MULTIPLAYER PERFORMATIVE PREDICTION GAMES WITH INCOMPLETE INFORMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper investigates multiplayer performative prediction games in settings with incomplete information. We formalize these games as *performative Bayesian games*, in which players strategically choose their decisions (or predictions) that jointly shape the resulting data distributions, while each player possesses private information unknown to others. A natural adaptive strategy for players is to iteratively (1) optimize their decisions based on current posterior beliefs about others' private information under the data distribution induced by the joint decision, and (2) update those beliefs upon observing others' realized decisions. We introduce the concept of a *performatively stable Bayesian equilibrium* as a desirable outcome: a decision profile and associated belief system such that the deployed decisions are mutually optimal given the data distributions they induce and the converged beliefs, eliminating the need for further update of either decisions or beliefs. Our main contributions are threefold. First, we provide a formal game-theoretic definition of performative Bayesian games. Second, we derive theoretical conditions under which the natural alternating best-response-and-belief-update dynamics converge to a performatively stable Bayesian equilibrium. Third, leveraging these theoretical results, we design an algorithm that guarantees convergence to such an equilibrium and minimizes the maximum deviation from the original equilibrium when the convergence conditions are violated. Experimental evaluations on crude oil trading and ride-sharing market validate our theoretical results and demonstrate the effectiveness of the proposed algorithm.

1 INTRODUCTION

Conventional supervised learning relies on the assumption that training and testing data distributions are identical. However, in real-world scenarios, a decision (or a model prediction) may trigger data distribution shifts, rendering the decision suboptimal on the shifted data distribution. For instance, a firm raising product prices to maximize revenue may induce a demand drop, thereby reducing its actual revenue and necessitating an update of its price. This phenomenon is termed performative prediction (Perdomo et al., 2020), which has been explored in a range of decision-dependent applications, including marketing (König et al., 2025), recommendation systems (Xiang et al., 2024), federated learning (Jin et al., 2024), and reinforcement learning (Pollatos et al., 2025).

In reality, the performative effect is rarely caused by a single decision, but by decisions of multiple players. In multiplayer performative prediction games (Narang et al., 2023; Wang et al., 2023; Zhong et al., 2026a), players optimize their decisions based on the data distributions influenced by the joint decisions. The relationship between joint decisions and induced data distribution relevant to player $i \in [n]$ is characterized by a distribution map, denoted as $\mathcal{D}_i(X)$, where $X = (x_1, x_2, \cdots, x_n)^T$ is the joint decision of all players, and x_i is the decision of player i. Each player must play the best response to others with respect to the decision-dependent data distribution, resulting in a more complex dynamic than single-player scenarios. For example, demand for a product may decrease due to a price increase for the product itself or a price decrease for competing products.

Existing works on multiplayer performative prediction games assume that all players possess complete information. However, in real-world applications (e.g., commodity trading and marketing), players make decisions based on private information, denoted as θ_i for player i, which is unobservable to others. With others' private information θ_j ($j \neq i$) being unknown, each player i must make

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

decisions based on not only the decision-dependent distribution $\mathcal{D}_i(X)$, but also a belief (probability distribution) μ about θ_j, denoted as $\mu(\theta_j)$. In this paper, we term such a multiplayer performative prediction game with incomplete information a *performative Bayesian game*, where players simultaneously deal with two challenges: distribution shift and incomplete information. Here, we use a performative Bertrand competition with incomplete information as an example of such a game:

Example 1.1. (Bertrand Competition (Vintila & Roman, 2021)) Suppose there are n firms, each setting a price $x_i \in \mathbb{R}$ to maximize revenue with demand $q_i = z_i + a_i x_i + \mathbf{a}_{-i} x_{-i}$, where $z_i \sim \mathcal{D}_i(X)$ is a demand parameter influenced by the joint decision X, $a_i \in \mathbb{R}$ and $\mathbf{a}_{-i} \in \mathbb{R}^{n-1}$ are publicly known parameters, and $x_{-i} \in \mathbb{R}^{n-1}$ is the price vector of competitors. The cost $\theta_i \in \mathbb{R}$ of each firm $i \in [n]$ is its private information, unknown to other firms j. Each firm i holds a belief $\mu(\theta_{-i})$ about the costs of its competitors. In this case, each firm $i \in [n]$ must anticipate the competitors' price x_j according to its belief $\mu_i(\theta_j)$, and maximize its revenue by $\min_{x_i} -(x_i - \theta_i)q_i$. For the parameter $z_i \sim \mathcal{D}_i(X)$, a possible data distribution map can be modeled as a logistic function (Phillips, 2021):

$$z_i \sim \mathcal{D}_i(X) = \frac{2z_{0,i}}{1 + e^{-(b_i x_i + \mathbf{b}_{-i} x_{-i})}} + w_i, \tag{1}$$

where $z_{0,i} \in \mathbb{R}$ is the initial demand parameter of firm i, w_i is a noise term, and $b_i \in \mathbb{R}$ and $\mathbf{b}_{-i} \in \mathbb{R}^{n-1}$ are unknown parameters controlling the shift induced by x_i and x_{-i}, respectively.

In such a performative Bayesian game, given the unknown $\mathcal{D}_i(X)$ and the observable decisions of others, all players alternate between two steps: (1) optimizing their decisions based on beliefs under data distributions induced by the joint decision, and (2) updating their beliefs based on the observation of others' decisions. We term this procedure a *repeated performative Bayesian game*. A desirable status is that deployed decisions are optimal for each player with respect to induced data distributions under converged beliefs, eliminating the need for further update of decisions or beliefs. In this paper, we call this pair of decisions and beliefs a *performatively stable Bayesian equilibrium*.

1.1 Our Contributions

This paper proposes a novel framework for performative Bayesian games, including formalization, theoretical analysis, and algorithm design. Our contributions are threefold:

- **Formalization of the Game-Theoretic Definition:** In Section 2, we formalize the performative Bayesian game. We then describe the procedure of the repeated performative Bayesian game, where all players optimize decisions based on beliefs under induced data distributions and update their beliefs based on observed decisions. Finally, we define the performatively stable Bayesian equilibrium, which is a Bayesian equilibrium with respect to induced data distributions under converged beliefs.
- **Theoretical Convergence Analysis:** To characterize the performative effect, we apply a gradient-based sensitivity measure for the repeated performative Bayesian game in Section 3.1, which quantifies the difference in loss function gradients under distinct decision-dependent distributions. Using this measure, we derive conditions under which a repeated performative Bayesian game converges to a performatively stable Bayesian equilibrium in Section 3.2 by investigating the dynamics of beliefs and decisions. We then extend these results to practical scenarios where only finite samples are available in Section 3.3.
- **Algorithm Design:** Leveraging our theoretical results, we design an algorithm for a repeated performative Bayesian game in Section 4. This algorithm guarantees convergence to a performatively stable Bayesian equilibrium by adapting to data distribution maps. When the original game fails to satisfy the convergence conditions, we show that our algorithm minimizes the maximum deviation of the equilibrium from the original equilibrium.

We perform numerical experiments using the Cournot competition in the crude oil trade and the Bertrand competition in the ride-share market in Section 5. Experimental results validate our theoretical results and demonstrate the effectiveness of our algorithm under various distribution maps.

1.2 Related Work

Performative Prediction. Existing research on single-agent performative prediction assumes that a data distribution shift is induced by a single decision (or a model), and focuses on two objec-

tives: performative stability and performative optimality. Performative stability minimizes loss under the induced data distribution, preventing further distribution shifts through repeated risk minimization (Perdomo et al., 2020; Bautiste et al., 2026; Zhong et al., 2026b) or repeated gradient descent (Mendler-Dünner et al., 2020; Piliouras & Yu, 2023; Li et al., 2025a). Performative optimality minimizes the performative loss, which integrates the data distribution map into the loss function, using convex optimization (Tsoy et al., 2025; Bracale et al., 2025) or gradient descent (Chen et al., 2024; Li et al., 2025b).

In real-world scenarios, performativity is widely observed in multiplayer settings, where multiple decisions shape the data distributions. In a competitive setting, each player minimizes individual loss via repeated retraining (Narang et al., 2023; Zhong et al., 2026a) or gradient descent (Cutler et al., 2024; Yan & Cao, 2024). In a cooperative setting, players jointly optimize a global loss through gradient descent (Wang et al., 2023; Li et al., 2022). However, existing studies rely on the complete information assumption, where each player has full knowledge of others. In reality, competitive scenarios (e.g., commodity trading and marketing) or even cooperative scenarios (e.g., federated learning (Jin et al., 2024)) cannot satisfy this assumption, since players do not share their private details. To enhance the practical applicability of multiplayer performative prediction games, we investigate the competitive setting in the presence of incomplete information, where each player possesses private information and must make a decision without full knowledge of competitors.

Bayesian Game. The incomplete information game is first formalized as the Bayesian game by Harsanyi (1967), where each player makes decisions based on beliefs and updates beliefs after observing others' decisions. Much research focuses on deriving equilibrium solutions to various Bayesian games (Huang et al., 2025; Li et al., 2024a). Recently, the Bayesian game framework has been applied to various machine learning domains, including reinforcement learning (Li et al., 2024b), federated learning (Hu et al., 2024), and large language model reasoning (Yi et al., 2025). However, existing Bayesian game research assumes fixed data distributions, which neglects the performative effect in real-world applications. In this paper, we investigate and analyze the Bayesian game under the performative setting.

2 Problem Formalization

2.1 Performative Bayesian Game

In a performative Bayesian game, each player seeks to minimize their own loss function under incomplete information with respect to decision-dependent data distributions. They compute a best response to the anticipated decisions of the others based on beliefs about the others' private information. For an index set $[n] = \{1, 2, \cdots, n\}$, each player $i \in [n]$ aims to minimize the loss function $\mathcal{L}_i(X)$ on player i's decision-dependent data distribution $\mathcal{D}_i(X)$ induced by the joint decision X:

$$\min_{x_i \in \mathcal{X}_i} \mathcal{L}_i(X), \text{ with } \mathcal{L}_i(X) := \mathop{\mathbb{E}}_{\substack{Z_i \sim \mathcal{D}_i(X) \\ \theta_{-i} \sim \mu(\theta_{-i}|h)}} \ell_i\left(Z_i, x_i, x_{-i}, \theta_i, \theta_{-i}\right), \tag{2}$$

where x_i is the decision vector of player i with dimension d_i, $\mathcal{X}_i \subset \mathbb{R}^{d_i}$ is the compact decision space of player i, $\mathcal{X} = \prod_{i\in[n]} \mathcal{X}_i \subset \mathbb{R}^d$ is the joint decision space of dimension $d = \sum_{i=1}^n d_i$, $X = (x_1, x_2, \cdots, x_n)^T \in \mathcal{X}$ is the joint decision vector of all players, $\theta_i \in \Theta_i$ is the type of player i that contains player i's private information, Θ_i is the compact type space of player i, $\theta_{-i} = (\theta_1, \cdots, \theta_{i-1}, \theta_{i+1}, \cdots, \theta_n)^T \in \Theta_{-i} = \prod_{j\neq i} \Theta_j$ is the type of all players except player i, $x_{-i} = (x_1, \cdots, x_{i-1}, x_{i+1}, \cdots, x_n)^T$ is the joint decision vector of all players except player i when their types are θ_{-i}, $\mathcal{D}_i(X)$ is player i's data distribution induced by joint decision vector X, $Z_i = \{z_{i1}, z_{i2}, \cdots, z_{im}\}$ is an i.i.d. sample set of m data points from $\mathcal{D}_i(X)$, $\mu = \prod_{i\in[n]} \mu(\theta_i)$ is the common independent probability distribution (belief) over all players' types conditioned on history h, implying $\mu(\theta_{-i}|\theta_i, h) = \mu(\theta_{-i}|h)$, $h = (X^{(0)}, X^{(1)}, \cdots, Z^{(0)}, Z^{(1)}, \cdots)$ is the sequence containing historical joint decisions and historical data samples, and $\ell_i(\cdot)$ is player i's loss function.

For notational convenience, we introduce $\mathcal{G}_\mu(X)$ to denote the static Bayesian game (without performative effect) over the data distribution induced by joint decision $X \in \mathcal{X}$ under the belief μ. This game comprises the loss functions of all n players and is defined as follows:

$$\mathcal{G}_\mu(X) := \left(\mathcal{L}_1(X), \mathcal{L}_2(X), \cdots, \mathcal{L}_n(X)\right)^T. \tag{3}$$

The joint decision $X^* = (x_1^*, x_2^*, \cdots, x_n^*)^T$ is a Bayesian equilibrium of $\mathcal{G}_\mu(X)$ if, for each player $i \in [n]$ with type θ_i, the decision x_i^* is a best response to the equilibrium decisions of other players x_{-i}^* based on the belief μ about the types of other players θ_{-i}. Specifically, for a given belief μ conditioned on history h and each player i with type θ_i, the following holds:

$$x_i^* \in \arg\min_{x_i \in \mathcal{X}_i} \mathcal{L}_i^*(X), \forall i \in [n], \text{with } \mathcal{L}_i^*(X) := \mathop{\mathbb{E}}_{\substack{Z_i \sim \mathcal{D}_i(X) \\ \theta_{-i} \sim \mu(\theta_{-i}|h)}} \ell_i\left(Z_i, x_i, x_{-i}^*, \theta_i, \theta_{-i}\right), \quad (4)$$

where $x_{-i}^* = (x_1^*, \cdots, x_{i-1}^*, x_{i+1}^*, \cdots, x_n^*)^T$ denotes the equilibrium decisions of all players except player i. We denote such a Bayesian equilibrium of $\mathcal{G}_\mu(X)$ with belief μ as $BE_\mu(X)$:

$$BE_\mu(X) := \left\{X^* \in \mathcal{X} : X^* \text{ is a Bayesian equilibrium of game } \mathcal{G}_\mu(X) \text{ with belief } \mu\right\}. \quad (5)$$

2.2 Repeated Performative Bayesian Game

We consider the standard settings in Bayesian games (Harsanyi, 1967; Fudenberg & Tirole, 1991b), in which players hold a common prior belief about other players' private information. Since data distribution maps are unknown to players and can be arbitrarily complex, players can only compute Bayesian equilibria under the current data distributions induced by the previous joint decision and based on the common belief μ. After observing all players' decisions, the common belief μ is updated. We define the repeated performative Bayesian game as follows:

Definition 2.1 (Repeated Performative Bayesian Game)**.** In a repeated performative Bayesian game, all players repeat two stages at each time step: (1) decision optimization and (2) belief update.

1. **Decision optimization:** The joint decision $X^{(t)}$ constitutes the Bayesian equilibrium under the current belief and the data distributions $\mathcal{D}_i(X^{(t-1)}), \forall i \in [n]$ induced by the previous joint decision $X^{(t-1)}$:
$$X^{(t)} \in BE_{\mu(\boldsymbol{\theta}|h^{(t)})}(X^{(t-1)}), \quad (6)$$
where $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_n)$ is the joint type, $h^{(t)} = (X^{(0)}, \cdots, X^{(t-1)}, Z^{(0)}, \cdots, Z^{(t-1)})$ contains joint decisions and data samples in previous time steps, $\{X^{(0)}, \cdots, X^{(t)}, \cdots\} \subseteq \mathcal{X}_{RPBG}$, $X^{(0)}$ is the initial joint decision inducing the initial data distributions, and $\mathcal{X}_{RPBG}$ denotes the compact convex hull of the joint decision sequence $\{X^{(t)}\}$ obtained during a repeated performative Bayesian game.

2. **Belief update:** The belief μ is common to all players and follows Bayesian updating after observing the joint decision $X^{(t)}$ (Harsanyi, 1967; Fudenberg & Tirole, 1991a). Specifically, the common belief is updated from $\mu(\boldsymbol{\theta}|h^{(t)})$ to $\mu(\boldsymbol{\theta}|h^{(t)}, X^{(t)})$ based on Bayes' rule. For player $i \in [n]$, history $h^{(t)}$, and decision $x_i^{(t)} \in \mathcal{X}_i$, for all $\theta_i \in \Theta_i$ with $\mu(\theta_i|h^{(t)}) > 0$ and $\varrho_i(x_i^{(t)}|h^{(t)}, \theta_i) > 0$, the updated belief is given by:
$$\mu\left(\theta_i|h^{(t)}, X^{(t)}\right) = \frac{\mu(\theta_i|h^{(t)})\varrho_i(x_i^{(t)}|h^{(t)}, \theta_i)}{\int_{\tilde{\theta}_i \in \Theta_i} \mu(\tilde{\theta}_i|h^{(t)})\varrho_i(x_i^{(t)}|h^{(t)}, \tilde{\theta}_i)\, d\tilde{\theta}_i}, \quad (7)$$
where $\varrho_i(x_i^{(t)}|h^{(t)}, \theta_i)$ is the probability that player i chooses decision $x_i^{(t)}$ given history $h^{(t)}$ and type θ_i, incorporating potential observation or decision noise.

2.3 Performatively Stable Bayesian Equilibrium

In a repeated performative Bayesian game, a desirable outcome is the performatively stable Bayesian equilibrium, where the common belief sequence $\mu^{(t)}$ converges to μ^{PS} and the joint decision X^{PS} is a Bayesian equilibrium under the induced data distribution $\mathcal{D}_i(X^{PS})$ for all $i \in [n]$, thereby eliminating the need to update decisions and beliefs. Formally, a performatively stable Bayesian equilibrium is defined as follows:

Definition 2.2 (Performatively Stable Bayesian Equilibrium)**.** A performatively stable Bayesian equilibrium of a repeated performative Bayesian game is a belief and joint decision pair (μ^{PS}, X^{PS}), where the common belief sequence $\mu^{(t)}$ converges to μ^{PS} and the joint decision X^{PS} is a Bayesian equilibrium with respect to the data distribution induced by X^{PS}:

$$X^{PS} \in BE_{\mu^{PS}}(X^{PS}). \quad (8)$$

3 THEORETICAL RESULTS

We present our main results regarding the convergence conditions for a performatively stable Bayesian equilibrium. First, to characterize $\mathcal{D}_i(X)$, we adopt a gradient-based sensitivity measure to quantify the influence of data distribution shifts on decision-making. Second, leveraging this measure, we derive the convergence conditions by analyzing the dynamics of both beliefs and decisions. Lastly, we extend convergence guarantees to the finite-sample setting.

3.1 GRADIENT-BASED SENSITIVITY MEASURE

In this subsection, we adopt the $\hat{\varepsilon}_i$-sensitivity measure for $\mathcal{D}_i(X)$, which was introduced by Zhong et al. (2026a) in the complete information setting. We extend this measure to the incomplete information setting by evaluating the gradient under the belief $\mu(\theta_{-i}|h)$.

Definition 3.1 ($\hat{\varepsilon}_i$-Sensitivity)**.** The data distribution map $\mathcal{D}_i(X)$ for each player $i \in [n]$ is $\hat{\varepsilon}_i$-sensitive if, for all joint decisions $Y, X, X' \in \mathcal{X}_{RPBG}$, the following inequality holds during a repeated performative Bayesian game:

$$\|\nabla_{y_i}\mathcal{L}_i(X) - \nabla_{y_i}\mathcal{L}_i(X')\|_2 \leq \hat{\varepsilon}_i\|X - X'\|_2, \tag{9}$$

where $\nabla_{y_i}\mathcal{L}_i(X) := \underset{\substack{Z_i\sim\mathcal{D}_i(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}}{\mathbb{E}} \nabla_{y_i}\ell_i\left(Z_i, y_i, y_{-i}, \theta_i, \theta_{-i}\right)$ with y_i, y_{-i} constituting Y.

Note that, in the incomplete information setting, the $\hat{\varepsilon}_i$-sensitivity of $\mathcal{D}_i(X)$ can be estimated by each player i, since all the elements in $\nabla_{y_i}\mathcal{L}_i(X)$ are known to player i and the measure is defined on the decision space $\mathcal{X}_{RPBG}$ that contains all joint decisions in a repeated performative Bayesian game. An approach for estimating $\hat{\varepsilon}_i$ is provided in Algorithm 1 (Section 4)[1].

3.2 THEORETICAL GUARANTEES FOR PERFORMATIVELY STABLE BAYESIAN EQUILIBRIA

Before deriving the convergence conditions for a performatively stable Bayesian equilibrium, we first introduce the definition of individual gradients and the assumption of α-strong monotonicity.

Definition 3.2 (Individual Gradients of Game $\mathcal{G}_\mu(\cdot)$)**.** The individual gradients of the joint decision Y, under the data distribution induced by X, are defined as:

$$\nabla_Y\mathcal{G}_\mu(X) := \left(\nabla_{y_1}\mathcal{L}_1(X), \cdots, \nabla_{y_n}\mathcal{L}_n(X)\right). \tag{10}$$

Assumption 3.3 (α-Strong Monotonicity)**.** A game $\mathcal{G}_\mu(X)$ is α-strongly monotone if, for any $X \in \mathcal{X}_{RPBG}$, all joint decisions $Y, Y' \in \mathcal{X}_{RPBG}$ and beliefs μ that arise in the repeated performative Bayesian game, the following holds:

$$\alpha\|Y - Y'\|_2^2 \leq \left(\nabla_Y\mathcal{G}_\mu(X) - \nabla_{Y'}\mathcal{G}_\mu(X)\right)^T(Y - Y'). \tag{11}$$

The α-strong monotonicity, together with the compact decision and type spaces, guarantees the existence and uniqueness of a Bayesian equilibrium in each iteration of a repeated performative Bayesian game (Fudenberg & Tirole, 1991a), thus ensuring tractable game dynamics. This condition is widely adopted in multiplayer performative prediction games (Narang et al., 2023; Cutler et al., 2024; Wang et al., 2023; Zhong et al., 2026a) and can be constructed in multi-agent systems (Wen et al., 2022). Note that α-strong monotonicity does not guarantee convergence to a performatively stable Bayesian equilibrium due to the data distribution shift caused by the performative effect.

With this assumption, we can now present three lemmas that investigate the convergence dynamics in a repeated performative Bayesian game.

Lemma 3.4. *The Bayesian updated beliefs converge almost surely to a limiting probability measure $\mu^{(\infty)}$ on Θ. Specifically, for every measurable set $\mathcal{A} \subseteq \Theta$, we have:*

$$|\mu^{(t)}(\mathcal{A}) - \mu^{(\infty)}(\mathcal{A})| \xrightarrow{a.s.} 0 \text{ as } t \to \infty, \tag{12}$$

where $\mu^{(t)}(\mathcal{A}) = \mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)})$ for $t = 1, 2, \cdots$.

[1]Note that $\hat{\varepsilon}_i$-sensitivity requires the gradients $\nabla_{y_i}\mathcal{L}_i(X)$ to be continuous in the joint decision X that induces the distributions.

Lemma 3.5. *Suppose Assumption 3.3 holds, the loss functions $\ell_i(\cdot)$ for all $i \in [n]$ are continuous on the product space $\mathcal{X}_i \times \mathcal{X}_{-i} \times \Theta$, and the belief μ follows Bayesian updating. Then for any fixed joint decision X, a sequence of Bayesian equilibrium mappings $\{BE_{\mu^{(0)}}(X),\ BE_{\mu^{(1)}}(X),\ \cdots, BE_{\mu^{(t)}}(X), \cdots\}$ satisfies:*

$$\left\| BE_{\mu^{(t)}}(X) - BE_{\mu^{(\infty)}}(X) \right\|_2 \xrightarrow{a.s.} 0 \text{ as } t \to \infty. \tag{13}$$

Lemma 3.6. *Suppose Assumption 3.3 holds, the data distribution maps $\mathcal{D}_i(X)$ for all $i \in [n]$ are $\hat{\varepsilon}_i$-sensitive, and the belief μ is fixed. Then for any joint decisions $X, X' \in \mathcal{X}_{RPBG}$, the Bayesian equilibrium map satisfies:*

$$\left\| BE_\mu(X) - BE_\mu(X') \right\|_2 \leq \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \left\| X - X' \right\|_2. \tag{14}$$

The proofs of Lemmas 3.4-3.6 are given in Appendices B-D. Note that Lemma 3.6 shows the dynamics under a fixed belief, where its proof parallels the complete-information convergence analysis in (Narang et al., 2023; Zhong et al., 2026a). Our main challenge is to establish the convergence of both the Bayesian updated beliefs and the decisions. Based on Lemmas 3.4-3.6, we provide the convergence conditions for a performatively stable Bayesian equilibrium.

Theorem 3.7. *Suppose Assumption 3.3 holds, the loss functions $\ell_i(\cdot)$ for all $i \in [n]$ are continuous on the product space $\mathcal{X}_i \times \mathcal{X}_{-i} \times \Theta$, data distribution maps $\mathcal{D}_i(X)$ for all $i \in [n]$ are $\hat{\varepsilon}_i$-sensitive, and $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$. Then we have the following:*

1. *The Bayesian updated beliefs converge almost surely to the limiting probability measure μ^{PS} on Θ. Specifically, for every measurable set $\mathcal{A} \subseteq \Theta$, we have:*

$$|\mu^{(t)}(\mathcal{A}) - \mu^{PS}(\mathcal{A})| \xrightarrow{a.s.} 0 \text{ as } t \to \infty, \tag{15}$$

   where $\mu^{(t)}(\mathcal{A})$ is as defined in Lemma 3.4, and μ^{PS} is as defined in Definition 2.2.

2. *The sequence of joint decisions $\{X^{(0)}, X^{(1)}, \cdots, X^{(t)}, \cdots\} \subseteq \mathcal{X}_{RPBG}$ satisfies:*

$$\left\| X^{(t)} - X^{PS} \right\|_2 \xrightarrow{a.s.} 0 \text{ as } t \to \infty, \tag{16}$$

   where X^{PS} is as defined in Definition 2.2.

The proof of Theorem 3.7 is given in Appendix E. Note that μ^{PS} does not imply that players' private information is fully revealed, since players with distinct types may make the same decision, making it impossible to identify true types even with a converged belief (Fudenberg & Tirole, 1991a).

3.3 Finite Sample Analysis

In this section, we extend our results to a more practical setting, where each player can only collect finite samples from the data distribution. We use the following standard variance condition (Narang et al., 2023; Wang et al., 2023; Li et al., 2025a) to establish the finite sample convergence.

Assumption 3.8 (Finite Variance)**.** There exists a constant $\sigma \geq 0$ satisfying:

$$\left\| \overline{\nabla_Y \mathcal{G}_\mu(X)} - \nabla_Y \mathcal{G}_\mu(X) \right\|_2 \leq \sigma \text{ for all } X, Y \in \mathcal{X}_{RPBG}, \tag{17}$$

where $\overline{\nabla_Y \mathcal{G}_\mu(X)} := (\overline{\nabla_{y_1} \mathcal{L}_1(X)}, \cdots, \overline{\nabla_{y_n} \mathcal{L}_n(X)})^T$ is the empirical average of the individual gradients, $\overline{\nabla_{y_i} \mathcal{L}_i(X)} := \frac{1}{m} \sum_{j=1}^m \mathop{\mathbb{E}}_{\theta_{-i} \sim \mu(\theta_{-i}|\theta_i,h)} \nabla_{y_i} \ell_i(z_{i_j}, y_i, y_{-i}, \theta_i, \theta_{-i})$ is the empirical average gradient of player i over m samples.

With this assumption, we can now establish the convergence conditions in finite sample scenarios.

Theorem 3.9. *Suppose Assumptions 3.3 and 3.8 hold, the loss functions $\ell_i(\cdot)$ for all $i \in [n]$ are continuous on the product space $\mathcal{X}_i \times \mathcal{X}_{-i} \times \Theta$, data distribution maps $\mathcal{D}_i(X)$ for all $i \in [n]$ are $\hat{\varepsilon}_i$-sensitive, and $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$. Then we have the following:*

1. *The Bayesian updated beliefs with finite samples converge almost surely to the limiting probability measure μ^{PS} on Θ. Specifically, for every measurable set $\mathcal{A} \subseteq \Theta$, we have:*

$$|\mu^{(t)}(\mathcal{A}) - \mu^{PS}(\mathcal{A})| \xrightarrow{a.s.} 0 \text{ as } t \to \infty, \tag{18}$$

where $\mu^{(t)}(\mathcal{A})$ is as defined in Lemma 3.4, and μ^{PS} is as defined in Definition 2.2.

2. *The sequence of joint decisions $\{X^{(0)}, X^{(1)}, \cdots, X^{(t)}, \cdots\} \subseteq \mathcal{X}_{RPBG}$ with finite samples satisfies:*

$$\|X^{(t)} - X^{PS}\|_2 \leq \frac{\sigma}{\alpha - \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}} \text{ as } t \to \infty, \tag{19}$$

where X^{PS} is as defined in Definition 2.2.

The proof of Theorem 3.9 can be found in Appendix F[2]. Theorem 3.9 guarantees finite-sample convergence, which informs practical algorithm design in the next section.

4 Algorithm Design

In this section, we design Algorithm 1 to achieve a performatively stable Bayesian equilibrium. For an original performative Bayesian game $\mathcal{G}_\mu(\cdot)$ with ψ-strong monotonicity, our algorithm requires each player to add a γ-strongly convex regularizer $\frac{\gamma}{2}\|x_i\|_2^2$ to loss function. This regularization ensures that the regularized game $\hat{\mathcal{G}}_\mu(\cdot)$ is α-strongly monotone with $\alpha = \psi + \gamma > \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$, where γ is tuned based on the measured $\hat{\varepsilon}_i$. For such strongly monotone games with compact decision and type spaces, the computation of a Bayesian equilibrium is tractable using numerical methods. According to Theorem 3.9, Algorithm 1 guarantees convergence to a performatively stable Bayesian equilibrium[3]. Detailed descriptions are provided in Appendix H.

Note that our algorithm guarantees convergence to a performatively stable Bayesian equilibrium of the regularized game $\hat{\mathcal{G}}_\mu(\cdot)$. Theorem 4.1 establishes that the quadratic regularizer $R(x_i) = \frac{\gamma}{2}\|x_i\|_2^2$ used in Algorithm 1 minimizes the upper bound of the distance between the original and regularized equilibria among all γ-strongly convex regularizers.

Theorem 4.1. *Let $\mathcal{R}$ be the family of all γ-strongly convex regularizers defined on the decision x_i for all $i \in [n]$. The distance $\|X^* - X^R\|_2$ has an upper bound that depends on the regularizer $R(\cdot) \in \mathcal{R}$, where X^* and X^R are the Bayesian equilibria of the original and regularized games, respectively. Among all choices of $R(\cdot) \in \mathcal{R}$, the quadratic regularizer $R(x_i) = \frac{\gamma}{2}\|x_i\|_2^2$ minimizes this upper bound. For the quadratic regularizer $R(x_i) = \frac{\gamma}{2}\|x_i\|_2^2$, we have:*

$$\|X^* - X^R\|_2 \leq \frac{\gamma \|X^*\|_2}{\psi + \gamma}. \tag{20}$$

The proof of Theorem 4.1 is provided in Appendix J.

5 Numerical Experiments

We perform numerical experiments in two performative Bayesian games: the Cournot competition in the crude oil trade and the Bertrand competition in the ride-share market[4]. To validate our theoretical analysis and the proposed algorithm, we set four values for the parameter c in Algorithm 1: $c = 0, 0.3$ (violating the convergence conditions), and $c = 1.1, 30$ (satisfying the conditions). In addition, we include three baselines to show the different dynamics in repeated performative Bayesian games: (1) Static Bayesian Equilibrium, (2) Complete Information Equilibrium, and (3) Heuristic Strategy. More details on experimental settings can be found in Appendix K.1. The performance comparison over wall-clock time is provided in Appendix K.2.

[2]Our finite sample analysis can be extended via other statistical tools. We provide a variant of Theorem 3.9 with the Central Limit Theorem in Appendix G, under the idealized large-sample assumption (as $m \to \infty$).

[3]Player i reports $\hat{\varepsilon}_i$ in line 11 of Algorithm 1. We analyze the truthfulness of this report in Appendix I.

[4]The source code of experiments is available at: 'anonymous.4open.science/r/PerformativeBayesianGame'.

Algorithm 1 Algorithm for Performatively Stable Bayesian Equilibrium: Procedures for Player $i \in [n]$

Input: Parameter $c \in (1, +\infty)$, initial time step $t = 0$,initial sensitivity $\hat{\varepsilon}_i \geq 0$, convergence thresholds ρ_1, ρ_2,initial joint decision $X^{(0)} \in \mathcal{X}_{RPBG}$, initial datasets $Z_i^{(0)} \sim \mathcal{D}_i(X^{(0)})$, initial belief $\mu(\boldsymbol{\theta})$, initial history $h^{(0)} = \{X^{(0)}, Z^{(0)}\}$, ψ-strongly monotone game $\mathcal{G} := (\mathcal{L}_1, \cdots, \mathcal{L}_n)$
Output: X^{PS}, μ^{PS} at a performatively stable Bayesian equilibrium
repeat
  $t \leftarrow t + 1$;
  $\gamma \leftarrow \max(0, c\sqrt{\sum_{k=1}^{n} \hat{\varepsilon}_k^2} - \psi)$;
  $\hat{\mathcal{L}}_i \leftarrow \mathop{\mathbb{E}}_{\substack{Z_i^{(t-1)} \sim \mathcal{D}_i(X^{(t-1)}) \\ \theta_{-i} \sim \mu(\theta_{-i}|h^{(t-1)})}} \ell_i(Z_i^{(t-1)}, x_i, x_{-i}, \theta_i, \theta_{-i}) + \frac{\gamma}{2}\|x_i\|_2^2$;
  $x_i^{(t)} \leftarrow \arg\min_{x_i} \hat{\mathcal{L}}_i(X^{(t-1)})$;
  $X^{(t)} \leftarrow (x_1^{(t)}, x_2^{(t)}, \cdots, x_n^{(t)})^T$;
  Collect data samples $Z_i^{(t)} \sim \mathcal{D}_i(X^{(t)})$;
  $\hat{\varepsilon}_i \leftarrow \max\left(\hat{\varepsilon}_i, \frac{\left\|\nabla_{x_i^{(t)}} \mathcal{L}_i(X^{(t)}) - \nabla_{x_i^{(t)}} \mathcal{L}_i(X^{(t-1)})\right\|_2}{\|X^{(t)} - X^{(t-1)}\|_2}\right)$;
  Report $\hat{\varepsilon}_i$;
  $\mu\left(\boldsymbol{\theta}|h^{(t)}, X^{(t)}, Z^{(t)}\right) \leftarrow \frac{\mu(\boldsymbol{\theta}|h^{(t)})\varrho(X^{(t)}|h^{(t)}, \boldsymbol{\theta})}{\int_{\tilde{\boldsymbol{\theta}} \in \Theta} \mu(\tilde{\boldsymbol{\theta}}|h^{(t)})\varrho(X^{(t)}|h^{(t)}, \tilde{\boldsymbol{\theta}})\, d\tilde{\boldsymbol{\theta}}}$;
  $h^{(t)} \leftarrow (h^{(t-1)}, X^{(t)}, Z^{(t)})$;
until $\|X^{(t)} - X^{(t-1)}\|_2 \leq \rho_1$ and $|\mu^{(t)}(\mathcal{A}) - \mu^{(t-1)}(\mathcal{A})| \leq \rho_2$
$X^{PS} \leftarrow X^{(t)}$;
$\mu^{PS} \leftarrow \mu\left(\boldsymbol{\theta}|h^{(t)}, X^{(t)}\right)$;

5.1 Cournot Competition in Crude Oil Trade

Game Abstraction.

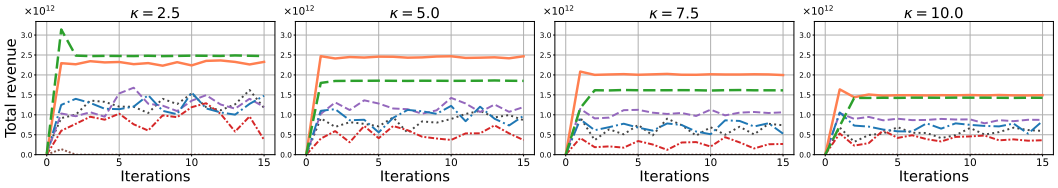
In the Cournot competition, $n = 28$ crude oil exporting countries set export quantity adjustments $x_i \in \mathbb{R}$ to maximize revenue[5]. For the initial export quantity $q_i \in \mathbb{R}$, the actual export quantity is $x_i + q_i$. The crude oil price P is given by a function $P = z - b\sum_{i=1}^{n}(q_i + x_i)$, where $z \sim \mathcal{D}(X)$, and $b \in \mathbb{R}$ is the price sensitivity parameter. The private cost $c_i \in \mathbb{R}$ of country i is sampled from a normal distribution $\mathcal{N}(\theta_i, \varsigma^2)$, where θ_i is the type of country i and ς^2 is the variance. The common prior belief is a uniform distribution $\mu(\theta_i) = \mathcal{U}[\theta_-, \theta_+]$. Each country i maximizes its revenue by $\min_{x_i} \mathbb{E}_{z \sim \mathcal{D}(X), \theta_{-i} \sim \mu(\theta_{-i}|h)}\left[-(P - c_i) \cdot (q_i + x_i)\right]$. The data distribution map $\mathcal{D}(X)$ is identical across all countries and is designed as:

$$z \sim \mathcal{D}(X) = z_0 - \kappa \cdot \sinh^{-1}\left(\sum_{i=1}^{n} x_i\right) + w, \tag{21}$$

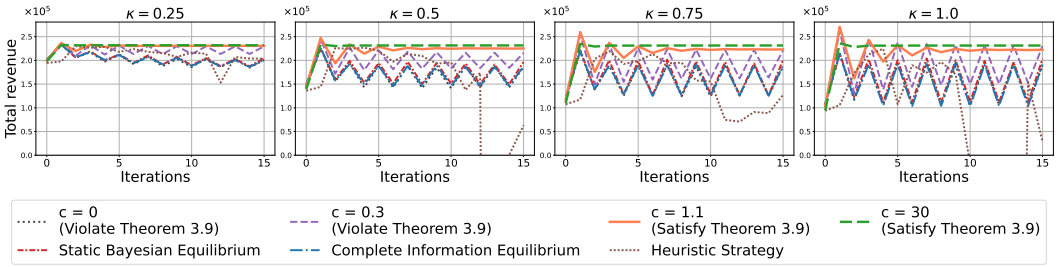
where z_0 is a parameter, κ controls the sensitivity of $\mathcal{D}(X)$, and w is the noise term. In our experiment, we set $\kappa = \{2.5, 5, 7.5, 10\}$, where larger values of κ correspond to higher sensitivities. Further details are provided in Appendix K.3.

Result Analysis. The total revenue over 15 iterations is shown in Fig. 1a. For $c = 1.1, 30$ (satisfying the convergence condition $c > 1$), a performatively stable Bayesian equilibrium is achieved, while for $c = 0, 0.3$ the revenue fluctuates and is suboptimal across all κ. Static Bayesian Equilibrium and Complete Information Equilibrium suffer from the distribution shift caused by performativity. Heuristic Strategy performs the worst because it fits a linear map to the non-linear distribution map $\mathcal{D}(X)$, leading to a mis-specified game structure. Additional results (Table 1 and Figs. 4-6) are provided in Appendix K.4.

[5]The dataset is available at: 'kaggle.com/datasets/toriqulstu/global-crude-petroleum-trade-1995-2021'.



(a) Total revenue in Cournot competition within crude oil trade with $\kappa = \{2.5, 5, 7.5, 10\}$



(b) Total revenue in Bertrand competition within ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1\}$

Figure 1: Experimental results with $c = 0, 0.3, 1.1, 30$ over 10 trials.

5.2 Bertrand Competition in Ride-Share Market

Game Abstraction. In the Bertrand competition, two companies $n = 2$, Lyft and Uber, compete across 11 locations in Boston[6] by setting price adjustments $x_i \in \mathbb{R}^{11}$ to maximize revenue. For the initial price p, the actual price is $x_i + p\mathbf{1}$ with $\mathbf{1} \in \mathbb{R}^{11}$ the all-ones vector. The demand for company i is given by $q_i = z_i + \mathbf{A}_i x_i + \mathbf{A}_{-i} x_{-i}$, where $z_i \sim \mathcal{D}_i(X)$, and $\mathbf{A}_i, \mathbf{A}_{-i} \in \mathbb{R}^{11\times 11}$ are price sensitivity matrices. The private cost $c_i \in \mathbb{R}^{11}$ is sampled from a normal distribution $\mathcal{N}_{11}(\theta_i, \mathbf{\Sigma})$, where $\theta_i \in \mathbb{R}^{11}$ is the type of company i and $\mathbf{\Sigma} \in \mathbb{R}^{11\times 11}$ is the covariance matrix. The common prior belief is a uniform distribution $\mu(\theta_i) = \mathcal{U}[\theta_-, \theta_+]^{11}$. Company i maximizes its revenue by $\min_{x_i} \mathbb{E}_{z_i\sim\mathcal{D}_i(X),\theta_{-i}\sim\mu(\theta_{-i}|h)} \left[-(p\mathbf{1} + x_i - c_i)^T q_i\right]$. The data distribution map $\mathcal{D}_i(X)$ is:

$$z_i \sim \mathcal{D}_i(X) = 2z_{0,i} \oslash \left(1 + e^{-(\mathbf{B}_i x_i + \mathbf{B}_{-i} x_{-i})}\right) + w_i, \tag{22}$$

where $z_{0,i} \in \mathbb{R}^{11}$ is the initial demand of company i, $\mathbf{B}_i, \mathbf{B}_{-i} \in \mathbb{R}^{11\times 11}$ are unknown to the companies, $\oslash$ denotes the Hadamard division, and w_i is the noise term. We set $\kappa = \{0.25, 0.5, 0.75, 1.0\}$, increasing the sensitivity as κ grows. For further details, please refer to Appendix K.5.

Result Analysis. The total revenue of the two companies over 15 iterations is shown in Fig. 1b. Consistent with the Cournot experiment, a performatively stable Bayesian equilibrium is achieved for $c = 1.1, 30$, while $c = 0, 0.3$ and the three baselines fluctuate and achieve low revenue. This demonstrates the robustness of the proposed algorithm in a higher-dimensional setting. Additional results (Table 2 and Figs. 7-10) are provided in Appendix K.6.

6 Conclusion

In this paper, we formalized and analyzed performative Bayesian games. With the proposed $\hat{\varepsilon}_i$-sensitivity measure, we established convergence conditions for a performatively stable Bayesian equilibrium for α-strongly monotone games in finite-sample scenarios. Accordingly, we designed an algorithm that ensures convergence to such an equilibrium. Numerical experiments validated our theoretical results and the effectiveness of the proposed algorithm. The potential limitations of our framework and future directions are discussed in Appendix L.

[6]The dataset is available at: 'kaggle.com/datasets/brllrb/uber-and-lyft-dataset-boston-ma'.

AI USE STATEMENT

In this work, we used generative AI tools for implementing methods. We have not used generative AI tools for developing theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs, proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, assisting with the translation of the manuscript, supporting qualitative and thematic data analysis, or interpreting results. Generating synthetic datasets and cleaning or reformatting datasets are not applicable to this work. We have reviewed all AI-assisted work. The AI-assisted code was verified and tested by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The source code for all experiments is available via the anonymous link provided in Section 5. The complete proofs of all theoretical results, together with the auxiliary definitions and lemmas used in the analysis, are provided in the appendix. The detailed experimental settings, including parameter settings, convergence thresholds, and the full game setups for the Cournot and Bertrand competitions, are described in the Appendices K.1, K.3, and K.5. The two public datasets used in the experiments are referenced in the footnotes of Section 5.

REFERENCES

Heinz Bauer. *Measure and integration theory*, volume 26 of *de Gruyter Studies in Mathematics*. de Gruyter, Berlin, June 2001.

Heinz H Bauschke and Patrick L Combettes. Correction to: convex analysis and monotone operator theory in hilbert spaces. In *Convex analysis and monotone operator theory in Hilbert spaces*, pp. C1–C4. Springer, 2020.

Javier Sanguino Bautiste, Thomas Kehrenberg, Carlos Rosety, Jose A. Lozano, and Novi Quadrianto. Performative prediction made practical. In *ICLR Blogposts*, 2026.

Claude Berge. *Topological Spaces*. Oliver & Boyd, 1963.

Joseph Bertrand. Review of "theorie mathematique de la richesse sociale" and of "recherches sur les principles mathematiques de la theorie des richesses.". *Journal de savants*, 67:499, 1883.

Daniele Bracale, Subha Maity, Yuekai Sun, and Moulinath Banerjee. Learning the distribution map in reverse causal performative prediction. In *International Conference on Artificial Intelligence and Statistics*, volume 258, pp. 973–981. PMLR, 2025.

Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

Yatong Chen, Wei Tang, Chien-Ju Ho, and Yang Liu. Performative prediction with bandit feedback: Learning through reparameterization. In *International Conference on Machine Learning*, volume 235, pp. 7298–7324. PMLR, 2024.

A.A. Cournot, T. Bacon, and I. Fisher. *Researches Into the Mathematical Principles of the Theory of Wealth*. Economic classics. Macmillan, 1897.

Joshua Cutler, Mateo Díaz, and Dmitriy Drusvyatskiy. Stochastic approximation with decision-dependent distributions: Asymptotic normality and optimality. *Journal of Machine Learning Research*, 25(90):1–49, 2024.

Drew Fudenberg and Jean Tirole. *Game theory*. MIT press, 1991a.

Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991b.

John C Harsanyi. Games with incomplete information played by "bayesian" players, i–iii part i. the basic model. *Management science*, 14(3):159–182, 1967.

Miao Hu, Wenzhuo Yang, Zhenxiao Luo, Xuezheng Liu, Yipeng Zhou, Xu Chen, and Di Wu. Autofl: A bayesian game approach for autonomous client participation in federated edge learning. *IEEE Transactions on Mobile Computing*, 23(1):194–208, 2024.

Kexin Huang, Ziqian Chen, Xue Wang, Chongming Gao, Jinyang Gao, Bolin Ding, and Xiang Wang. Learning bayesian nash equilibrium in auction games via approximate best response. In *International Conference on Machine Learning*, 2025.

John K Hunter and Bruno Nachtergaele. *Applied analysis*. World Scientific, 2001.

Kun Jin, Tongxin Yin, Zhongzhu Chen, Zeyu Sun, Xueru Zhang, Yang Liu, and Mingyan Liu. Performative federated learning: A solution to model-dependent and heterogeneous distribution shifts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 12938–12946, 2024.

Achim Klenke. *Probability theory: a comprehensive course*. Springer, 2008.

Gunnar König, Hidde Fokkema, Timo Freiesleben, Celestine Mendler-Dünner, and Ulrike von Luxburg. Performative validity of recourse explanations. In *Advances in Neural Information Processing Systems*, 2025.

Kai Li, Hang Xu, Haobo Fu, Qiang Fu, and Junliang Xing. Automatically designing counterfactual regret minimization algorithms for solving imperfect-information games. *Artificial Intelligence*, 337:104232, 2024a.

Qiang Li, Chung-Yiu Yau, and Hoi-To Wai. Multi-agent performative prediction with greedy deployment and consensus seeking agents. In *Advances in Neural Information Processing Systems*, volume 35, pp. 38449–38460, 2022.

Qiang Li, Michal Yemini, and Hoi To Wai. Clipped SGD algorithms for performative prediction: Tight bounds for stochastic bias and remedies. In *International Conference on Machine Learning*, 2025a.

Simin Li, Jun Guo, Jingqiao Xiu, Ruixiao Xu, Xin Yu, Jiakai Wang, Aishan Liu, Yaodong Yang, and Xianglong Liu. Byzantine robust cooperative multi-agent reinforcement learning as a bayesian game. In *International Conference on Learning Representations*, 2024b.

Victor Li, Baiting Chen, Yuzhen Mao, Qi Lei, and Zhun Deng. Performative risk control: Calibrating models for reliable deployment under performativity. In *Advances in Neural Information Processing Systems*, 2025b.

Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. In *Advances in Neural Information Processing Systems*, volume 33, pp. 4929–4939, 2020.

Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research*, 24(202):1–56, 2023.

Karl Pearson. Contributions to the mathematical theory of evolution. *Proceedings of the Royal Society of London*, 54:329–333, 1893.

Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, volume 119, pp. 7599–7609. PMLR, 2020.

Robert L. Phillips. *Pricing and revenue optimization, Second Edition*. Stanford University Press, Redwood City, 2021. ISBN 9781503614260.

Georgios Piliouras and Fang-Yi Yu. Multi-agent performative prediction: From global stability and optimality to chaos. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 1047–1074, 2023.

Vasilis Pollatos, Debmalya Mandal, and Goran Radanovic. On corruption-robustness in performative reinforcement learning. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence*. AAAI Press, 2025. ISBN 978-1-57735-897-8.

R Tyrrell Rockafellar and Roger JB Wets. *Variational analysis*. Springer, 1998.

W Rudin. *Principles of mathematical analysis*. McGraw-Hill, United States, 3rd edition, 1976. Includes Index.

Michael Spence. Job market signaling. In *Uncertainty in economics*, pp. 281–306. Elsevier, 1978.

Nikita Tsoy, Ivan Kirev, Negin Rahimiyazdi, and Nikola Konstantinov. On the impact of performative risk minimization for binary random variables. In *International Conference on Machine Learning*, 2025.

Alexandra Vintila and Mihai Daniel Roman. Bertrand competition under incomplete information. In *International Conference on Economics and Social Sciences*, pp. 441, 2021.

Xiaolu Wang, Chung-Yiu Yau, and Hoi To Wai. Network effects in performative prediction games. In *International Conference on Machine Learning*, volume 202, pp. 36514–36540, 2023.

Muning Wen, Jakub Kuba, Runji Lin, Weinan Zhang, Ying Wen, Jun Wang, and Yaodong Yang. Multi-agent reinforcement learning is a sequence modeling problem. In *Advances in Neural Information Processing Systems*, volume 35, pp. 16509–16521, 2022.

David Williams. *Probability with martingales*. Cambridge university press, 1991.

Zhichen Xiang, Hongke Zhao, Chuang Zhao, Ming He, and Jianping Fan. Performative debias with fair-exposure optimization driven by strategic agents in recommender systems. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3507–3517, 2024.

Wenjing Yan and Xuanyu Cao. Decentralized noncooperative games with coupled decision-dependent distributions. In *Advances in Neural Information Processing Systems*, volume 37, pp. 25592–25627, 2024.

Xie Yi, Zhanke Zhou, Chentao Cao, Qiyu Niu, Tongliang Liu, and Bo Han. From debate to equilibrium: Belief-driven multi-agent LLM reasoning via bayesian nash equilibrium. In *International Conference on Machine Learning*, 2025.

Guangzheng Zhong, Yang Liu, and Jiming Liu. Towards performatively stable equilibria in decision-dependent games for arbitrary data distribution maps. *Machine Learning*, 115(8):197, 2026a.

Guangzheng Zhong, Yang Liu, Ruichen Liu, and Jiming Liu. Performative prediction in the wild: Adapting to arbitrary data distribution maps. *Machine Learning*, 115(3):42, 2026b.

Appendix

A Auxiliary Definitions, Lemma and Theorems

Theorem A.1 (Tower Property (Williams, 1991)). *If $\mathcal{H} \subseteq \mathcal{G}$, then the following holds almost surely:*

$$\mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{H}] = \mathbb{E}[X|\mathcal{H}]. \tag{23}$$

Definition A.2 (Martingale (Williams, 1991)). A sequence $Y_1, Y_2, \cdots$ is said to be a martingale with respect to another sequence $X_1, X_2, \cdots$ if for all n, the following holds almost surely:

$$\begin{aligned} &\mathbb{E}[|Y_n|] < \infty, \\ &\mathbb{E}[Y_{n+1}|X_1, X_2, \cdots, X_n] = Y_n. \end{aligned} \tag{24}$$

Theorem A.3 (Martingale Convergence Theorem (Williams, 1991)). *If X_t is a martingale with $|X| < M < \infty$ for some M and all t, then there exists a random variable X such that X_t converges to X almost surely (with probability 1), i.e.,:*

$$|X_n - X| \xrightarrow{a.s.} 0 \text{ as } n \to \infty. \tag{25}$$

Theorem A.4 (Berge's Maximum Theorem (Berge, 1963)). *Let $f : \mathcal{X} \times \Theta \to \mathbb{R}$ be a continuous function on the product space $\mathcal{X} \times \Theta$, and $C : \Theta \rightrightarrows \mathcal{X}$ be a compact-valued correspondence. Define the value function $f^* : \Theta \to \mathbb{R}$:*

$$f^*(\theta) = \max\{f(x, \theta) : x \in C(\theta)\}, \tag{26}$$

and the maximizer function $x^* : \Theta \rightrightarrows \mathcal{X}$:

$$x^*(\theta) = \arg\max\{f(x, \theta) : x \in C(\theta)\} = \{x \in C(\theta) : f(x, \theta) = f^*(\theta)\}. \tag{27}$$

If $C(\theta)$ is continuous at θ and the maximizer is unique for all $\theta \in \Theta$, then the value function $f^*(\theta)$ is continuous, and the maximizer function $x^*(\theta)$ is continuous.

Theorem A.5 (Portmanteau Theorem (Klenke, 2008)). *Let E be a metric space and let $\{\mu_n\}$ be a bounded sequence of non-negative probability measures, where $\{\mu_n\} = \mu_1, \mu_2, \cdots$ with each $\mu_n \in \mathcal{M}_{\leq 1}(E)$, and $\mathcal{M}_{\leq 1}(E)$ is the set of all probability measures defined on E with non-negative values and total probability at most 1. The following are equivalent:*

1. *$\mu_n \xrightarrow{w} \mu$ (i.e., μ_n converges weakly to μ as $n \to \infty$),*
2. *$\int_E f d\mu_n \xrightarrow{n\to\infty} \int_E f d\mu$ for all bounded continuous functions $f : E \to \mathbb{R}$.*

Theorem A.6 (Heine–Cantor Theorem (Rudin, 1976)). *Let f be a continuous mapping of a compact metric space X into a metric space Y. Then f is uniformly continuous on X.*

Definition A.7 (Uniform continuity (Rudin, 1976)). Let f be a mapping of a metric space X into a metric space Y. We say that f is uniformly continuous on X if for every $\eta > 0$ there exists a $\delta > 0$ such that

$$d_Y(f(p), f(q)) < \eta, \tag{28}$$

for all p and q in X for which $d_X(p, q) < \delta$. Here, d_Y and d_X refer to the distances in X and Y, respectively.

Definition A.8 (Equicontinuity (Rudin, 1976)). A family $\mathcal{F}$ of complex functions f defined on a set E in a metric space X is said to be equicontinuous on E if for every $\eta > 0$ there exists a $\delta > 0$ such that

$$|f(x) - f(y)| < \eta, \tag{29}$$

whenever $d(x, y) < \delta, x \in E, y \in E$, and $f \in \mathcal{F}$. Here, d denotes the metric of X.

Lemma A.9 (Equicontinuity and Uniform Convergence (Rudin, 1976)). *Suppose $\{f_n\}$ is an equicontinuous sequence of functions on a compact set K, and $\{f_n\}$ converges pointwise on K, then $\{f_n\}$ converges uniformly on K.*

B PROOF OF LEMMA 3.4

Proof. Define $\mathcal{A}$ as a measurable subset of type space Θ, i.e., $\mathcal{A} \subseteq \Theta$, and consider the indicator function $I_{\mathcal{A}}(\boldsymbol{\theta})$ for the measurable set $\mathcal{A}$ which is the random variable given by:

$$I_{\mathcal{A}}(\boldsymbol{\theta}) = \begin{cases} 1 & \text{if } \boldsymbol{\theta} \in \mathcal{A} \\ 0 & \text{if } \boldsymbol{\theta} \notin \mathcal{A} \end{cases}. \tag{30}$$

We have $\mu(\boldsymbol{\theta} \in \mathcal{A}) = \mathbb{E}[I_{\mathcal{A}}(\boldsymbol{\theta})]$, i.e., the probability of $\boldsymbol{\theta}$ in the measurable set $\mathcal{A}$ is the expectation of the indicator function $I_{\mathcal{A}}(\boldsymbol{\theta})$. Since $h^{(t)} \subseteq h^{(t+1)}$, according to the tower property (Theorem A.1), we have almost surely:

$$\mathbb{E}[\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t+1)})|h^{(t)}] = \mathbb{E}[\mathbb{E}[I_{\mathcal{A}}(\boldsymbol{\theta})|h^{(t+1)}]|h^{(t)}] = \mathbb{E}[I_{\mathcal{A}}(\boldsymbol{\theta})|h^{(t)}] = \mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)}), \tag{31}$$

which holds for any measurable set $\mathcal{A} \subseteq \Theta$.

Thus, according to the definition of martingale (Definition A.2), the sequence of beliefs in any measurable set $\{\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)})\}$ is a martingale for any measurable set $\mathcal{A} \subseteq \Theta$ with respect to the history sequence $\{h^{(t)}\}$.

Since $0 \leq \mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)}) \leq 1$, the martingale is bounded.

By the Martingale Convergence Theorem (Theorem A.3), there exists a random variable $\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(\infty)})$ such that:

$$|\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)}) - \mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(\infty)})| \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty. \tag{32}$$

■

C PROOF OF LEMMA 3.5

Proof. Consider a function $f(x_i, x_{-i}, \theta_i, \theta_{-i}) = -\mathop{\mathbb{E}}_{Z_i \sim \mathcal{D}_i(Y)} \ell_i(Z_i, x_i, x_{-i}, \theta_i, \theta_{-i})$ over a fixed distribution $\mathcal{D}_i(Y)$ induced by a joint decision $Y \in \mathcal{X}$. Here, Y is an arbitrary fixed joint decision denoted as X in Lemma 3.5, which induces the data distribution. In this proof, we use Y to avoid confusion with other decisions.

Since the loss function $\ell_i(\cdot)$ is continuous on the product space $\mathcal{X}_i \times \mathcal{X}_{-i} \times \Theta$, the function $f(x_i, x_{-i}, \theta_i, \theta_{-i})$ is also continuous on this product space.

Since $\mathcal{X}_i$ is a compact set, we define the decision set correspondence as $C(x_{-i}, \theta_i, \theta_{-i}) = \mathcal{X}_i$, indicating that the decision space $\mathcal{X}_i$ remains the same for all x_{-i}, θ_i and θ_{-i}. Then, we define the maximum of the function $f(x_i, x_{-i}, \theta_i, \theta_{-i})$ as $f^*(x_{-i}, \theta_i, \theta_{-i}) = \max\{f(x_i, x_{-i}, \theta_i, \theta_{-i}) : x_i \in \mathcal{X}_i\}$, which is achieved by the best response of player i with type θ_i to any decisions of others x_{-i}. The maximizer function $x_i^*(\theta_i) : \Theta_i \rightrightarrows \mathcal{X}_i$ is defined as $x_i^*(\theta_i) = \arg\max\{f(x_i, x_{-i}, \theta_i, \theta_{-i}) : x_i \in \mathcal{X}_i\} = \{x_i \in \mathcal{X}_i : f(x_i, x_{-i}, \theta_i, \theta_{-i}) = f^*(x_{-i}, \theta_i, \theta_{-i})\}$ for any fixed x_{-i} and θ_{-i}. Since the game is strongly monotone, the maximizer $x_i^*(\theta_i)$ (best response) exists and is unique for all types θ_i, others' decisions x_{-i} and others' types θ_{-i}.

Note that the maximization is pointwise over $\mathcal{X}_i$ for each fixed type θ_i. By Berge's Maximum Theorem (Theorem A.4), we can conclude that player i's best response $x_i^*(\theta_i)$ is continuous with respect to type θ_i for all $i \in [n]$, for any fixed x_{-i} and θ_{-i}.

Consider the loss function $\ell_i\left(Z_i, x_i, x_{-i}^*, \theta_i, \theta_{-i}\right)$, where x_{-i}^* is the Bayesian equilibrium decision vector of other players with types θ_{-i}. Since the loss function $\ell_i(\cdot)$ is continuous with respect to x_{-i}^*, x_{-i}^* is continuous with respect to $\theta_{-i} \in \Theta_{-i}$, and Θ_{-i} is a compact space, each loss function $\ell_i(\cdot)$ is continuous with respect to θ_{-i} and bounded on Θ_{-i}, for all $i \in [n]$.

According to Lemma 3.4, for every measurable set $\mathcal{A} \subseteq \Theta$, we have $|\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)}) - \mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(\infty)})| \xrightarrow{\text{a.s.}} 0$ as $t \to \infty$. The belief sequence $\{\mu(\theta_{-i}|h^{(t)})\}$ is a bounded sequence of non-negative probability measures on Θ_{-i} and almost surely converges to $\mu^{(\infty)}(\theta_{-i})$ as $t \to \infty$, which entails the weak convergence of the beliefs. According to Portmanteau Theorem (Theorem A.5),

since the function $\mathop{\mathbb{E}}_{Z_i\sim\mathcal{D}_i(Y)} \ell_i\left(Z_i, x_i, x^*_{-i}, \theta_i, \theta_{-i}\right)$ is continuous and bounded on the compact set Θ_{-i}, the following holds:

$$\left| \mathop{\mathbb{E}}_{\substack{Z_i\sim\mathcal{D}_i(Y)\\ \theta_{-i}\sim\mu(\theta_{-i}|h^{(t)})}} \ell_i\left(Z_i, x_i, x^*_{-i}, \theta_i, \theta_{-i}\right) - \mathop{\mathbb{E}}_{\substack{Z_i\sim\mathcal{D}_i(Y)\\ \theta_{-i}\sim\mu(\theta_{-i}|h^{(\infty)})}} \ell_i\left(Z_i, x_i, x^*_{-i}, \theta_i, \theta_{-i}\right) \right| \xrightarrow{\text{a.s.}} 0 \text{ as } t\to\infty. \tag{33}$$

Now, we denote $g_i^{(t)}(x_i) = \mathop{\mathbb{E}}_{\substack{Z_i\sim\mathcal{D}_i(Y)\\ \theta_{-i}\sim\mu(\theta_{-i}|h^{(t)})}} \ell_i\left(Z_i, x_i, x^*_{-i}, \theta_i, \theta_{-i}\right)$ and $x_i^{*(t)} = \arg\min_{x_i} g_i^{(t)}(x_i)$. Since the game $\mathcal{G}_\mu(Y)$ is α-strongly monotone, we consider two joint decisions X, X' that differ only in player i's decision in Eq. (11), i.e., $\|X - X'\|_2 = \|x_i - x_i'\|_2$. In this case, we have:

$$X - X' = (x_1 - x_1, \cdots, x_i - x_i', \cdots, x_n - x_n)^T = (\mathbf{0}, x_i - x_i', \mathbf{0})^T. \tag{34}$$

Then, the inner product on the right-hand side of Eq. (11) is:

$$\left(\nabla_X \mathcal{G}_\mu(Y) - \nabla_{X'}\mathcal{G}_\mu(Y)\right)^T (X - X') = \left(\nabla_{x_i} g_i^{(t)}(x_i) - \nabla_{x_i} g_i^{(t)}(x_i')\right)^T (x_i - x_i'). \tag{35}$$

Therefore, the α-strong monotonicity condition holds for each loss function gradient with respect to x_i for $i \in [n]$:

$$\alpha\|x_i - x_i'\|_2^2 \le \left(\nabla_{x_i} g_i^{(t)}(x_i) - \nabla_{x_i} g_i^{(t)}(x_i')\right)^T (x_i - x_i'). \tag{36}$$

Since a function that is α-strongly monotone in gradient is α-strongly convex (Bauschke & Combettes, 2020), we have:

$$\begin{aligned} g_i^{(\infty)}(x_i^{*(t)}) &\ge g_i^{(\infty)}(x_i^{*(\infty)}) + \frac{\alpha}{2}\|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2 + \nabla_{x_i} g_i^{(\infty)}(x_i^{*(\infty)})^T(x_i^{*(t)} - x_i^{*(\infty)}),\\ g_i^{(t)}(x_i^{*(\infty)}) &\ge g_i^{(t)}(x_i^{*(t)}) + \frac{\alpha}{2}\|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2 + \nabla_{x_i} g_i^{(t)}(x_i^{*(t)})^T(x_i^{*(\infty)} - x_i^{*(t)}). \end{aligned} \tag{37}$$

Since $x_i^{*(\infty)}$ minimizes $g_i^{(\infty)}(x_i)$ and $x_i^{*(t)}$ minimizes $g_i^{(t)}(x_i)$, by the first-order optimality conditions (Bubeck et al., 2015), we have $\nabla_{x_i} g_i^{(\infty)}(x_i^{*(\infty)})^T(x_i^{*(t)} - x_i^{*(\infty)}) \ge 0$ and $\nabla_{x_i} g_i^{(t)}(x_i^{*(t)})^T(x_i^{*(\infty)} - x_i^{*(t)}) \ge 0$. Combining the two inequalities in Eq. (37), we have:

$$\begin{aligned} &g_i^{(\infty)}(x_i^{*(t)}) - g_i^{(t)}(x_i^{*(t)}) + g_i^{(t)}(x_i^{*(\infty)}) - g_i^{(\infty)}(x_i^{*(\infty)})\\ \ge&\alpha\|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2 + \nabla_{x_i} g_i^{(\infty)}(x_i^{*(\infty)})^T(x_i^{*(t)} - x_i^{*(\infty)}) + \nabla_{x_i} g_i^{(t)}(x_i^{*(t)})^T(x_i^{*(\infty)} - x_i^{*(t)})\\ \ge&\alpha\|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2. \end{aligned} \tag{38}$$

Since the loss function $\ell_i(\cdot)$ is continuous on the compact product space $\mathcal{X}_i \times \Theta_{-i}$, the loss function is uniformly continuous on $\mathcal{X}_i \times \Theta_{-i}$ (Theorem A.6). Since Θ_{-i} is a compact space, for any $\eta > 0$, there exists a $\delta > 0$ that guarantee the uniform continuity of loss function $\ell_i(\cdot)$ for any $\theta_{-i} \in \Theta_{-i}$ (Definition A.7). Since $g_i^{(t)}(x_i)$ is continuous on the compact space $\mathcal{X}_i$, this function is uniformly continuous on $\mathcal{X}_i$. Since $g_i^{(t)}(x_i)$ is the expectation of loss function $\ell_i(\cdot)$ over $\theta_{-i} \in \Theta_{-i}$, the uniform continuity of $g_i^{(t)}(x_i)$ holds for any belief $\mu(\theta_{-i}|h^{(t)})$ and is independent of t. Thus, we have that the sequence $\{g_i^{(t)}(x_i)\}$ is equicontinuous (Definition A.8).

According to Eq. (33), we have $|g_i^{(t)}(x_i) - g_i^{(\infty)}(x_i)| \xrightarrow{\text{a.s.}} 0$ as $t \to \infty$, for any x_i, which is pointwise convergence.

Since $\{g_i^{(t)}(x_i)\}$ is an equicontinuous sequence of functions on a compact set $\mathcal{X}_i$, and $\{g_i^{(t)}(x_i)\}$ converges pointwise on $\mathcal{X}_i$, we have $\{g_i^{(t)}(x_i)\}$ converges uniformly on $\mathcal{X}_i$ (Lemma A.9). Thus, we have $|g_i^{(\infty)}(x_i^{*(t)}) - g_i^{(t)}(x_i^{*(t)})| \xrightarrow{\text{a.s.}} 0$ and $|g_i^{(t)}(x_i^{*(\infty)}) - g_i^{(\infty)}(x_i^{*(\infty)})| \xrightarrow{\text{a.s.}} 0$ as $t \to \infty$.

Given the strong monotonicity and loss function continuity hold for all $i \in [n]$, we have:

$$\begin{aligned} & g_i^{(\infty)}(x_i^{*(t)}) - g_i^{(t)}(x_i^{*(t)}) + g_i^{(t)}(x_i^{*(\infty)}) - g_i^{(\infty)}(x_i^{*(\infty)}) \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty, \\ \Longrightarrow & \alpha \|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2 \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty, \\ \Longrightarrow & \|x_i^{*(t)} - x_i^{*(\infty)}\|_2 \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty. \end{aligned} \tag{39}$$

Thus, for $\|BE_{\mu^{(t)}}(Y) - BE_{\mu^{(\infty)}}(Y)\|_2$, we have:

$$\begin{aligned} & \|BE_{\mu^{(t)}}(Y) - BE_{\mu^{(\infty)}}(Y)\|_2^2 = \sum_{i=1}^{n} \|x_i^{*(t)} - x_i^{*(\infty)}\|_2^2 \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty, \\ \Longrightarrow & \|BE_{\mu^{(t)}}(Y) - BE_{\mu^{(\infty)}}(Y)\|_2 \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty, \end{aligned} \tag{40}$$

where Y is the fixed joint decision denoted X in Lemma 3.5, and $\mu^{(t)} = \mu(\boldsymbol{\theta}|h^{(t)})$. ■

D Proof of Lemma 3.6

Proof. Let $Y = BE_\mu(X)$ and $Y' = BE_\mu(X')$ in this proof. Given an α-strongly monotone game $\mathcal{G}_\mu(\cdot)$ (as defined in Eq. (11)) and the definition of $BE_\mu(X)$ in Eq. (5), we have:

$$\alpha \|Y - Y'\|_2^2 \leq \left(\nabla_Y \mathcal{G}_\mu(X) - \nabla_{Y'} \mathcal{G}_\mu(X)\right)^T (Y - Y'). \tag{41}$$

By the first-order optimality conditions (Bubeck et al., 2015), we have:

$$\begin{aligned} \nabla_Y \mathcal{G}_\mu(X)^T (Y - Y') \leq 0, \\ \nabla_{Y'} \mathcal{G}_\mu(X')^T (Y - Y') \geq 0, \end{aligned} \tag{42}$$

and thus we have:

$$\nabla_Y \mathcal{G}_\mu(X)^T (Y - Y') \leq \nabla_{Y'} \mathcal{G}_\mu(X')^T (Y - Y'). \tag{43}$$

Combining Eqs. (41) and (43) and applying the Cauchy-Schwarz inequality (Hunter & Nachtergaele, 2001), we obtain:

$$\begin{aligned} & \alpha \|Y - Y'\|_2^2 \\ \leq & \left(\nabla_Y \mathcal{G}_\mu(X) - \nabla_{Y'} \mathcal{G}_\mu(X)\right)^T (Y - Y') \\ \leq & \left(\nabla_{Y'} \mathcal{G}_\mu(X') - \nabla_{Y'} \mathcal{G}_\mu(X)\right)^T (Y - Y') \\ \leq & \|\nabla_{Y'} \mathcal{G}_\mu(X') - \nabla_{Y'} \mathcal{G}_\mu(X)\|_2 \|Y - Y'\|_2 \\ \Longrightarrow & \alpha \|Y - Y'\|_2 \leq \|\nabla_{Y'} \mathcal{G}_\mu(X') - \nabla_{Y'} \mathcal{G}_\mu(X)\|_2 \end{aligned} \tag{44}$$

From the definition of individual gradients (Definition 3.2) and the $\hat{\varepsilon}_i$-sensitivity (Definition 3.1), we have:

$$\begin{aligned} & \|\nabla_{Y'} \mathcal{G}_\mu(X') - \nabla_{Y'} \mathcal{G}_\mu(X)\|_2^2 \\ = & \sum_{i=1}^{n} \left\| \mathop{\mathbb{E}}_{\substack{Z_i' \sim \mathcal{D}_i(X') \\ \theta_{-i} \sim \mu(\theta_{-i})}} \nabla_{y_i'} \ell_i \left(Z_i', y_i', y_{-i}', \theta_i, \theta_{-i}\right) - \mathop{\mathbb{E}}_{\substack{Z_i \sim \mathcal{D}_i(X) \\ \theta_{-i} \sim \mu(\theta_{-i})}} \nabla_{y_i'} \ell_i \left(Z_i, y_i', y_{-i}', \theta_i, \theta_{-i}\right) \right\|_2^2 \\ \leq & \sum_{i=1}^{n} \hat{\varepsilon}_i^2 \|X - X'\|_2^2 \\ \Longrightarrow & \|\nabla_{Y'} \mathcal{G}_\mu(X') - \nabla_{Y'} \mathcal{G}_\mu(X)\|_2 \leq \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2} \|X - X'\|_2. \end{aligned} \tag{45}$$

Combining Eqs. (44) and (45), we have:

$$\|BE_\mu(X) - BE_\mu(X')\|_2 = \|Y - Y'\|_2 \leq \frac{\sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}}{\alpha} \|X - X'\|_2. \tag{46}$$

■

E PROOF OF THEOREM 3.7

Proof. The first statement can be obtained directly from Lemma 3.4.

We prove the second statement as follows. According to the definition of repeated performative Bayesian game (Definition 2.1) and the definition of performatively stable Bayesian equilibrium (Definition 2.2), we have:

$$\begin{aligned} X^{(t+1)} &= BE_{\mu^{(t)}}(X^{(t)}), \\ X^{PS} &= BE_{\mu^{PS}}(X^{PS}). \end{aligned} \tag{47}$$

And we have:

$$\|X^{(t+1)} - X^{PS}\|_2 = \|BE_{\mu^{(t)}}(X^{(t)}) - BE_{\mu^{PS}}(X^{PS})\|_2 \tag{48}$$

With the triangle inequality, we have:

$$\begin{aligned} &\|BE_{\mu^{(t)}}(X^{(t)}) - BE_{\mu^{PS}}(X^{PS})\|_2 \\ \leq &\|BE_{\mu^{(t)}}(X^{(t)}) - BE_{\mu^{(t)}}(X^{PS})\|_2 + \|BE_{\mu^{(t)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \end{aligned} \tag{49}$$

From Lemma 3.5, we have:

$$\|BE_{\mu^{(t)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \xrightarrow{a.s.} 0 \text{ as } t \to \infty. \tag{50}$$

From Lemma 3.6, we have:

$$\|BE_{\mu^{(t)}}(X^{(t)}) - BE_{\mu^{(t)}}(X^{PS})\|_2 \leq \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\|X^{(t)} - X^{PS}\|_2. \tag{51}$$

As a result, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have:

$$\|X^{(t)} - X^{PS}\|_2 \leq \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\|X^{(t-1)} - X^{PS}\|_2 + \|BE_{\mu^{(t-1)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2. \tag{52}$$

We set $\beta = \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} < 1$, then we have:

$$\|X^{(t)} - X^{PS}\|_2 \leq \beta^t \|X^{(0)} - X^{PS}\|_2 + \sum_{s=0}^{t-1} \beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2. \tag{53}$$

Let $\eta > 0$. Since $\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \to 0$ almost surely as $s \to \infty$ by Lemma 3.5, there exists S such that $\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \leq \frac{\eta(1-\beta)}{2}$ for all $s \geq S$. Then we have:

$$\sum_{s=S}^{t-1} \beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \leq \frac{\eta(1-\beta)}{2}\sum_{s=S}^{t-1}\beta^{t-1-s} \leq \frac{\eta(1-\beta)}{2}\cdot\frac{1}{1-\beta} = \frac{\eta}{2}, \tag{54}$$

For $\beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2$ with $s < S$, the number of these terms is finite, and $\beta^{t-1-s} \to 0$ for each fixed s as $t \to \infty$. Thus, these terms tend to 0. In particular, for all sufficiently large t with $s < S$, $\sum_{s=0}^{S-1} \beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2$ is smaller than $\frac{\eta}{2}$. This implies that for all sufficiently large t,

$$\sum_{s=0}^{t-1} \beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \leq \frac{\eta}{2} + \frac{\eta}{2} = \eta. \tag{55}$$

Since $\eta > 0$ can be chosen arbitrarily small, we have $\sum_{s=0}^{t-1} \beta^{t-1-s}\|BE_{\mu^{(s)}}(X^{PS}) - BE_{\mu^{PS}}(X^{PS})\|_2 \to 0$ as $t \to \infty$ almost surely.

Meanwhile, since $\beta < 1$, we have $\beta^t \to 0$ and $\beta^t\|X^{(0)} - X^{PS}\|_2 \to 0$.

As a result, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we can conclude that:

$$\left\|X^{(t)} - X^{PS}\right\|_2 = \|BE_{\mu^{(t-1)}}(X^{(t-1)}) - BE_{\mu^{PS}}(X^{PS})\|_2 \xrightarrow{a.s.} 0 \text{ as } t \to \infty \tag{56}$$

■

F PROOF OF THEOREM 3.9

Proof. We first prove the first statement as follows. In the finite sample scenario, the only difference from the true distribution scenario lies in the data set. The Z_i in the history $h^{(t)}$ denotes the finite sample set from the induced data distribution, while the Bayesian update rule for beliefs remains identical to that in the true distribution scenario. The Bayesian updated beliefs remain a martingale. Thus, Lemma 3.4 still holds in the finite sample scenario. According to Lemma 3.4, the Bayesian updated beliefs under the finite sample setting converge almost surely to the limiting probability measure μ^{PS} on Θ. Specifically, for every measurable set $\mathcal{A} \subseteq \Theta$, we have:

$$|\mu(\boldsymbol{\theta} \in \mathcal{A}|h^{(t)}) - \mu^{PS}(\mathcal{A})| \xrightarrow{\text{a.s.}} 0 \text{ as } t \to \infty. \tag{57}$$

We prove the second statement as follows. By applying the triangle inequality, we have:

$$\begin{aligned} &\|\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{PS})\|_2 \\ \leq &\|\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})\|_2 \\ &+ \|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)}) - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{PS})\|_2, \end{aligned} \tag{58}$$

where $BE_{\mu^{(t)}}(X^{(t)})$ is the joint decision of Bayesian equilibrium given belief $\mu^{(t)}$ of the game $\mathcal{G}_{\mu^{(t)}}(X^{(t)})$ under the finite sample setting, $\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})}$ represents the average of individual gradients at the joint decision $BE_{\mu^{(t)}}(X^{(t)})$, $\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})$ is the expected gradient of the game at $BE_{\mu^{(t)}}(X^{(t)})$ with respect to the true data distributions at time step t, and $\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{PS})$ is the expected gradient of the game at $BE_{\mu^{(t)}}(X^{(t)})$ with respect to the stable distributions $\mathcal{D}_i(X^{PS}), i \in [n]$.

From Theorem 3.7, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we know that $X^{PS} \in \mathcal{X}_{RPBG}$ almost surely. As stated in Definition 3.1, Y, X, and X' can be any joint decisions in the set $\mathcal{X}_{RPBG}$. Therefore, by setting $Y = BE_{\mu^{(t)}}(X^{(t)})$, $X = X^{(t)}$, and $X' = X^{PS}$, we have almost surely:

$$\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)}) - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{PS})\|_2 \leq \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2. \tag{59}$$

According to Assumption 3.8, the following inequality holds:

$$\|\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})\|_2 \leq \sigma. \tag{60}$$

Combining Eqs. (59)-(60), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have almost surely:

$$\|\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{PS})\|_2 \leq \sigma + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2. \tag{61}$$

Since X^{PS} and $BE_{\mu^{(t)}}(X^{(t)})$ are the Bayesian equilibria of the game $\mathcal{G}_{\mu^{PS}}(X^{PS})$ and $\mathcal{G}_{\mu^{(t)}}(X^{(t)})$, respectively, and the game $\mathcal{G}_{\mu}(\cdot)$ is monotone, according to the first-order optimality condition (Bubeck et al., 2015), the following inequalities hold:

$$\begin{aligned} (X^{PS} - BE_{\mu^{(t)}}(X^{(t)}))^T\overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)}}(X^{(t)})} \geq 0, \\ (BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\nabla_{X^{PS}}\mathcal{G}_{\mu^{PS}}(X^{PS}) \geq 0. \end{aligned} \tag{62}$$

Combining the two inequalities in Eq. (62), we have:

$$\begin{aligned}
&(X^{PS} - BE_{\mu^{(t)}}(X^{(t)}))^T \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} + (BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \nabla_{X^{PS}} \mathcal{G}_{\mu^{PS}}(X^{PS}) \\
=&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{X^{PS}} \mathcal{G}_{\mu^{PS}}(X^{PS}) - \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} \right) \\
=&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{X^{PS}} \mathcal{G}_{\mu^{PS}}(X^{PS}) - \nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) \right) + \\
&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} \right) \\
\geq& 0.
\end{aligned} \tag{63}$$

The last inequality in Eq. (63) can be rewritten as follows:

$$\begin{aligned}
&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \nabla_{X^{PS}} \mathcal{G}_{\mu^{PS}}(X^{PS}) \right) \\
\leq&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} \right).
\end{aligned} \tag{64}$$

Since the game $\mathcal{G}_\mu(\cdot)$ is α-strongly monotone, we have:

$$\begin{aligned}
&\alpha \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2^2 \\
\leq&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T \left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \nabla_{X^{PS}} \mathcal{G}_{\mu^{PS}}(X^{PS}) \right).
\end{aligned} \tag{65}$$

Multiplying both sides of Eq. (61) by $\| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2$ and applying the Cauchy-Schwarz inequality (Hunter & Nachtergaele, 2001), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have almost surely:

$$\begin{aligned}
&\left(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \right)^T \left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} \right) \\
\leq& \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2 \| \nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{PS}) - \overline{\nabla_{BE_{\mu^{(t)}}(X^{(t)})} \mathcal{G}_{\mu^{(t)}}(X^{(t)})} \|_2 \\
\leq& \left(\sigma + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2} \| X^{(t)} - X^{PS} \|_2 \right) \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2,
\end{aligned} \tag{66}$$

Combining Eqs. (64)-(66), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have almost surely:

$$\begin{aligned}
&\alpha \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2^2 \leq \left(\sigma + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2} \| X^{(t)} - X^{PS} \|_2 \right) \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2 \\
\Longrightarrow& \| BE_{\mu^{(t)}}(X^{(t)}) - X^{PS} \|_2 \leq \frac{\sigma}{\alpha} + \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \| X^{(t)} - X^{PS} \|_2.
\end{aligned} \tag{67}$$

Since $X^{(t+1)} = BE_{\mu^{(t)}}(X^{(t)})$, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have almost surely:

$$\begin{aligned}
\| X^{(t+1)} - X^{PS} \|_2 &\leq \frac{\sigma}{\alpha} + \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \| X^{(t)} - X^{PS} \|_2 \\
&= \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \right)^t \| X^{(1)} - X^{PS} \|_2 + \frac{\sigma}{\alpha} \sum_{k=0}^{t-1} \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \right)^k.
\end{aligned} \tag{68}$$

The sum of the geometric series is:

$$\sum_{k=0}^{t-1} \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \right)^k = \frac{1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \right)^t}{1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha} \right)}. \tag{69}$$

Since $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have the following as $t \to \infty$:

$$\left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t \|X^{(1)} - X^{PS}\|_2 \to 0,$$
$$\frac{\sigma\left(1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t\right)}{\alpha\left(1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)\right)} \to \frac{\sigma}{\alpha - \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}. \tag{70}$$

Thus, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have almost surely:

$$\|X^{(t)} - X^{PS}\|_2 \leq \frac{\sigma}{\alpha - \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}} \text{ as } t \to \infty. \tag{71}$$

■

G VARIANT OF THEOREM 3.9

We provide a variant of Theorem 3.9 by the Central Limit Theorem. We first quantify the discrepancy between empirical and expected individual gradients based on the Central Limit Theorem in Lemma G.1. Leveraging this lemma, we derive the finite-sample convergence condition for a performatively stable Bayesian equilibrium in Theorem G.2.

G.1 DISCREPANCY BETWEEN EMPIRICAL AND EXPECTED INDIVIDUAL GRADIENTS VIA CENTRAL LIMIT THEOREM

Lemma G.1. *Consider the empirical average individual gradients $\overline{\nabla_Y \mathcal{G}_{\mu,m}(X)}$ estimated over m samples of a game $\mathcal{G}_\mu(X)$:*

$$\overline{\nabla_Y \mathcal{G}_{\mu,m}(X)} := \left(\frac{1}{m}\sum_{j=1}^m \mathop{\mathbb{E}}_{\theta_{-1}\sim\mu(\theta_{-1}|h)} \nabla_{y_1} \ell_1(z_{1_j}, y_1, y_{-1}, \theta_1, \theta_{-1}), \cdots, \frac{1}{m}\sum_{j=1}^m \mathop{\mathbb{E}}_{\theta_{-n}\sim\mu(\theta_{-n}|h)} \nabla_{y_n} \ell_n(z_{n_j}, y_n, y_{-n}, \theta_n, \theta_{-n})\right)^T. \tag{72}$$
$$z_{1_j}\sim\mathcal{D}_1(X) \qquad\qquad z_{n_j}\sim\mathcal{D}_n(X)$$

Then for any $\delta > 0$, with probability at least $F_{\chi^2(d)}(\frac{m\delta^2}{\sigma'})$, the following inequality holds:

$$\left\|\overline{\nabla_Y \mathcal{G}_{\mu,m}(X)} - \nabla_Y \mathcal{G}_\mu(X)\right\|_2 \leq \delta, \tag{73}$$

where $\nabla_Y \mathcal{G}_\mu(X)$ is the expected individual gradients, $F_{\chi^2(d)}(\cdot)$ denotes the cumulative distribution function of the χ^2 distribution with d degrees of freedom, and $\sigma' = \lambda_{\max}(\mathbf{\Sigma})$ denotes the largest eigenvalue of $\mathbf{\Sigma}$, the covariance matrix of the single-sample individual gradients.

Proof. We decompose the loss function of player i into two components:

$$\ell_i(Z_i, y_i, y_{-i}, \theta_i, \theta_{-i}) = \ell_{i1}(Z_i, y_i, y_{-i}, \theta_i, \theta_{-i}) + \ell_{i2}(y_i, y_{-i}, \theta_i, \theta_{-i}), \tag{74}$$

where $\ell_{i1}(Z_i, y_i, y_{-i}, \theta_i, \theta_{-i})$ represents the Z_i-dependent component, and $\ell_{i2}(y_i, y_{-i}, \theta_i, \theta_{-i})$ represents the Z_i-independent component. We then define the following gradients:

- $\nabla_{y_i}\ell_{i\ z_{i_j}\sim\mathcal{D}_i} = \nabla_{y_i}\ell_{i1}(z_{i_j}, y_i, y_{-i}, \theta_i, \theta_{-i}) + \nabla_{y_i}\ell_{i2}(y_i, y_{-i}, \theta_i, \theta_{-i})$: the gradient of the loss function for a single sample z_{i_j} from data distribution $\mathcal{D}_i$.
- $\nabla_{y_i}\ell_{i\ Z_{i_m}\sim\mathcal{D}_i} = \frac{1}{m}\sum_{z_{i_j}\in Z_{i_m}} \nabla_{y_i}\ell_{i1}(z_{i_j}, y_i, y_{-i}, \theta_i, \theta_{-i}) + \nabla_{y_i}\ell_{i2}(y_i, y_{-i}, \theta_i, \theta_{-i})$: the average gradient of the loss function over a set of m i.i.d. samples from data distribution $\mathcal{D}_i$, where Z_{i_m} is the sample set containing these m samples, and z_{i_j} is the j^{th} sample in set Z_{i_m}.
- $\nabla_{y_i}\ell_{i\ \mathcal{D}_i} = \mathop{\mathbb{E}}_{Z_i\sim\mathcal{D}_i} \nabla_{y_i}\ell_i(Z_i, y_i, y_{-i}, \theta_i, \theta_{-i})$: the expected gradient of the loss function over the underlying, ground-truth distribution $\mathcal{D}_i$.

- $\nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}} = \left(\nabla_{y_1} \ell_{1\ z_{1_j} \sim \mathcal{D}_1}, \cdots, \nabla_{y_n} \ell_{n\ z_{n_j} \sim \mathcal{D}_n} \right)^T$: the individual gradients of the game for samples z_{i_j} from data distribution $\mathcal{D}_i, i \in [n]$, where each player i obtains a sample z_{i_j}.

- $\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} = \left(\nabla_{y_1} \ell_{1\ Z_{1_m} \sim \mathcal{D}_1}, \cdots, \nabla_{y_n} \ell_{n\ Z_{n_m} \sim \mathcal{D}_n} \right)^T$: the average individual gradients of the game over a set of m i.i.d. samples from data distribution $\mathcal{D}_i, i \in [n]$, where each player i obtains m samples.

- $\nabla_Y \mathcal{G}_{\mu, \mathcal{D}} = \left(\nabla_{y_1} \ell_{1\ \mathcal{D}_1}, \cdots, \nabla_{y_n} \ell_{n\ \mathcal{D}_n} \right)^T$: the expected individual gradients of the game over the underlying, ground-truth distribution $\mathcal{D}_i, i \in [n]$.

Then, the covariance matrix of $\nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}}, j = 1, 2, \cdots, m$ is given by:

$$\begin{aligned} \mathbf{\Sigma} &= \mathbf{Cov} \left[\nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}}, \nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}} \right] \\ &= \mathbb{E} \left[\left(\nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \left(\nabla_Y \mathcal{G}_{\mu, z_{\bullet_j} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)^T \right], \end{aligned} \tag{75}$$

where $\mathbf{\Sigma}$ is the covariance matrix.

According to the Central Limit Theorem (Bauer, 2001), we have:

$$\sqrt{m} \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \sim \mathcal{N}_d(\mathbf{0}, \mathbf{\Sigma}), \tag{76}$$

where $\mathcal{N}_d(\mathbf{0}, \mathbf{\Sigma})$ denotes the d-dimensional zero-mean Gaussian distribution with covariance matrix $\mathbf{\Sigma}$. The eigendecomposition of $\mathbf{\Sigma}$ yields:

$$\mathbf{\Sigma} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T, \tag{77}$$

where $\mathbf{U}$ is the $d \times d$ matrix whose j^{th} column is the j^{th} eigenvector $\mathbf{u}_j$ of $\mathbf{\Sigma}$, and $\mathbf{\Lambda}$ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues of $\mathbf{\Sigma}$: $\mathbf{\Lambda}_{jj} = \lambda_j$.

Let $\mathbf{a} = \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{U}^T \sqrt{m} \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)$, we have:

$$\begin{aligned} & \mathbf{a} \sim \mathcal{N}_d(\mathbf{0}, \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{U}^T \mathbf{\Sigma} \mathbf{U} \mathbf{\Lambda}^{-\frac{1}{2}}) \\ \implies & \mathbf{a} \sim \mathcal{N}_d(\mathbf{0}, \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{U}^T \left(\mathbf{U} \mathbf{\Lambda} \mathbf{U}^T \right) \mathbf{U} \mathbf{\Lambda}^{-\frac{1}{2}}) \\ \implies & \mathbf{a} \sim \mathcal{N}_d(\mathbf{0}, \mathbf{I}_d), \end{aligned} \tag{78}$$

where $\mathbf{I}_d$ denotes the $d \times d$ identity matrix, and the vector $\mathbf{a}$ follows the standard d-dimensional multivariate normal distribution $\mathcal{N}_d(\mathbf{0}, \mathbf{I}_d)$. Accordingly, the quadratic form $\mathbf{a}^T \mathbf{a}$ follows a χ^2 distribution with d degrees of freedom (Pearson, 1893):

$$\begin{aligned} \mathbf{a}^T \mathbf{a} &= m \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)^T \mathbf{U} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{U}^T \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \\ &= m \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)^T \mathbf{\Sigma}^{-1} \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \\ &\sim \chi^2(d). \end{aligned} \tag{79}$$

Then, for any positive real number τ, we have:

$$P \left[m \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)^T \mathbf{\Sigma}^{-1} \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \leq \tau \right] = F_{\chi^2(d)}(\tau), \tag{80}$$

where $F_{\chi^2(d)}(\tau)$ denotes the cumulative distribution function of the χ^2 distribution with d degrees of freedom, evaluated at τ.

For any real symmetric matrix $\mathbf{B} \in \mathbb{R}^{d \times d}$ and vector $\mathbf{v} \in \mathbb{R}^d$, the inequality $\mathbf{v}^T \mathbf{B} \mathbf{v} \geq \lambda_{min}(\mathbf{B}) \|\mathbf{v}\|_2^2$ holds, where $\lambda_{min}(\mathbf{B})$ denotes the smallest eigenvalue of $\mathbf{B}$. Therefore, we have:

$$\begin{aligned} & m \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right)^T \mathbf{\Sigma}^{-1} \left(\nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \right) \\ \geq & m \lambda_{min}(\mathbf{\Sigma}^{-1}) \| \nabla_Y \mathcal{G}_{\mu, Z_{\bullet_m} \sim \mathcal{D}} - \nabla_Y \mathcal{G}_{\mu, \mathcal{D}} \|_2^2. \end{aligned} \tag{81}$$

Using this inequality, we can derive the following:

$$\begin{aligned}
&P\left[m\lambda_{min}(\mathbf{\Sigma}^{-1})\|\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\|_2^2\le\tau\right]\\
=&P\left[\|\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\|_2^2\le\frac{\tau\lambda_{max}(\mathbf{\Sigma})}{m}\right]\\
=&P\left[\|\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\|_2\le\sqrt{\frac{\tau\lambda_{max}(\mathbf{\Sigma})}{m}}\right]\\
\ge&P\left[m\left(\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\right)^T\mathbf{\Sigma}^{-1}\left(\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\right)\le\tau\right]\\
=&F_{\chi^2(d)}(\tau).
\end{aligned}\tag{82}$$

Let $\delta=\sqrt{\frac{\tau\lambda_{max}(\mathbf{\Sigma})}{m}}$ and $\sigma'=\lambda_{max}(\mathbf{\Sigma})$, then $\tau=\frac{m\delta^2}{\sigma'}$. When each player i samples m samples, we can state that the following inequality holds with probability at least $F_{\chi^2(d)}(\frac{m\delta^2}{\sigma'})$:

$$\|\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}-\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}\|_2\le\delta.\tag{83}$$

Consider the game on the data distributions induced by the joint decision X, we can set $\nabla_Y\mathcal{G}_{\mu,Z_{\bullet_m}\sim\mathcal{D}}=\overline{\nabla_Y\mathcal{G}_{\mu,m}(X)}$ and $\nabla_Y\mathcal{G}_{\mu,\mathcal{D}}=\nabla_Y\mathcal{G}_{\mu}(X)$ to complete the proof. ■

G.2 Finite Sample Analysis with Central Limit Theorem

Theorem G.2. *Suppose Assumption 3.3 holds, the loss functions $\ell_i(\cdot)$ for all $i\in[n]$ are continuous on the product space $\mathcal{X}_i\times\mathcal{X}_{-i}\times\Theta$, data distribution maps $\mathcal{D}_i(X)$ for all $i\in[n]$ are $\hat{\varepsilon}_i$-sensitive, and $\alpha>\sqrt{\sum_{i=1}^n\hat{\varepsilon}_i^2}$. Then we have the following:*

1. *The Bayesian updated beliefs with finite samples converge almost surely to a limiting probability measure μ^{PS} on Θ. Specifically, for every measurable set $\mathcal{A}\subseteq\Theta$, we have:*

$$|\mu^{(t)}(\mathcal{A})-\mu^{PS}(\mathcal{A})|\xrightarrow{a.s.}0\text{ ,as }t\to\infty,\tag{84}$$

   where $\mu^{(t)}(\mathcal{A})$ is as defined in Lemma 3.4, and μ^{PS} is as defined in Definition 2.2.

2. *The sequence of joint decisions $\{X^{(0)},X^{(1)},\cdots,X^{(t)},\cdots\}\subseteq\mathcal{X}_{RPBG}$ with finite samples satisfies the following with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$:*

$$\|X^{(t)}-X^{PS}\|_2\le\frac{\delta}{\alpha-\sqrt{\sum_{i=1}^n\hat{\varepsilon}_i^2}}\text{ as }t\to\infty,\tag{85}$$

   where X^{PS} is as defined in Definition 2.2, δ is the parameter determining the convergence radius, m is the collected sample size in each time step, and $\sigma'=\lambda_{\max}(\mathbf{\Sigma})$ denotes the largest eigenvalue of $\mathbf{\Sigma}$.

Proof. We first prove the first statement as follows. In the finite sample scenario, the only difference from the true distribution scenario lies in the data set. The Z_i in the history $h^{(t)}$ denotes the finite sample set from the induced data distribution, while the Bayesian update rule for beliefs remains identical to that in the true distribution scenario. The Bayesian updated beliefs remain a martingale. Thus, Lemma 3.4 still holds in the finite sample scenario. According to Lemma 3.4, the Bayesian updated beliefs under the finite sample setting converge almost surely to the limiting probability measure μ^{PS} on Θ. Specifically, for every measurable set $\mathcal{A}\subseteq\Theta$, we have:

$$|\mu(\boldsymbol{\theta}\in\mathcal{A}|h^{(t)})-\mu^{PS}(\mathcal{A})|\xrightarrow{\text{a.s.}}0\text{ as }t\to\infty.\tag{86}$$

We prove the second statement as follows. By applying the triangle inequality, we have:

$$\begin{aligned}
&\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet_m}\sim\mathcal{D}(X^{(t)})}-\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}\|_2\\
\le&\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet_m}\sim\mathcal{D}(X^{(t)})}-\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{(t)})}\|_2\\
&+\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{(t)})}-\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}\|_2,
\end{aligned}\tag{87}$$

where m is the number of samples each player collects at each time step, $BE_{\mu^{(t)}}(X^{(t)})$ is the joint decision of Bayesian equilibrium given belief $\mu^{(t)}$ of the game $\mathcal{G}_{\mu^{(t)}}(X^{(t)})$ under the finite sample setting, $Z_{\bullet_m} \sim \mathcal{D}(X^{(t)})$ denotes the set of m samples drawn from the data distribution $\mathcal{D}_i(X^{(t)})$ for all $i \in [n]$, $\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}$ represents the average of individual gradients at the joint decision $BE_{\mu^{(t)}}(X^{(t)})$, computed over m collected samples, $\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{(t)})}$ is the expected gradient of the game at $BE_{\mu^{(t)}}(X^{(t)})$ with respect to the true data distribution at time step t, and $\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}$ is the expected gradient of the game at $BE_{\mu^{(t)}}(X^{(t)})$ with respect to the stable distribution $\mathcal{D}_i(X^{PS}), i \in [n]$.

From Theorem 3.7, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we know that $X^{PS} \in \mathcal{X}_{RPBG}$ almost surely. As stated in Definition 3.1, Y, X, and X' can be any joint decisions in the set $\mathcal{X}_{RPBG}$. Therefore, by setting $Y = BE_{\mu^{(t)}}(X^{(t)})$, $X = X^{(t)}$, and $X' = X^{PS}$, we have almost surely:

$$\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}\|_2 \leq \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2. \quad (88)$$

According to Lemma G.1, the following inequality holds with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$:

$$\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{(t)})}\|_2 \leq \delta. \quad (89)$$

Combining Eqs. (88)-(89), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have:

$$\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}\|_2 \leq \delta + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2, \quad (90)$$

with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$.

Since X^{PS} and $BE_{\mu^{(t)}}(X^{(t)})$ are the Bayesian equilibria of the game $\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})}$ and $\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}$, respectively, and the game $\mathcal{G}_\mu(\cdot)$ is monotone, according to the first-order optimality condition (Bubeck et al., 2015), the following inequalities hold:

$$\begin{aligned}(X^{PS} - BE_{\mu^{(t)}}(X^{(t)}))^T\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})} \geq 0,\\ (BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})} \geq 0.\end{aligned} \quad (91)$$

Combining the two inequalities in Eq. (91), we have:

$$\begin{aligned}&(X^{PS} - BE_{\mu^{(t)}}(X^{(t)}))^T\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})} + (BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})}\\ =&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}\right)\\ =&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})}\right) +\\ &(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}\right)\\ \geq&0.\end{aligned} \quad (92)$$

The last inequality in Eq. (92) can be rewritten as follows:

$$\begin{aligned}&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})}\right)\\ \leq&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}\right).\end{aligned} \quad (93)$$

Since the game $\mathcal{G}_\mu(\cdot)$ is α-strongly monotone, we have:

$$\begin{aligned}&\alpha\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2^2\\ \leq&(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS})^T\left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{X^{PS}}\mathcal{G}_{\mu^{PS},\mathcal{D}(X^{PS})}\right).\end{aligned} \tag{94}$$

Multiplying both sides of Eq. (90) by $\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2$ and applying Cauchy-Schwarz inequality (Hunter & Nachtergaele, 2001), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have:

$$\begin{aligned}&\left(BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\right)^T\left(\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}\right)\\ \leq&\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2\|\nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},\mathcal{D}(X^{PS})} - \nabla_{BE_{\mu^{(t)}}(X^{(t)})}\mathcal{G}_{\mu^{(t)},Z_{\bullet m}\sim\mathcal{D}(X^{(t)})}\|_2\\ \leq&\left(\delta + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2\right)\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2,\end{aligned} \tag{95}$$

with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$.

Combining Eqs. (93)-(95), when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have:

$$\begin{aligned}&\alpha\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2^2 \leq \left(\delta + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2\right)\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2\\ \Longrightarrow&\|BE_{\mu^{(t)}}(X^{(t)}) - X^{PS}\|_2 \leq \frac{\delta + \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}\|X^{(t)} - X^{PS}\|_2}{\alpha},\end{aligned} \tag{96}$$

with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$.

Since $X^{(t+1)} = BE_{\mu^{(t)}}(X^{(t)})$, when $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have:

$$\begin{aligned}\|X^{(t+1)} - X^{PS}\|_2 &\leq \frac{\delta}{\alpha} + \frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\|X^{(t)} - X^{PS}\|_2\\ &= \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t\|X^{(1)} - X^{PS}\|_2 + \frac{\delta}{\alpha}\sum_{k=0}^{t-1}\left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^k,\end{aligned} \tag{97}$$

with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$.

The sum of the geometric series is:

$$\sum_{k=0}^{t-1}\left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^k = \frac{1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t}{1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)}. \tag{98}$$

Since $\alpha > \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, we have the following as $t \to \infty$:

$$\begin{aligned}&\left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t\|X^{(1)} - X^{PS}\|_2 \to 0,\\ &\frac{\delta\left(1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)^t\right)}{\alpha\left(1 - \left(\frac{\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}{\alpha}\right)\right)} \to \frac{\delta}{\alpha - \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}}.\end{aligned} \tag{99}$$

Thus, when $\alpha > \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$, we have:

$$\|X^{(t)} - X^{PS}\|_2 \leq \frac{\delta}{\alpha - \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}} \text{ as } t \to \infty, \tag{100}$$

with probability at least $F_{\chi^2(d)}\left(\frac{m\delta^2}{\sigma'}\right)$. ■

H Details of Algorithm Description

The algorithm leverages the convergence conditions derived in Theorem 3.9, which requires the game to be α-strongly monotone with $\alpha > \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$. To achieve this, our algorithm uses a designed parameter $c \in (1, +\infty)$ to set $\gamma = \max(0, c\sqrt{\sum_{k=1}^{n} \hat{\varepsilon}_k^2} - \psi)$, ensuring the convergence condition is satisfied. In each time step t, each player $i \in [n]$ first updates the regularization coefficient $\gamma = \max(0, c\sqrt{\sum_{k=1}^{n} \hat{\varepsilon}_k^2} - \psi)$ based on the input parameter $c \in (1, +\infty)$, its own latest $\hat{\varepsilon}_i$ estimates, and latest reported $\hat{\varepsilon}_j$ for $j \neq i$ from the previous iteration (line 5). Then, for the ψ-strongly monotone game $\mathcal{G} := (\mathcal{L}_1, \cdots, \mathcal{L}_n)$, each player i adds a uniform regularization term $\frac{\gamma}{2}\|x_i\|_2^2$ to its loss function, constructing an α-strongly monotone game $\hat{\mathcal{G}} := \left(\hat{\mathcal{L}}_1, \cdots, \hat{\mathcal{L}}_n\right)$ with $\alpha = \psi + \gamma$ (line 6). Following this, each player i calculates its decision $x_i^{(t)}$ in a Bayesian equilibrium of the game $\hat{\mathcal{G}}(X^{(t-1)})$ under the data distributions induced by $X^{(t-1)}$ (line 7). For the strongly monotone game $\hat{\mathcal{G}}(X^{(t-1)})$ with compact decision and type spaces, the computation of a Bayesian equilibrium is tractable using numerical methods. After deploying the joint decision $X^{(t)} = (x_1^{(t)}, x_2^{(t)}, \cdots, x_n^{(t)})^T$ (line 8), each player i collects samples $Z_i^{(t)}$ from the distributions induced by the joint decision $X^{(t)}$ (line 9). Each player i then estimates the sensitivity parameter $\hat{\varepsilon}_i$ from the induced data samples $Z_i^{(t)}$ (line 10) and reports the estimated $\hat{\varepsilon}_i$ (line 11). Subsequently, each player i updates their belief μ via the Bayes' rule (line 12). The procedure repeats until it converges to a performatively stable Bayesian equilibrium.

By doing so, our algorithm ensures that the game is α-strongly monotone with $\alpha > \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$, satisfying the convergence condition in Theorem 3.9. Therefore, our algorithm guarantees convergence to a performatively stable Bayesian equilibrium X^{PS} with finite samples.

The algorithm is designed based on the convergence conditions in Theorem 3.9 and is compatible with any ψ-strongly monotone game that guarantees the existence of a Bayesian equilibrium at each time step. The value of ψ may depend on the players' types, which can be unknown to players. Nevertheless, ψ is common knowledge in some classic incomplete information games, such as Cournot competition (Cournot et al., 1897) and Bertrand competition (Bertrand, 1883), enabling the direct implementation of our algorithm in these scenarios. For games with unknown ψ but with a separating equilibrium (e.g., Spence's job-market signaling model (Spence, 1978)), the true types of players are fully revealed once beliefs converge, and ψ can be verified through the belief update process, making our algorithm applicable.

The algorithm is a protocol that fully rational players follow. By Theorem 3.9, the repeated performative Bayesian game may not converge to a performatively stable Bayesian equilibrium when $\psi \leq \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$, leading to high losses for all players. Thus, under the common knowledge of the convergence conditions and rationality, all players will adopt the regularization, since it ensures convergence to a performatively stable Bayesian equilibrium.

In our algorithm, each $\hat{\varepsilon}_i$ is a running maximum, which is monotone non-decreasing and bounded over the compact set $\mathcal{X}_{RPBG}$, ensuring its convergence. Specifically, if the current estimate is smaller than the true $\hat{\varepsilon}_i$, the convergence condition is not satisfied, causing the decisions to fluctuate and the estimated $\hat{\varepsilon}_i$ to increase until the condition holds.

I Truthfulness of Reported $\hat{\varepsilon}_i$

By the definition of $\hat{\varepsilon}_i$-sensitivity (Definition 3.1), calculating $\hat{\varepsilon}_i$ requires the player i's private information θ_i, which may not be fully revealed through belief updates. However, selecting the appropriate γ-strongly convex regularizer depends on each $\hat{\varepsilon}_i$. To address this, our algorithm requires

each player to report their calculated $\hat{\varepsilon}_i$ during each iteration of the repeated performative Bayesian game.

The truthfulness of reported $\hat{\varepsilon}_i$ is crucial for convergence. We demonstrate that any rational player will report the true $\hat{\varepsilon}_i$. Let $\tilde{\varepsilon}_i$ denote player i's reported $\hat{\varepsilon}_i$. If $\hat{\varepsilon}_i > \tilde{\varepsilon}_i$, then $\sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2} > \sqrt{\sum_{i=1}^n \tilde{\varepsilon}_i^2}$. Consequently, the algorithm selects a $\tilde{\gamma}$-strongly convex regularizer such that the game becomes α'-strongly monotone with $\alpha' > \sqrt{\sum_{i=1}^n \tilde{\varepsilon}_i^2}$. If $\alpha' < \sqrt{\sum_{i=1}^n \hat{\varepsilon}_i^2}$, the convergence condition is violated, and convergence to a performatively stable Bayesian equilibrium cannot be guaranteed. In this case, all players cannot have stable outcomes.

Conversely, if $\tilde{\varepsilon}_i > \hat{\varepsilon}_i$, the algorithm uses a larger $\tilde{\gamma}$, potentially leading to a large bias from the original equilibrium. As shown in Theorem 4.1, the upper bound of $\|X^* - X^R\|_2$ increases with γ. Although the convergence is guaranteed, the resulting equilibrium may be suboptimal for all players due to the excessive regularization. Therefore, rational players do not over-report $\hat{\varepsilon}_i$.

Furthermore, a player can report a large $\hat{\varepsilon}_i$ to motivate other players to apply a stronger regularization while employing a weaker one itself. However, under the common knowledge of rationality, all players would have the same incentive to do so, leading to a large γ that increases losses for all players. In addition, if any player uses a γ different from the one provided by the algorithm, it would disrupt the belief updates, since the observed decisions do not align with the equilibrium decisions.

Consequently, truthfully reporting $\hat{\varepsilon}_i$ ensures convergence to an equilibrium and that no player has an incentive to deviate.

In summary, our algorithm ensures that rational players have no incentive to misreport $\hat{\varepsilon}_i$, since reporting a small $\hat{\varepsilon}_i$ may lead to non-convergence, while reporting a large $\hat{\varepsilon}_i$ may lead to suboptimal outcomes. Consequently, our proposed algorithm guarantees truthful reporting of $\hat{\varepsilon}_i$, thereby ensuring convergence to a performatively stable Bayesian equilibrium.

J Proof of Theorem 4.1

Proof. In this proof, we use Y to denote the fixed joint decision that induces the data distribution, while X^* and X^R denote the Bayesian equilibria of the original and regularized games, respectively.

For any γ-strongly convex regularization term $R(\cdot)$, the regularized game becomes:

$$\hat{\mathcal{G}}_\mu^R(Y) := \left(\begin{array}{c} \mathop{\mathbb{E}}_{\substack{Z_1 \sim \mathcal{D}_1(Y) \\ \theta_{-1} \sim \mu(\theta_{-1}|h)}} \left[\ell_1(Z_1, x_1, x_{-1}, \theta_1, \theta_{-1}) + R(x_1)\right], \cdots, \\ \mathop{\mathbb{E}}_{\substack{Z_n \sim \mathcal{D}_n(Y) \\ \theta_{-n} \sim \mu(\theta_{-n}|h)}} \left[\ell_n(Z_n, x_n, x_{-n}, \theta_n, \theta_{-n}) + R(x_n)\right] \end{array} \right). \tag{101}$$

Since the regularized game $\hat{\mathcal{G}}_\mu^R(Y)$ is $(\psi+\gamma)$-strongly monotone, the individual gradient $\nabla_X \hat{\mathcal{G}}_\mu^R(Y)$ satisfies:

$$(\nabla_{X^*} \hat{\mathcal{G}}_\mu^R(Y) - \nabla_{X^R} \hat{\mathcal{G}}_\mu^R(Y))^T (X^* - X^R) \geq (\psi + \gamma)\|X^* - X^R\|_2^2, \tag{102}$$

where X^* is the Bayesian equilibrium of the original game $\mathcal{G}_\mu(Y)$, and X^R is the Bayesian equilibrium of the regularized game $\hat{\mathcal{G}}_\mu^R(Y)$.

Applying the Cauchy-Schwarz inequality (Hunter & Nachtergaele, 2001) to Eq. (102), we obtain:

$$\begin{aligned} &\|\nabla_{X^*} \hat{\mathcal{G}}_\mu^R(Y) - \nabla_{X^R} \hat{\mathcal{G}}_\mu^R(Y)\|_2 \|X^* - X^R\|_2 \\ &\geq (\nabla_{X^*} \hat{\mathcal{G}}_\mu^R(Y) - \nabla_{X^R} \hat{\mathcal{G}}_\mu^R(Y))^T (X^* - X^R) \\ \geq &(\psi + \gamma)\|X^* - X^R\|_2^2, \\ \Longrightarrow &\|\nabla_{X^*} \hat{\mathcal{G}}_\mu^R(Y) - \nabla_{X^R} \hat{\mathcal{G}}_\mu^R(Y)\|_2 \geq (\psi + \gamma)\|X^* - X^R\|_2. \end{aligned} \tag{103}$$

By the first-order optimality condition (Bubeck et al., 2015), we have:

$$\begin{aligned} &\nabla_{X^*}\mathcal{G}_\mu(Y) = \mathbf{0} \Longrightarrow \nabla_{X^*}\hat{\mathcal{G}}_\mu^R(Y) - \nabla_{X^*}R(X^*) = \mathbf{0}, \\ &\nabla_{X^R}\hat{\mathcal{G}}_\mu^R(Y) = \mathbf{0}, \end{aligned} \tag{104}$$

where $\nabla_{X^*}R(X^*)$ is the gradient of the regularization term $R(\cdot)$ at X^*.

Substituting $\nabla_{X^*}\hat{\mathcal{G}}_\mu^R(Y)$ and $\nabla_{X^R}\hat{\mathcal{G}}_\mu^R(Y)$ into Eq. (103) yields:

$$\|\nabla_{X^*}R(X^*)\|_2 \geq (\psi+\gamma)\|X^* - X^R\|_2 \Longrightarrow \frac{\|\nabla_{X^*}R(X^*)\|_2}{(\psi+\gamma)} \geq \|X^* - X^R\|_2. \tag{105}$$

Considering the gradient of $R(\cdot)$ is L-Lipschitz continuous, we have:

$$\begin{aligned} &\|\nabla_X R(X) - \nabla_{X'}R(X')\|_2 \leq L\|X - X'\|_2, \\ \Longrightarrow &\|\nabla_X R(X)\|_2 = \|\nabla_X R(X) - \nabla_{\mathbf{0}}R(\mathbf{0}) + \nabla_{\mathbf{0}}R(\mathbf{0})\|_2 \\ &\leq \|\nabla_X R(X) - \nabla_{\mathbf{0}}R(\mathbf{0})\|_2 + \|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2 \leq L\|X\|_2 + \|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2. \end{aligned} \tag{106}$$

Combining Eqs. (105) and (106), we have:

$$\frac{L\|X^*\|_2 + \|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2}{(\psi+\gamma)} \geq \frac{\|\nabla_{X^*}R(X^*)\|_2}{(\psi+\gamma)} \geq \|X^* - X^R\|_2. \tag{107}$$

This indicates that the distance between the original and regularized equilibria, i.e., $\|X^* - X^R\|_2$, is upper bounded by $\frac{L\|X^*\|_2 + \|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2}{(\psi+\gamma)}$.

For γ-strongly convex $R(\cdot)$, by definition:

$$\begin{aligned} R(X) - R(X') &\geq \nabla_{X'}R(X')^T(X - X') + \frac{\gamma}{2}\|X - X'\|_2^2, \\ R(X') - R(X) &\geq \nabla_X R(X)^T(X' - X) + \frac{\gamma}{2}\|X - X'\|_2^2. \end{aligned} \tag{108}$$

Summing these two inequalities yields:

$$0 \geq (\nabla_{X'}R(X') - \nabla_X R(X))^T(X - X') + \gamma\|X - X'\|_2^2. \tag{109}$$

By the Cauchy-Schwarz inequality (Hunter & Nachtergaele, 2001) and L-Lipschitz continuity of the gradient:

$$\begin{aligned} &(\nabla_X R(X) - \nabla_{X'}R(X'))^T(X - X') \\ \leq &\|\nabla_{X'}R(X') - \nabla_X R(X)\|_2\|X - X'\|_2 \\ \leq &L\|X - X'\|_2^2, \\ \Longrightarrow &(\nabla_{X'}R(X') - \nabla_X R(X))^T(X - X') \geq -L\|X - X'\|_2^2. \end{aligned} \tag{110}$$

Combining Eqs. (109) and (110):

$$0 \geq -L\|X - X'\|_2^2 + \gamma\|X - X'\|_2^2 \Longrightarrow L \geq \gamma. \tag{111}$$

According to Eqs. (107) and (111), the upper bound of $\|X^* - X^R\|_2$, i.e., $\frac{L\|X^*\|_2 + \|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2}{(\psi+\gamma)}$, is minimized when: (1) L is minimized, i.e., $L = \gamma$, and (2) $\|\nabla_{\mathbf{0}}R(\mathbf{0})\|_2 = 0$. Clearly, the proposed quadratic regularizer $R(X) = \frac{\gamma}{2}\|X\|_2^2$ satisfies these two conditions, as it is γ-Lipschitz in gradient, and has $\nabla R(\mathbf{0}) = \mathbf{0}$, yielding the tightest upper bound $\frac{\gamma\|X^*\|_2}{(\psi+\gamma)}$. Thus, $R(X) = \frac{\gamma}{2}\|X\|_2^2$ minimizes the upper bound of the distance $\|X^* - X^R\|_2$ among all γ-strongly convex regularizers. Specifically, for $R(X) = \frac{\gamma}{2}\|X\|_2^2$, we have:

$$\|X^* - X^R\|_2 \leq \frac{\gamma\|X^*\|_2}{(\psi+\gamma)}. \tag{112}$$

■

K EXPERIMENTAL DETAILS

K.1 ADDITIONAL EXPERIMENTAL SETTINGS

Theorem 3.9 guarantees the convergence if $\alpha > \sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2}$. In our algorithm, this condition is satisfied by setting the parameter $c \in (1, +\infty)$ and applying γ-strongly convex regularization with $\gamma = \max(0, c\sqrt{\sum_{i=1}^{n} \hat{\varepsilon}_i^2} - \psi)$ for a ψ-strongly monotone game. To validate our theoretical analysis and the proposed algorithm, we set four values for the parameter c in Algorithm 1, where $c = 0$ and $c = 0.3$ violate the convergence conditions in Theorem 3.9, and $c = 1.1$ and 30 satisfy the convergence conditions in Theorem 3.9.

In addition, we include three baselines without using our algorithm to show the different dynamics in performative Bayesian games: (1) Static Bayesian Equilibrium (SBE), where players optimize decisions based on current data distributions and fixed prior beliefs without belief updates; (2) Complete Information Equilibrium (CIE), where players optimize decisions based on complete information and current data distributions; and (3) Heuristic Strategy (HS), where players fit linear data distribution maps (a heuristic rule from (Narang et al., 2023)) using historical decision-distribution pairs, and substitute the fitted maps into loss functions to compute the equilibrium decisions with updated beliefs.

At each time step t, we show the total revenue achieved by the joint decision $X^{(t)}$ on the induced distribution $\mathcal{D}_i(X^{(t)})$. For each setting, we run 15 time steps for each experiment and report the average results over 10 trials. We apply the same convergence thresholds $\rho_1 = \rho_2$ in our experiments, specified as $\rho_1 = \rho_2 = 0.1\|X^{(t)}\|_2$ for the Cournot competition and $\rho_1 = \rho_2 = 0.01$ for the Bertrand competition.

K.2 COMPARISON OF WALL-CLOCK TIME

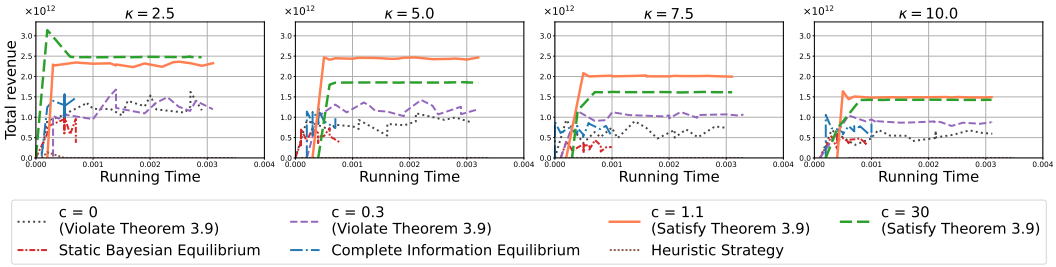


Figure 2: Total revenue for our algorithm with $c = 0, 0.3, 1.1, 30$ and three baselines in Cournot competition within crude oil trade with $\kappa = \{2.5, 5, 7.5, 10\}$ over 10 trials across wall-clock time.

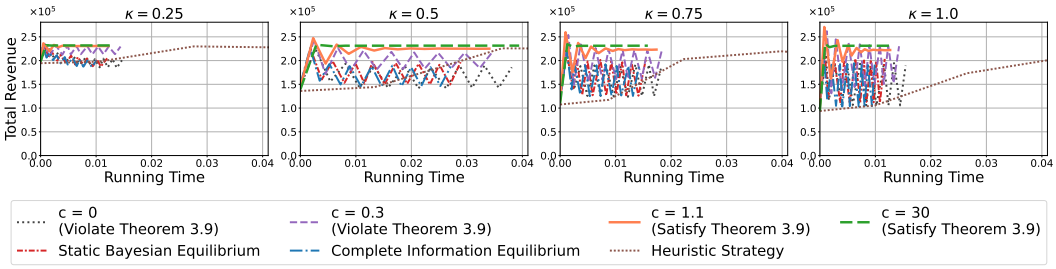


Figure 3: Total revenue for our algorithm with $c = 0, 0.3, 1.1, 30$ and three baselines in Bertrand competition within the ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1.0\}$ over 10 trials across wall-clock time.

We compare the wall-clock time of our algorithm with $c = 0, 0.3, 1.1, 30$ and the three baselines over 15 iterations in Figs. 2-3. For our algorithm with $c = 1.1, 30$, which satisfies the convergence condition $c > 1$, a performatively stable Bayesian equilibrium is successfully achieved within 0.015

seconds in both experiments. In contrast, our algorithm with $c = 0, 0.3$ (which violates the convergence condition $c > 1$), as well as the three baselines, fail to converge. These results not only validate our theoretical findings but also demonstrate the practical effectiveness of our algorithm in terms of wall-clock time.

Since all approaches were executed for a fixed budget of 15 iterations, we compare their running times over this common number of iterations. The Heuristic Strategy takes the longest time, as it must fit the data distribution maps and solve for the equilibrium using the learned maps. By comparison, the Complete Information Equilibrium and the Static Bayesian Equilibrium exhibit the shortest running times, which is not surprising: neither updates beliefs nor adapts the game structure in response to changes in the data distribution maps.

K.3 Cournot Competition Details

In the Cournot competition in crude oil trade with incomplete information, $n = 28$ crude oil exporting countries, covering 99% of the 2021 global exports[7], adjust crude oil export quantities to maximize revenue. These countries can only adjust their export quantities, and the price per barrel is determined by the global market. Utilizing 2021 crude oil trade data, we denote each country i's export quantity as q_i. The per barrel cost c_i of country i is sampled from a normal distribution $\mathcal{N}(\theta_i, \varsigma^2)$ at each time step, where θ_i is the type of country i and $\varsigma^2 = 10$ is the publicly known variance. The common prior belief about θ_i is a uniform distribution over $[\theta_-, \theta_+]$, i.e., $\mu(\theta_i) = \mathcal{U}[\theta_-, \theta_+]$, where $\theta_- = 10, \theta_+ = 50$. The bounds θ_- and θ_+ are the same for all countries. The crude oil price follows a publicly known linear inverse demand function $P = z - b\sum_{i=1}^{n}(q_i + x_i)$, where x_i is the decision representing export adjustment for country i, $q_i + x_i$ is the resulting export quantity, $z \sim \mathcal{D}(X)$ is a demand intercept parameter influenced by the joint decision X, and $b = 1 \times 10^{-8}$ quantifies the price sensitivity to the total supply. Each country i minimizes its expected loss by adjusting exports x_i:

$$\begin{aligned} &\mathbb{E}_{\substack{z\sim\mathcal{D}(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}} \ell_i(x_i, x_{-i}, z) = \mathbb{E}_{\substack{z\sim\mathcal{D}(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}} [-(P - c_i)\cdot(q_i + x_i)] \\ =&\mathbb{E}_{\substack{z\sim\mathcal{D}(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}} \left[c_i \cdot (q_i + x_i) - (z - b\sum_{j=1}^{n}(q_j + x_j))\cdot(q_i + x_i) \right], \end{aligned} \tag{113}$$

where $c_i \cdot (q_i + x_i)$ is the cost of country i, and $(q_i + x_i)\cdot(z - b\sum_{j=1}^{n}(q_j + x_j))$ is the revenue of country i. The strong monotonicity parameter of the game is the smallest eigenvalue of the Jacobian matrix $\mathcal{J}$ of the individual gradient $\nabla_Y \mathcal{G}_\mu(X)$ (Rockafellar & Wets, 1998). The eigenvalues of $\mathcal{J}$ in this game are $b(n + 1)$ and b. Given that $n = 28$ in our experiment, the smallest eigenvalue is b, which implies that $\psi = b$. The data distribution map $\mathcal{D}(X)$, which is identical across all countries, implies that adjustments in total exports induce a nonlinear transformation of the parameter z:

$$z \sim \mathcal{D}(X) = z_0 - \kappa \cdot \sinh^{-1}\left(\sum_{i=1}^{n} x_i\right) + w, \tag{114}$$

where $z_0 = 500$ is a base parameter, κ controls the sensitivity of the data distribution map $\mathcal{D}(X)$, $\sinh^{-1}(\cdot)$ denotes the inverse hyperbolic sine transformation, and $w \sim \mathcal{N}(0, 1)$ is the noise term. In our experiment, we set $\kappa = \{2.5, 5, 7.5, 10\}$ to simulate different performative effects on the price function $P(X)$, where larger values of κ correspond to higher sensitivities. The decision-making process of this game can be found in Appendix K.3.1. The belief update process of this game is provided in Appendix K.3.2.

[7]The crude oil trade dataset is available at: 'kaggle.com/datasets/toriqulstu/global-crude-petroleum-trade-1995-2021'.

K.3.1 Decision Making in Each Iteration of Cournot Competition

To solve this equilibrium at each iteration, we first calculate the best response of each country by the first-order optimality condition:

$$\begin{aligned}\nabla_{x_i}\ell_i &= -(-bq_i + z - b\sum_{j\neq i}^{n}(q_j + x_j) - bq_i - 2bx_i - c_i) \\ &= -(z - b\sum_{j\neq i}^{n}(q_j + x_j) - 2bq_i - 2bx_i - c_i) = 0 \\ \Longrightarrow x_i &= \frac{z - b\sum_{j\neq i}^{n}(q_j + x_j) - 2bq_i - c_i}{2b}\end{aligned} \tag{115}$$

Since $x_j, j \neq i$ is unknown to country i, the best response of country i is to take the expectation of $x_j, j \neq i$ according to the belief of cost $\mu(\theta_j), j \neq i$:

$$x_i^* = \frac{z - b\sum_{j\neq i}^{n}(q_j + \mathbb{E}[x_j(\theta_j)]) - 2bq_i - c_i}{2b}, \tag{116}$$

where the quantity $x_i(\theta_i), \forall i \in [n]$ is considered as a function of type θ_i:

$$x_i(\theta_i) = \frac{z - b\sum_{j\neq i}^{n}(q_j + \mathbb{E}[x_j(\theta_j)]) - 2bq_i}{2b} - \frac{1}{2b}c_i \text{ ,where } c_i \sim \mathcal{N}(\theta_i, \varsigma^2). \tag{117}$$

For country i, the optimal decision of country $j, j \neq i$ is a function of expectation of cost $\mathbb{E}[c_j] = \mathbb{E}[\theta_j]$:

$$\mathbb{E}[x_j(\theta_j)] = \frac{z - b\sum_{k\neq j}^{n}(q_k + \mathbb{E}[x_k(\theta_k)]) - 2bq_j}{2b} - \frac{1}{2b}\mathbb{E}[\theta_j]. \tag{118}$$

All countries except country j assume the optimal decision of country j is $\mathbb{E}[x_j(\theta_j)]$, for all $j \in [n]$. Thus, we solve $\mathbb{E}[x_i(\theta_i)], \forall i \in [n]$ with a system of linear equations:

$$\begin{aligned}2b\mathbb{E}[x_1(\theta_1)] + b\sum_{j\neq 1}^{n}\mathbb{E}[x_j(\theta_j)] &= z - \mathbb{E}[\theta_1] - b\sum_{j\neq 1}^{n} q_j - 2bq_1 \\ 2b\mathbb{E}[x_2(\theta_2)] + b\sum_{j\neq 2}^{n}\mathbb{E}[x_j(\theta_j)] &= z - \mathbb{E}[\theta_2] - b\sum_{j\neq 2}^{n} q_j - 2bq_2 \\ &\vdots \\ 2b\mathbb{E}[x_n(\theta_n)] + b\sum_{j\neq n}^{n}\mathbb{E}[x_j(\theta_j)] &= z - \mathbb{E}[\theta_n] - b\sum_{j\neq n}^{n} q_j - 2bq_n\end{aligned} \tag{119}$$

Specifically, we have:

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{120}$$

where

$$\mathbf{A} = \begin{bmatrix} 2b & b & \cdots & b \\ b & 2b & \cdots & b \\ \vdots & \vdots & \ddots & \vdots \\ b & b & \cdots & 2b \end{bmatrix}, \mathbf{x} = \begin{bmatrix} \mathbb{E}[x_1(\theta_1)] \\ \mathbb{E}[x_2(\theta_2)] \\ \vdots \\ \mathbb{E}[x_n(\theta_n)] \end{bmatrix}, \mathbf{b} = \begin{bmatrix} z - \mathbb{E}[\theta_1] - b\sum_{j\neq 1}^{n} q_j - 2bq_1 \\ z - \mathbb{E}[\theta_2] - b\sum_{j\neq 2}^{n} q_j - 2bq_2 \\ \vdots \\ z - \mathbb{E}[\theta_n] - b\sum_{j\neq n}^{n} q_j - 2bq_n \end{bmatrix} \tag{121}$$

And the $\mathbb{E}[x_i(\theta_i)]$ can be solved by

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}. \tag{122}$$

Finally, the Bayesian equilibrium can be obtained by the first optimality condition in Eq. (116):

$$x_i^* = \frac{z - b\sum_{j\neq i}^{n}(q_j + \mathbb{E}[x_j(\theta_j)]) - 2bq_i - c_i}{2b}, \tag{123}$$

In this game, our algorithm applies regularization to the actual total quantity, q_i+x_i, by incorporating the quadratic term $\frac{\gamma}{2}\|q_i + x_i\|_2^2$. This ensures that the regularization accounts for both the initial quantity q_i and the adjustment x_i. Since $\|q_i + x_i\|_2^2$ is γ-strongly convex in x_i, this regularization satisfies the convergence condition in Theorem 3.9.

K.3.2 Belief update in each iteration

For all countries, the optimal decision can be observed:

$$x_i^* = \frac{z - b\sum_{j\neq i}^{n}(q_j + \mathbb{E}[x_j(\theta_j)]) - 2bq_i - c_i}{2b}, \tag{124}$$

Since $\mathbb{E}[x_j(\theta_j)], \forall j \neq i$ is publicly known, the only unknown variable is c_i. As a result, after observing x_i^*, the cost of country i at time step t can be calculated by:

$$c_i^{(t)} = z - b\sum_{j\neq i}^{n}(q_j + \mathbb{E}[x_j(\theta_j)]) - 2bq_i - 2bx_i^*. \tag{125}$$

The cost of each country $i \in [n]$ is sampled from a normal distribution $c_i \sim \mathcal{N}(\theta_i, \varsigma^2)$. The expectation of this normal distribution θ_i is the type of country i. For other countries $j \neq i$, it is only known that θ_i is uniformly distributed in $[\theta_-, \theta_+]$.

After each country i decides a quantity, the cost $c_i^{(t)}$ can be learned by other countries at time step t. In this case, the belief at time step $t = 0$ is $\mu(\theta_i) = \frac{1}{\theta_+ - \theta_-}$. Then, given θ_i and observed cost $c_i^{(1)}, \cdots, c_i^{(t)}$, the likelihood function (the probability density of player i having costs $c_i^{(1)}, \cdots, c_i^{(t)}$) is

$$\varrho_i(c_i^{(1)}, \cdots, c_i^{(t)}|\theta_i) = \prod_{k=1}^{t}\frac{1}{\sqrt{2\pi\varsigma^2}}e^{-\frac{(c_i^{(k)}-\theta_i)^2}{2\varsigma^2}} = \frac{1}{\prod_{k=1}^{t}\sqrt{2\pi\varsigma^2}}e^{-\frac{\sum_{k=1}^{t}(c_i^{(k)}-\theta_i)^2}{2\varsigma^2}}, \tag{126}$$

where $\frac{1}{\sqrt{2\pi\varsigma^2}}e^{-\frac{(c_i^{(k)}-\theta_i)^2}{2\varsigma^2}}$ is the density function of normal distribution $\mathcal{N}(\theta_i, \varsigma^2)$.

Consider the average of cost $\bar{c}_i = \frac{1}{t}\sum_{k=1}^{t} c_i^{(k)}$, we have:

$$\begin{aligned}
\sum_{k=1}^{t}(c_i^{(k)} - \theta_i)^2 &= \sum_{k=1}^{t}\Big((c_i^{(k)} - \bar{c}_i) + (\bar{c}_i - \theta_i)\Big)^2 \\
&= \sum_{k=1}^{t}\Big((c_i^{(k)} - \bar{c}_i)^2 - 2(c_i^{(k)} - \bar{c}_i)(\bar{c}_i - \theta_i) + (\bar{c}_i - \theta_i)^2\Big) \\
&= \sum_{k=1}^{t}(c_i^{(k)} - \bar{c}_i)^2 - 2\sum_{k=1}^{t}(c_i^{(k)} - \bar{c}_i)(\bar{c}_i - \theta_i) + \sum_{k=1}^{t}(\bar{c}_i - \theta_i)^2 \\
&= \sum_{k=1}^{t}(c_i^{(k)} - \bar{c}_i)^2 - 2(\sum_{k=1}^{t} c_i^{(k)} - t\bar{c}_i)(\bar{c}_i - \theta_i) + t(\bar{c}_i - \theta_i)^2 \\
&= \sum_{k=1}^{t}(c_i^{(k)} - \bar{c}_i)^2 + t(\bar{c}_i - \theta_i)^2,
\end{aligned} \tag{127}$$

Then, the belief is updated according to Bayes's rule as follows:

$$\begin{aligned}
\mu\left(\theta_i|c_i^{(1)},\cdots,c_i^{(t)}\right) &= \frac{\mu(\theta_i)\varrho_i(c_i^{(1)},\cdots,c_i^{(t)}|\theta_i)}{\int_{\theta_-}^{\theta_+}\mu(\tilde{\theta}_i)\varrho_i(c_i^{(1)},\cdots,c_i^{(t)}|\tilde{\theta}_i)d\tilde{\theta}_i} \\
&= \frac{\frac{1}{\theta_+-\theta_-}\varrho_i(c_i^{(1)},\cdots,c_i^{(t)}|\theta_i)}{\int_{\theta_-}^{\theta_+}\frac{1}{\theta_+-\theta_-}\varrho_i(c_i^{(1)},\cdots,c_i^{(t)}|\tilde{\theta}_i)d\tilde{\theta}_i} \\
&= \frac{\frac{1}{\prod_{k=1}^t\sqrt{2\pi\varsigma^2}}e^{-\frac{\sum_{k=1}^t(c_i^{(k)}-\theta_i)^2}{2\varsigma^2}}}{\int_{\theta_-}^{\theta_+}\frac{1}{\prod_{k=1}^t\sqrt{2\pi\varsigma^2}}e^{-\frac{\sum_{k=1}^t(c_i^{(k)}-\tilde{\theta}_i)^2}{2\varsigma^2}}d\tilde{\theta}_i} \\
&= \frac{e^{-\frac{\sum_{k=1}^t(c_i^{(k)}-\bar{c}_i)^2+t(\bar{c}_i-\theta_i)^2}{2\varsigma^2}}}{\int_{\theta_-}^{\theta_+}e^{-\frac{\sum_{k=1}^t(c_i^{(k)}-\bar{c}_i)^2+t(\bar{c}_i-\theta_i)^2}{2\varsigma^2}}d\tilde{\theta}_i} \\
&= \frac{e^{-\frac{t}{2\varsigma^2}(\theta_i-\bar{c}_i)^2}}{\int_{\theta_-}^{\theta_+}e^{-\frac{t}{2\varsigma^2}(\tilde{\theta}_i-\bar{c}_i)^2}d\tilde{\theta}_i}.
\end{aligned} \tag{128}$$

And the expectation of θ_i is

$$\begin{aligned}
\mathbb{E}[\theta_i|c_i^{(1)},\cdots,c_i^{(t)}] &= \int_{\theta_-}^{\theta_+}\theta_i\mu\left(\theta_i|c_i^{(1)},\cdots,c_i^{(t)}\right)d\theta_i \\
&= \frac{\int_{\theta_-}^{\theta_+}\theta_i e^{-\frac{t}{2\varsigma^2}(\theta_i-\bar{c}_i)^2}d\theta_i}{\int_{\theta_-}^{\theta_+}e^{-\frac{t}{2\varsigma^2}(\tilde{\theta}_i-\bar{c}_i)^2}d\tilde{\theta}_i}.
\end{aligned} \tag{129}$$

Set $\xi=\frac{\sqrt{t}(\theta_i-\bar{c}_i)}{\varsigma}, a_-^{(t)}=\frac{\sqrt{t}(\theta_--\bar{c}_i)}{\varsigma}, a_+^{(t)}=\frac{\sqrt{t}(\theta_+-\bar{c}_i)}{\varsigma}$, then $d\theta_i=\frac{\varsigma}{\sqrt{t}}d\xi$.

Then, we have:

$$\begin{aligned}
\int_{\theta_-}^{\theta_+}e^{-\frac{t}{2\varsigma^2}(\tilde{\theta}_i-\bar{c}_i)^2}d\tilde{\theta}_i &= \frac{\varsigma}{\sqrt{t}}\int_{a_-^{(t)}}^{a_+^{(t)}}e^{-\frac{\xi^2}{2}}d\xi \\
&= \frac{\varsigma}{\sqrt{t}}\cdot\sqrt{2\pi}\int_{a_-^{(t)}}^{a_+^{(t)}}\frac{1}{\sqrt{2\pi}}e^{-\frac{\xi^2}{2}}d\xi \\
&= \frac{\sqrt{2\pi}\varsigma}{\sqrt{t}}(\Phi(a_+^{(t)})-\Phi(a_-^{(t)})),
\end{aligned} \tag{130}$$

where $\Phi(v)=\int_{-\infty}^{v}\frac{1}{\sqrt{2\pi}}e^{-\frac{\xi^2}{2}}d\xi$ is the standard normal cumulative distribution function.

And we can also have:

$$\begin{aligned}
\int_{\theta_-}^{\theta_+}\theta_i e^{-\frac{t}{2\varsigma^2}(\theta_i-\bar{c}_i)^2}d\theta_i &= \frac{\varsigma}{\sqrt{t}}\int_{a_-^{(t)}}^{a_+^{(t)}}(\bar{c}_i+\frac{\varsigma\xi}{\sqrt{t}})e^{-\frac{\xi^2}{2}}d\xi \\
&= \frac{\varsigma\bar{c}_i}{\sqrt{t}}\int_{a_-^{(t)}}^{a_+^{(t)}}e^{-\frac{\xi^2}{2}}d\xi+\frac{\varsigma^2}{t}\int_{a_-^{(t)}}^{a_+^{(t)}}\xi e^{-\frac{\xi^2}{2}}d\xi \\
&= \frac{\sqrt{2\pi}\varsigma\bar{c}_i}{\sqrt{t}}(\Phi(a_+^{(t)})-\Phi(a_-^{(t)}))+\frac{\varsigma^2}{t}\left[-e^{-\frac{\xi^2}{2}}\right]_{a_-^{(t)}}^{a_+^{(t)}} \\
&= \frac{\sqrt{2\pi}\varsigma\bar{c}_i}{\sqrt{t}}(\Phi(a_+^{(t)})-\Phi(a_-^{(t)}))+\frac{\varsigma^2}{t}\left(-e^{-\frac{(a_+^{(t)})^2}{2}}+e^{-\frac{(a_-^{(t)})^2}{2}}\right) \\
&= \frac{\sqrt{2\pi}\varsigma\bar{c}_i}{\sqrt{t}}(\Phi(a_+^{(t)})-\Phi(a_-^{(t)}))+\frac{\sqrt{2\pi}\varsigma^2}{t}(\phi(a_-^{(t)})-\phi(a_+^{(t)})),
\end{aligned} \tag{131}$$

where $\phi(\xi) = \frac{1}{\sqrt{2\pi}} e^{-\frac{\xi^2}{2}}$ is the standard normal density function.

Then, we have:

$$\begin{aligned}\mathbb{E}[\theta_i | c_i^{(1)}, \cdots, c_i^{(t)}] &= \frac{\frac{\sqrt{2\pi}\varsigma\bar{c}_i}{\sqrt{t}}(\Phi(a_+^{(t)}) - \Phi(a_-^{(t)})) + \frac{\sqrt{2\pi}\varsigma^2}{t}(\phi(a_-^{(t)}) - \phi(a_+^{(t)}))}{\frac{\sqrt{2\pi}\varsigma}{\sqrt{t}}(\Phi(a_+^{(t)}) - \Phi(a_-^{(t)}))} \\ &= \bar{c}_i + \frac{\varsigma}{\sqrt{t}} \frac{\phi(a_-^{(t)}) - \phi(a_+^{(t)})}{\Phi(a_+^{(t)}) - \Phi(a_-^{(t)})}.\end{aligned} \tag{132}$$

K.4 Additional Experimental Results on Cournot Competition

Table 1: Total revenue for our algorithm with $c = 0, 0.3, 1.1, 30$ and three baselines in Cournot competition within crude oil trade, with κ being the parameter in $\mathcal{D}(X)$. For each method, the average result and standard deviation over 10 trials are reported. For each setting, the best result is highlighted in **bold**, and the second-best result is *italicized*.

	$\kappa = 2.5$	$\kappa = 5.0$	$\kappa = 7.5$	$\kappa = 10.0$
$c = 0$ (Violate Thm. 3.9)	$1.17 \times 10^{12} \pm 5.63 \times 10^{11}$	$8.66 \times 10^{11} \pm 4.73 \times 10^{11}$	$7.38 \times 10^{11} \pm 4.02 \times 10^{11}$	$5.95 \times 10^{11} \pm 3.70 \times 10^{11}$
$c = 0.3$ (Violate Thm. 3.9)	$1.20 \times 10^{12} \pm 5.00 \times 10^{11}$	$1.19 \times 10^{12} \pm 2.58 \times 10^{11}$	$1.06 \times 10^{12} \pm 1.90 \times 10^{11}$	$8.73 \times 10^{11} \pm 8.15 \times 10^{10}$
$c = 1.1$ (Satisfy Thm. 3.9)	*$2.33 \times 10^{12} \pm 2.72 \times 10^{11}$*	**$2.47 \times 10^{12} \pm 4.06 \times 10^{10}$**	**$2.00 \times 10^{12} \pm 3.86 \times 10^{10}$**	**$1.49 \times 10^{12} \pm 2.90 \times 10^{10}$**
$c = 30$ (Satisfy Thm. 3.9)	**$2.47 \times 10^{12} \pm 9.40 \times 10^{10}$**	*$1.85 \times 10^{12} \pm 7.50 \times 10^{10}$*	*$1.61 \times 10^{12} \pm 7.17 \times 10^{10}$*	*$1.43 \times 10^{12} \pm 6.94 \times 10^{10}$*
SBE	$1.00 \times 10^{12} \pm 6.73 \times 10^{11}$	$3.75 \times 10^{11} \pm 3.92 \times 10^{11}$	$1.56 \times 10^{11} \pm 3.89 \times 10^{11}$	$5.26 \times 10^{11} \pm 5.09 \times 10^{11}$
CIE	$8.30 \times 10^{11} \pm 5.97 \times 10^{11}$	$8.93 \times 10^{11} \pm 3.06 \times 10^{11}$	$7.70 \times 10^{11} \pm 3.60 \times 10^{11}$	$7.17 \times 10^{11} \pm 3.65 \times 10^{11}$
HS	$0 \times 10^{0} \pm 0 \times 10^{0}$	$0 \times 10^{0} \pm 0 \times 10^{0}$	$0 \times 10^{0} \pm 0 \times 10^{0}$	$0 \times 10^{0} \pm 0 \times 10^{0}$

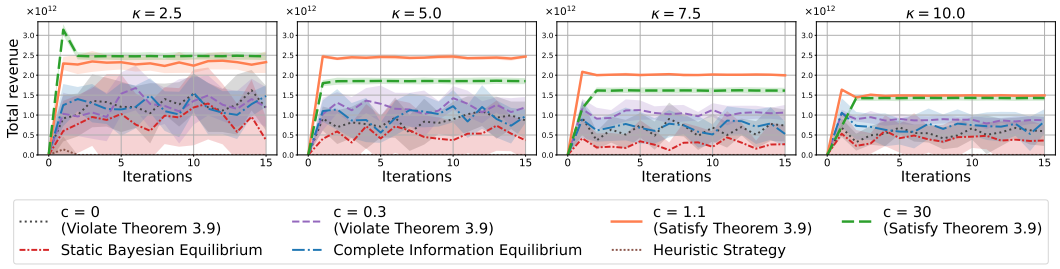


Figure 4: Total revenue and the standard deviation with $c = 0, 0.3, 1.1, 30$ in Cournot competition within crude oil trade with $\kappa = \{2.5, 5, 7.5, 10\}$ over 10 trials.

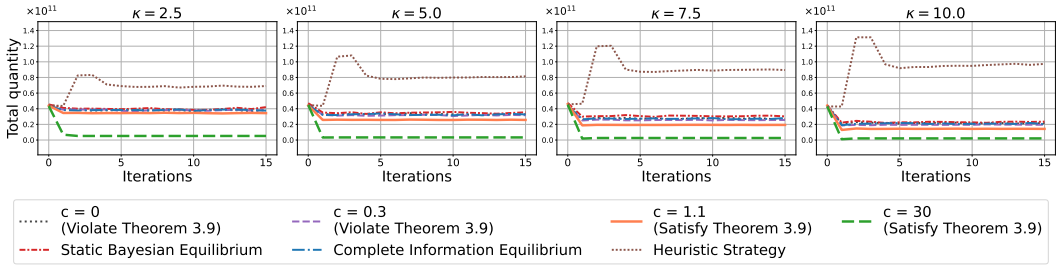


Figure 5: Total quantity with $c = 0, 0.3, 1.1, 30$ in Cournot competition within crude oil trade with $\kappa = \{2.5, 5, 7.5, 10\}$ over 10 trials.

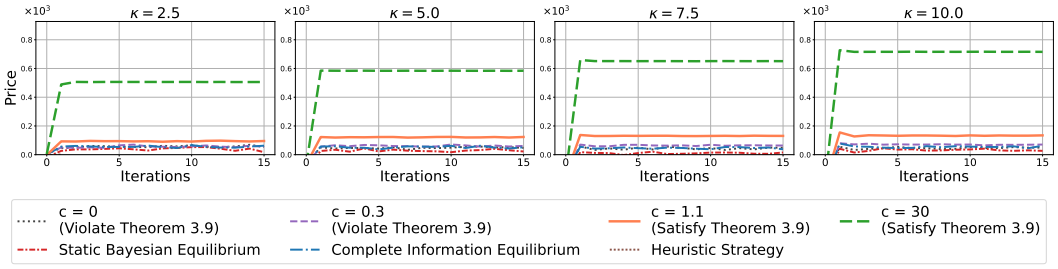


Figure 6: Crude oil price with $c = 0, 0.3, 1.1, 30$ in Cournot competition within crude oil trade with $\kappa = \{2.5, 5, 7.5, 10\}$ over 10 trials.

K.5 Bertrand Competition Details

In the Bertrand competition within the ride-share market under incomplete information, two companies ($n = 2$), Lyft and Uber, set optimal prices to maximize revenue, where each company's pricing decisions influence its own demand and that of its competitors. Specifically, when a company raises its prices, it reduces demand for its own services while increasing demand for its competitors' services. In this game, two companies adjust prices to maximize revenue across 11 locations in Boston[8]. The dataset's ride-share records are grouped into price intervals $[10, 15), [15, 20), [20, 25), [25, 30), [30, +\infty)$, represented by their lower bounds ($p = 10, 15, 20, 25, 30$) (Narang et al., 2023). The game is conducted independently in each price interval. Records are assigned to intervals according to their actual prices. For a given interval p, company i sets its price adjustments $x_i \in \mathbb{R}^{11}$ for the 11 locations, with each element in x_i representing the price adjustment in the corresponding location. Consequently, the actual price vector becomes $x_i + p\mathbf{1}$, where $\mathbf{1}$ is an all-ones vector in $\mathbb{R}^{11}$. For each interval, we calculate the initial demand for each company across 11 locations, denoted as $z_{0,i} \in \mathbb{R}^{11}$. The demand for company i is given by the publicly known function $q_i = z_i + \mathbf{A}_i x_i + \mathbf{A}_{-i} x_{-i}$, where $z_i \sim \mathcal{D}_i(X) \in \mathbb{R}^{11}$ is a demand parameter influenced by the joint decision X, and $\mathbf{A}_i$ and $\mathbf{A}_{-i} \in \mathbb{R}^{11\times 11}$ are publicly known matrices that quantify the demand sensitivity to the price adjustment with $\mathbf{A}_i \preceq 0, \mathbf{A}_{-i} \succeq 0$. The cost vector $c_i \in \mathbb{R}^{11}$ of each company i at each time step is sampled from a normal distribution $\mathcal{N}_{11}(\theta_i, \mathbf{\Sigma})$, where $\theta_i \in \mathbb{R}^{11}$ is the type of company i, $\mathbf{\Sigma} = \mathbf{I}_{11}$ is the publicly known covariance matrix, and $\mathbf{I}_{11}$ is the 11×11 identity matrix. The common prior belief about θ_i is a uniform distribution over $[\theta_-, \theta_+]$, i.e., $\mu(\theta_i) = \mathcal{U}[\theta_-, \theta_+]^{11}$, where $\theta_- = 2$ and $\theta_+ = 10$. In this experiment, the bounds θ_- and θ_+ are the same for both companies. Each company i aims to maximize its revenue in the price interval by minimizing the loss function:

$$\mathbb{E}_{\substack{z_i\sim\mathcal{D}_i(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}} \ell_i = \mathbb{E}_{\substack{z_i\sim\mathcal{D}_i(X)\\ \theta_{-i}\sim\mu(\theta_{-i}|h)}} \left[-(p\mathbf{1} + x_i - c_i)^T (z_i + \mathbf{A}_i x_i + \mathbf{A}_{-i} x_{-i})\right]. \tag{133}$$

When ride-share prices are extremely high or low, the demand is insensitive to price changes. When prices are within a moderate range, price adjustments can lead to a significant shift in demand. The data distribution map is defined as:

$$z_i \sim \mathcal{D}_i(X) = \frac{2z_{0,i}}{1 + e^{-(\mathbf{B}_i x_i + \mathbf{B}_{-i} x_{-i})}} + w_i, \tag{134}$$

where $z_{0,i}$ is the initial demand of company i across 11 locations, $\mathbf{B}_i, \mathbf{B}_{-i} \in \mathbb{R}^{11\times 11}$ are matrices controlling the sensitivity of data distribution map that are unknown to companies, and $w_i \sim \mathcal{N}_{11}(0, 0.01\mathbf{I}_{11})$ is the noise term. We generate matrices $\mathbf{B}_i$ and $\mathbf{B}_{-i}$ as $\mathbf{B}_i = diag(\mathcal{N}_{11}(-\kappa, (\frac{\kappa}{5})^2)) + E$ and $\mathbf{B}_{-i} = diag(\mathcal{N}_{11}(\frac{\kappa}{2}, (\frac{\kappa}{10})^2)) + E$, respectively, where $E_{ij} \sim \mathcal{N}(0, 0.01)$, and the diagonal entries are truncated to ensure $\mathbf{B}_i \preceq 0$ and $\mathbf{B}_{-i} \succeq 0$. In our experiment, we vary $\kappa = \{0.25, 0.5, 0.75, 1.0\}$, increasing the sensitivity as κ grows. The decision-making process of this game is provided in K.5.1. The belief update process of this game is provided in K.5.2.

[8]The Lyft and Uber Dataset is available at 'kaggle.com/datasets/brllrb/uber-and-lyft-dataset-boston-ma'.

K.5.1 Decision Making in Each Iteration of Bertrand Competition

To solve the equilibrium, we first derive the best response of each company using the first-order optimality condition. Since the type of competitor θ_{-i} is unknown to each company, its decision depends on the expectation of θ_{-i}, i.e., $\mathbb{E}[x_{-i}(\theta_{-i})]$. The first-order condition for company i is:

$$\begin{aligned} &\nabla_{x_i}\ell_i = -z_i - (\mathbf{A}_i + \mathbf{A}_i^T)x_i - \mathbf{A}_{-i}\mathbb{E}[x_{-i}(\theta_{-i})] - p\mathbf{A}_i^T\mathbf{1} + \mathbf{A}_i^T c_i = 0 \\ \Longrightarrow &(\mathbf{A}_i + \mathbf{A}_i^T)x_i = -z_i - \mathbf{A}_{-i}\mathbb{E}[x_{-i}(\theta_{-i})] - p\mathbf{A}_i^T\mathbf{1} + \mathbf{A}_i^T c_i \end{aligned} \tag{135}$$

For company i, the expected cost of the competitor equals the expected value of its type, i.e., $\mathbb{E}[c_{-i}] = \mathbb{E}[\theta_{-i}]$. To obtain the expected price adjustments $\mathbb{E}[x_i(\theta_i)]$ for both firms, we solve the following system of equations:

$$\begin{aligned} (\mathbf{A}_1 + \mathbf{A}_1^T)\mathbb{E}[x_1(\theta_1)] &= -z_1 - \mathbf{A}_{-1}\mathbb{E}[x_2(\theta_2)] - \mathbf{A}_1^T p\mathbf{1} + \mathbf{A}_1^T\mathbb{E}[\theta_1] \\ (\mathbf{A}_2 + \mathbf{A}_2^T)\mathbb{E}[x_2(\theta_2)] &= -z_2 - \mathbf{A}_{-2}\mathbb{E}[x_1(\theta_1)] - \mathbf{A}_2^T p\mathbf{1} + \mathbf{A}_2^T\mathbb{E}[\theta_2] \end{aligned} \tag{136}$$

Then, the optimal decision for player i is:

$$x_i^* = (\mathbf{A}_i + \mathbf{A}_i^T)^{-1}(-z_i - \mathbf{A}_{-i}\mathbb{E}[x_{-i}(\theta_{-i})] - \mathbf{A}_i^T p\mathbf{1} + \mathbf{A}_i^T c_i). \tag{137}$$

K.5.2 Belief update in each iteration

For all companies, the optimal decision can be observed:

$$x_i^* = (\mathbf{A}_i + \mathbf{A}_i^T)^{-1}(-z_i - \mathbf{A}_{-i}\mathbb{E}[x_{-i}(\theta_{-i})] - \mathbf{A}_i^T p\mathbf{1} + \mathbf{A}_i^T c_i). \tag{138}$$

Since $\mathbb{E}[x_j(\theta_j)], \forall j \neq i$ is publicly known, the only unknown variable is c_i. As a result, after observing x_i^*, the cost of company i at time step t can be calculated by:

$$c_i^{(t)} = \left(\mathbf{A}_i^T\right)^{-1}\left((\mathbf{A}_i + \mathbf{A}_i^T)x_i^* + z_i + \mathbf{A}_{-i}\mathbb{E}[x_{-i}(\theta_{-i})] + \mathbf{A}_i^T p\mathbf{1}\right) \tag{139}$$

Since the initial belief structure and the sampling method for cost c_i are identical to those in the Cournot competition, the belief update process remains the same. The remaining steps follow the procedure detailed in K.3.2, starting from Eq. (126).

K.6 Additional Experimental Results on Bertrand Competition

Table 2: Total revenue for our algorithm with $c = 0, 0.3, 1.1, 30$ and three baselines in Bertrand competition within the ride-share market, with κ being the parameter in $\mathcal{D}_i(X)$. For each method, the average result and standard deviation over 10 trials are reported. On each setting, the best result is highlighted in **bold**, and the second-best result is *italicized*.

	κ = 0.25	κ = 0.5	κ = 0.75	κ = 1.0
$c = 0$ (Violate Thm. 3.9)	$2.01 \times 10^5 \pm 9.66 \times 10^3$	$1.86 \times 10^5 \pm 1.43 \times 10^4$	$1.87 \times 10^5 \pm 1.73 \times 10^4$	$1.89 \times 10^5 \pm 1.46 \times 10^4$
$c = 0.3$ (Violate Thm. 3.9)	$2.30 \times 10^5 \pm 5.45 \times 10^3$	$2.15 \times 10^5 \pm 1.16 \times 10^4$	$2.21 \times 10^5 \pm 1.93 \times 10^4$	*$2.30 \times 10^5 \pm 2.09 \times 10^4$*
$c = 1.1$ (Satisfy Thm. 3.9)	*$2.31 \times 10^5 \pm 6.29 \times 10^2$*	*$2.25 \times 10^5 \pm 6.19 \times 10^2$*	*$2.23 \times 10^5 \pm 7.93 \times 10^2$*	$2.22 \times 10^5 \pm 7.76 \times 10^2$
$c = 30$ (Satisfy Thm. 3.9)	$\mathbf{2.32 \times 10^5 \pm 4.60 \times 10^1}$	$\mathbf{2.31 \times 10^5 \pm 4.51 \times 10^1}$	$\mathbf{2.31 \times 10^5 \pm 6.38 \times 10^1}$	$\mathbf{2.31 \times 10^5 \pm 6.04 \times 10^1}$
SBE	$2.05 \times 10^5 \pm 1.11 \times 10^4$	$1.97 \times 10^5 \pm 1.43 \times 10^4$	$2.01 \times 10^5 \pm 1.53 \times 10^4$	$2.03 \times 10^5 \pm 1.65 \times 10^4$
CIE	$2.01 \times 10^5 \pm 1.10 \times 10^4$	$1.87 \times 10^5 \pm 1.61 \times 10^4$	$1.89 \times 10^5 \pm 1.60 \times 10^4$	$1.91 \times 10^5 \pm 1.59 \times 10^4$
HS	$2.04 \times 10^5 \pm 9.20 \times 10^3$	$6.31 \times 10^4 \pm 1.54 \times 10^5$	$1.27 \times 10^5 \pm 7.94 \times 10^4$	$2.92 \times 10^4 \pm 1.53 \times 10^5$

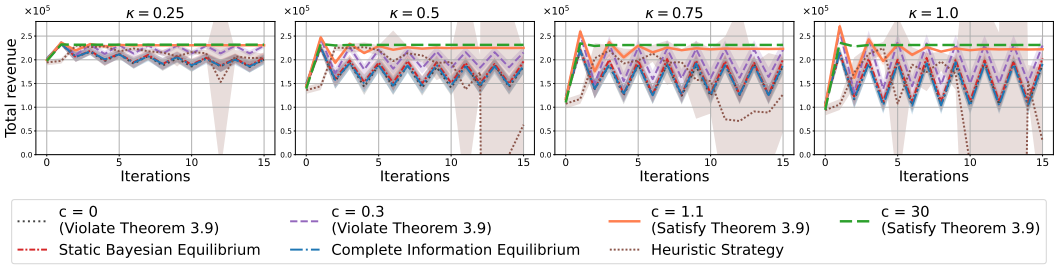


Figure 7: Total revenue and the standard deviation with $c = 0, 0.3, 1.1, 30$ in Bertrand competition within ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1.0\}$ over 10 trials.

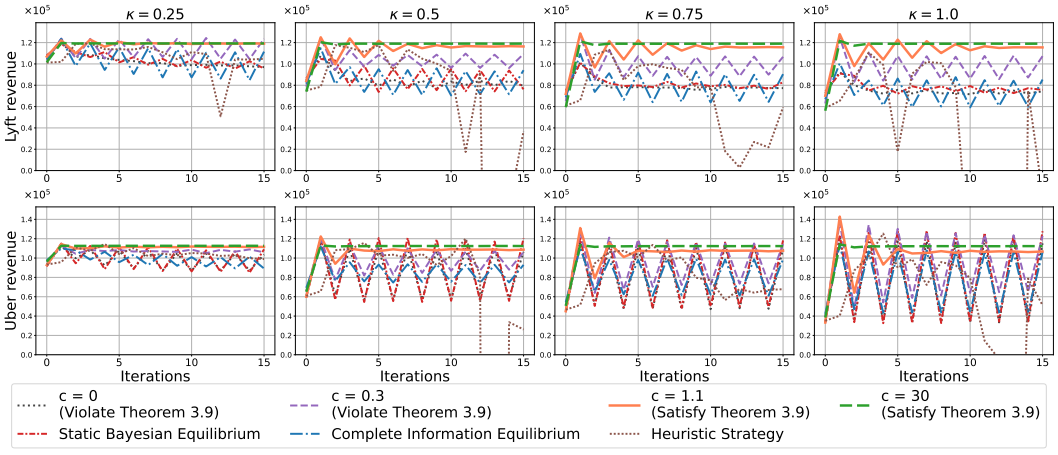


Figure 8: Revenues of each company with $c = 0, 0.3, 1.1, 30$ in Bertrand competition within ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1.0\}$ over 10 trials.

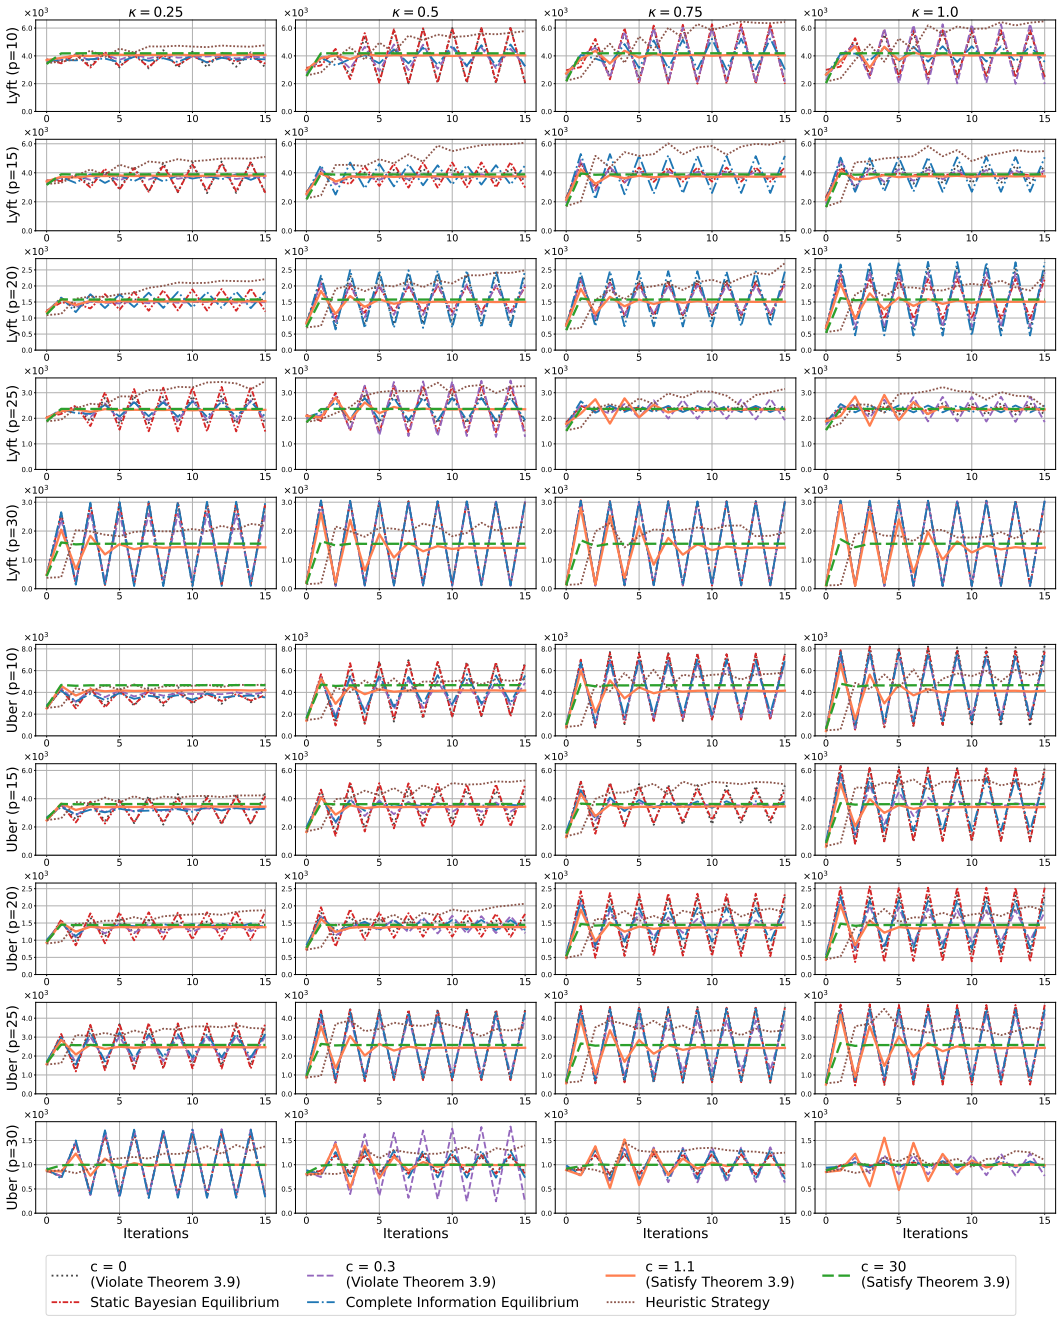


Figure 9: Demands of each company in each price interval with $c = 0, 0.3, 1.1, 30$ in Bertrand competition within ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1.0\}$ over 10 trials.

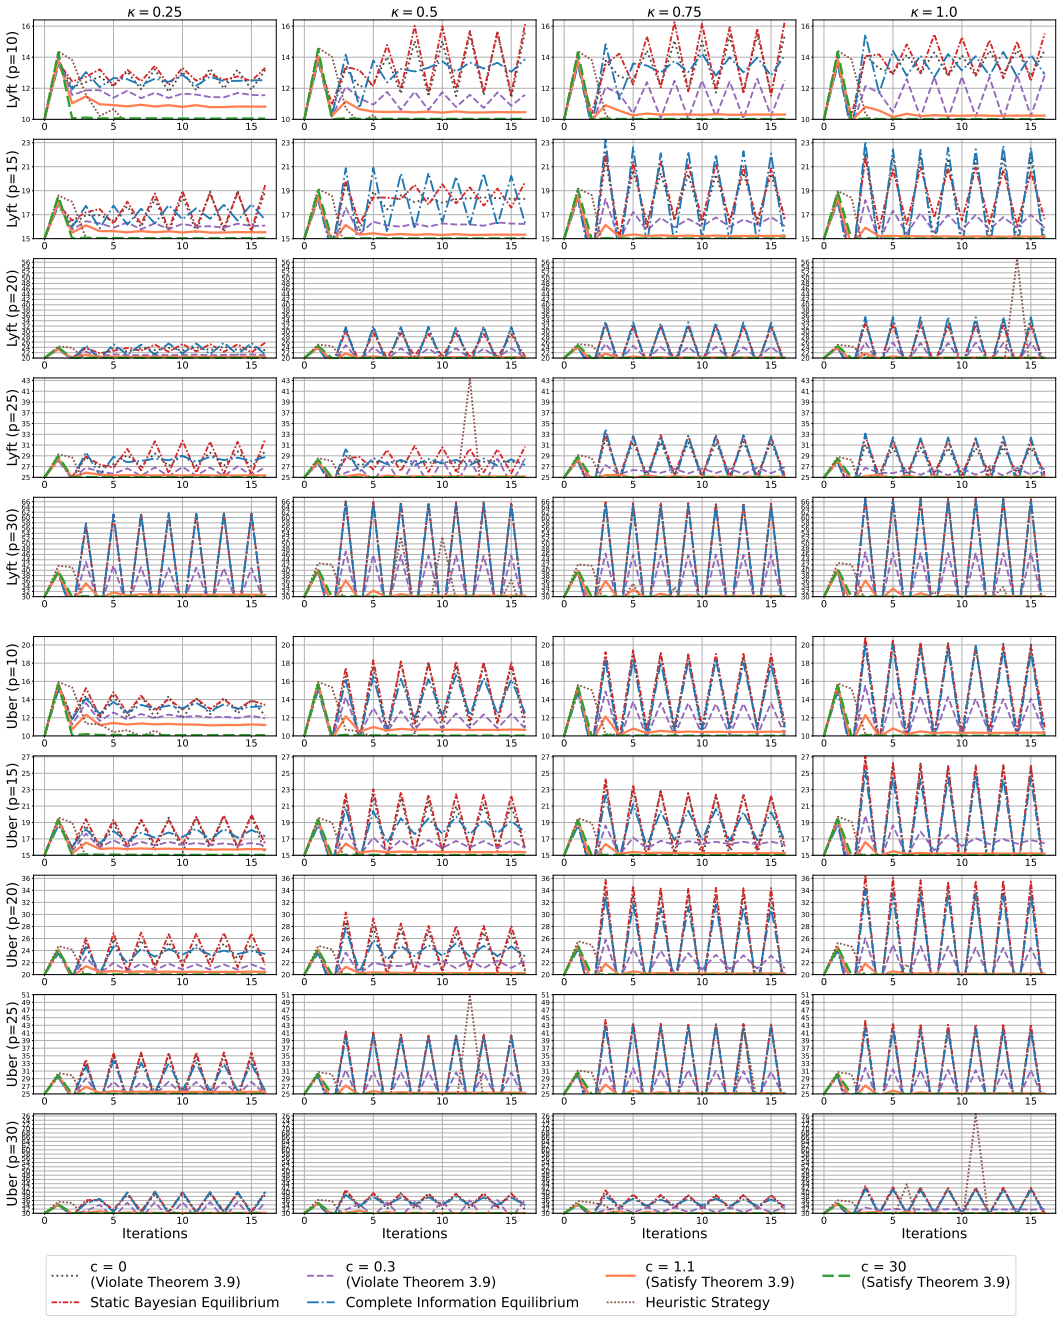


Figure 10: Prices of each company in each price interval with $c = 0, 0.3, 1.1, 30$ within Bertrand competition in ride-share market with $\kappa = \{0.25, 0.5, 0.75, 1.0\}$ over 10 trials.

L Potential Limitations and Future Directions

We point out several future directions for extending our framework. First, the strong monotonicity assumption may not hold in complex scenarios. Without this condition, a Bayesian equilibrium based on deterministic decisions may not exist. To relax this assumption, we will incorporate the mixed-strategy framework. Each player optimizes a probability distribution over its decision space and samples a decision from this distribution. This framework ensures equilibrium existence, enabling extension to non-monotone games. Second, the incomplete information setting can be extended by considering unobservable decisions, where players cannot update their beliefs through

others' decisions. To solve for equilibrium, we will use the belief update rule based on the realized loss value and shifted data distributions, both of which are influenced by others' decisions. Finally, the performative effect may also influence the players' private information in practice. For example, the cost of each company is potentially affected by the others' quantities, since a large production may lead to an increase in the price of raw materials. To address this, we will extend our $\hat{\varepsilon}_i$-sensitivity to quantify shifts in private information and will adapt our theoretical analysis and algorithm design accordingly.