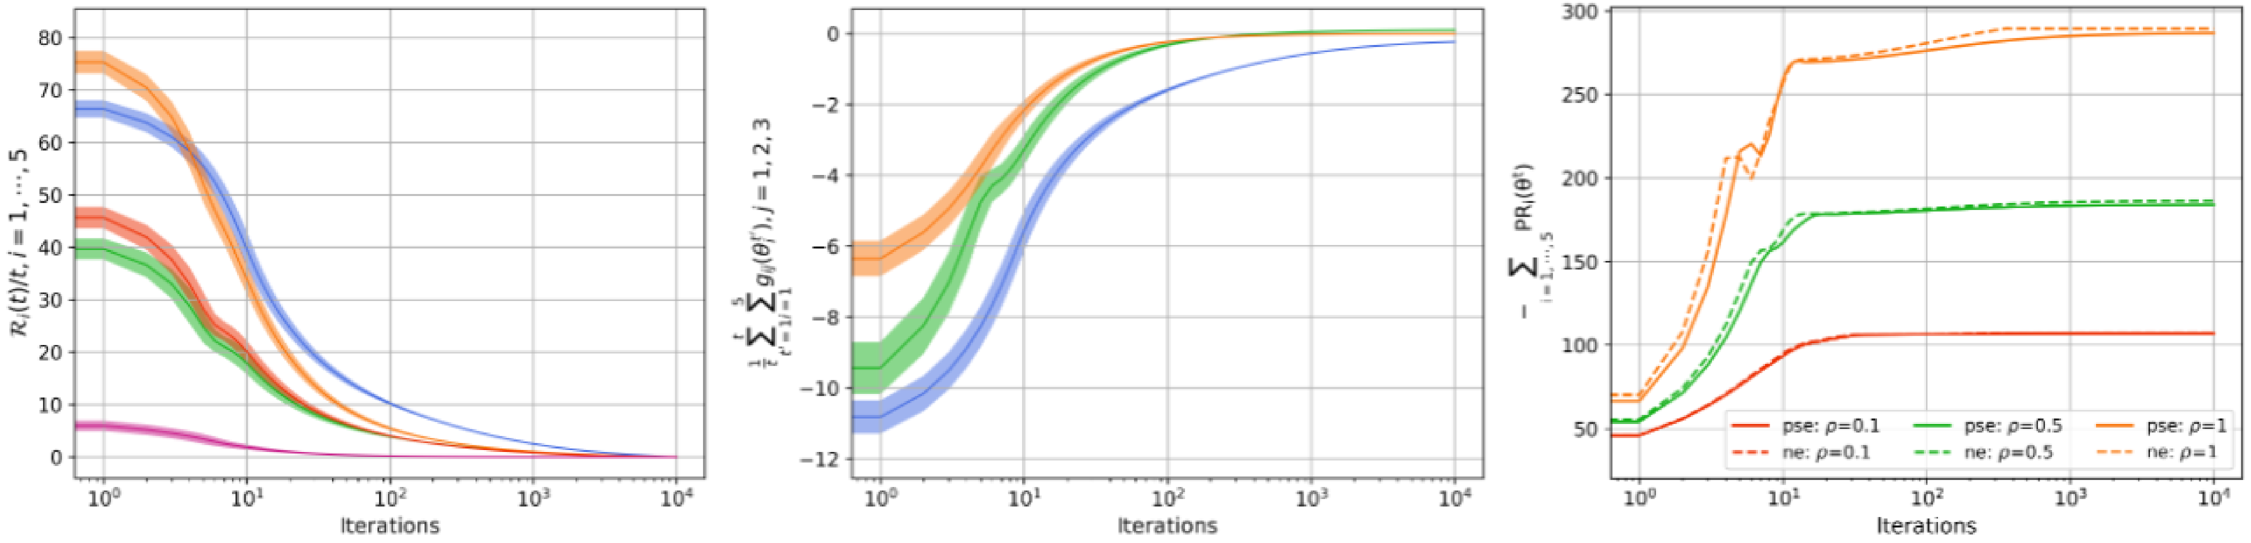

Figure 1: (a) $\frac{\mathrm{Reg}(t)}{t\cdot\mathrm{Reg}(1)}$; (b) $\frac{\mathrm{Vio}(t)}{t\cdot|\mathrm{Vio}(1)|}$; (c) $|\mathrm{PR}(\boldsymbol{\theta}^{\mathrm{pse}}) - \mathrm{PR}(\boldsymbol{\theta}^{\mathrm{ne}})|$.

- **Performative regret:** $\mathcal{R}_i(T) := \sum_{t=1}^{T}\left(\mathbb{E}_{\xi_i\sim\mathcal{D}_i(\boldsymbol{\theta}^{\mathrm{pse}})}[J_i(\boldsymbol{\xi}_i;\boldsymbol{\theta}_i^t,\boldsymbol{\theta}_{-i}^{\mathrm{pse}})] - \mathrm{PR}_i(\boldsymbol{\theta}^{\mathrm{pse}})\right)$.
- **Constraint violation:** $\mathcal{R}_g(T) := \left\|\left[\sum_{t=1}^{T}\sum_{i=1}^{n}\boldsymbol{g}_i(\boldsymbol{\theta}_i^t)\right]_+\right\|_2$.

➢ Both $\mathcal{R}_i(\mathrm{t})/T$ and $\mathcal{R}_g(\mathrm{t})/T$ converge sublinearly to 0 as the iterations increase, verifying the effectiveness of Algorithm 1 for handling data performativity.

➢ The total revenue -$\sum_{i=1}^{n}\mathcal{R}_i(\mathrm{t})$ of all firms at the PSE closely approaches that of the NE, verifying the effectiveness of the PSE solution.

Thank you for listening!