A Data-Driven Distributionally Robust Game Using Wasserstein Distance

Guanze Peng(✉), Tao Zhang, and Quanyan Zhu

Department of Electrical and Computer Engineering, Tandon School of Engineering,
New York University, Brooklyn, NY 11201, USA
{guanze.peng,tz636,quanyan.zhu}@nyu.edu

Abstract. This paper studies a special class of games, which enables the players to leverage the information from a dataset to play the game. However, in an adversarial scenario, the dataset may not be trustworthy. We propose a distributionally robust formulation to introduce robustness against the worst-case scenario and tackle the curse of the optimizer. By applying Wasserstein distance as the distribution metric, we show that the game considered in this work is a generalization of the robust game and data-driven empirical game. We also show that as the number of data points in the dataset goes to infinity, the game considered in this work boils down to a Nash game. Moreover, we present the proof of the existence of distributionally robust equilibria and a tractable mathematical programming approach to solve for such equilibria.

Keywords: Data-driven optimization · Distributionally robust game · Mathematical programming

1 Introduction

In the past decade, game theory as a powerful mathematical tool has been used by researchers to analyze security issues in Cyber-physical systems (CPS) [14], Internet-of-Things [5], cloud computing [20], etc. As the advancements in data analysis, attackers can deploy more sophisticated attacks using information from the dataset [4,15,16]. The dataset can be log files, connection histories, or server deployment diagrams. The defender can also use statistical methods to defend herself from these attacks. The classical game theory approach does not capture this data-driven feature of modern security concerns. Thus, there are potentials in combining data science and game theory to further the analysis of the case where the players extract information from data to play the game.

With reference to Fig. 1, consider the following cyber security scenario: both the attacker and the defender have the access to an open-source dataset. Both of

This research is partially supported by awards ECCS-1847056, CNS-1544782, CNS-2027884, and SES-1541164 from National Science of Foundation (NSF), and grant W911NF-19-1-0041 from Army Research Office (ARO).

Q. Zhu et al. (Eds.): GameSec 2020, LNCS 12513, pp. 405–421, 2020.
https://doi.org/10.1007/978-3-030-64793-3_22

them aim to improve their performance by using this dataset. Nevertheless, fully trusting this dataset is not plausible as the dataset can be either incomplete or sometimes intentionally poisoned. Mathematically speaking, blindly extracting information from a dataset in an empirical fashion oftentimes will result in an overoptimistic result. We propose a distributionally robust game framework capture the balance between optimism and conservativeness. In this work, we assume that all the players have the same uncertainty of the game, i.e., there is no information privately possessed by any players. We also assume that the uncertainty can be characterized by a random variable. Each player faces a distributionally robust optimization problem and is robust to the worst-case distribution of the uncertainty parameter in the system model.

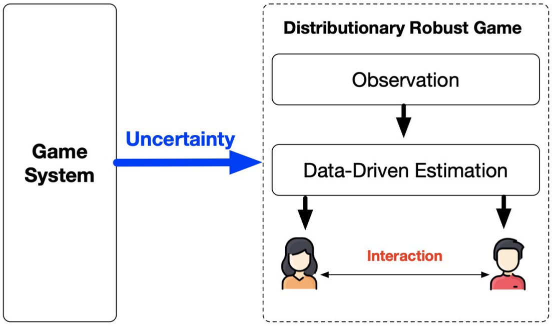

Fig. 1. A block diagram of the interaction between the attacker and the defender.

Our contributions are summarized as follows: we first define a data-driven empirical game (EG). A data-driven EG involves players estimating the distribution of the uncertainty parameter in an empirical way, and the players are able to learn the true distribution asymptotically. The empirical players suffer from the curse of the optimizer and oftentimes are too optimistic. Therefore, we propose a data-driven distributionally robust game framework to combat the overoptimism, while making sure that players are too not conservative as in robust games. We identify the relations between the proposed game with existing games. We define a special class of equilibria, which is termed *distributionally robust equilibrium (DRE)*. As the ambiguity in distribution can be characterized by a robustness parameter, this DRE can potentially simplify the distributionally robust mechanism design problem. Besides, as the ambiguity is generated by a dataset, the DRE considered in this work is endowed with the data-driven feature, which allows the possibility of sequential mechanism design. We show that when the robustness parameter goes to zero, the game boils down to an *empirical game* which is a Nash game. And when the robustness parameter goes to infinity, the game becomes a classical robust game. Then, we prove the existence of the DRE using Kakutani's fixed point theorem. Finally, we present a mathematical programming to solve for DRE.

Our work is closely related to [4,11], in which the authors provide the performance guarantees and tractable formulations for a data-driven distributionally robust optimization problem using the Wasserstein metric. Moreover, as the equilibrium concept considered here falls in the category of *Knightian equilibrium*, our work is related to [9] as well. There are also numerous papers on distributionally robust game theory [3,12,17], in which the ambiguity sets are not data-driven. Our work can be considered as a generalization of robust game [1], where the authors focus on the distribution-free setting.

Section 3 reviews the robust game theory. Section 4 develops a data-driven game model in which players utilize the information from data empirically. In Sect. 5, we first motive the formulation of the data-driven distributionally robust game. Then, we formally define such games, prove the existence of the equilibrium, and provide a tractable mathematical programming approach to solve for such equilibria. In Sect. 6, we use a bimatrix game as a toy example to validate the convergence result. Finally, Sect. 7 concludes the paper and points out the possible directions for future work.

2 Preliminaries

Let $\xi \in \Xi \subseteq \mathbb{R}^m$ be a random variable, where $m \in \mathbb{Z}_+$. Let $\mathcal{M}(\Xi)$ be the space of all probability distributions $\mathbb{Q}$ supported on Ξ with $\mathbb{E}_{\mathbb{Q}}[\|\xi\|] = \int_{\Xi} \|\xi\| \mathbb{Q}(d\xi) < \infty$. Here, $\|\cdot\|$ represents an arbitrary norm on $\mathbb{R}^m$.

Definition 1. *(Wasserstein Distance)* *[19] The Wasserstein metric $d : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \to \mathbb{R}_+$ is defined via*

$$d(\mathbb{Q}_1, \mathbb{Q}_2) = \inf_{\pi \in \Pi} \left\{ \int_{(\xi_1;\xi_2) \in \Xi \times \Xi} \|\xi_1 - \xi_2\| \pi(d\xi_1, d\xi_2) \right\}$$

for all measures $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\Xi)$, where Π the space of all the joint distributions of ξ_1 and ξ_2 with marginals $\mathbb{Q}_1$ and $\mathbb{Q}_2$, respectively.

Theorem 1. *(Kakutani's Fixed-Point Theorem)* *[8] If $x \to \phi(x)$ has an upper semicontinuous point-to-set mapping of an r-dimensional closed simplex $\mathcal{S}$ in to $\omega_i(\mathcal{S})$, then $\exists\ x_0 \in \mathcal{S}$, such that $x_0 \in \phi(x_0)$.*

3 Robust Game

Consider an incomplete-information game (I-game) with a finite set of players $\mathcal{I} = \{1, 2, ..., N\}$ and a finite set of actions $\mathcal{A}_i \in \mathbb{R}^{A_i}$ for each player i, where $A_i \in \mathbb{Z}_+$. As mention by Harsanyi in [6], the incompetence of information induced by uncertainty can be summarized and embedded in the cost functions (objective functions, cost matrix). We denote the uncertain parameter by $\xi \in \Xi$. For player i, we define his cost functions as $C_i(a_i, a_{-i}; \xi) : \prod_{i \in \mathcal{I}} \mathcal{A}_i \times \Xi \to \mathbb{R}$, where $a_{-i} := (a_1, ..., a_{i-1}, a_{i+1}, ..., a_N)$ is the collection of other players' actions. Note that all the players considered in this work are minimizers. Moreover, we assume that the uncertainty set Ξ is finite.

In [2,6], with the assumptions common prior and common knowledge of rationality, we can transform an I-game to a complete-information game (C-game), which is commonly known as a Bayesian game. However, the Bayesian game fails to characterize the case where common prior or stochastic information of the uncertainty is unavailable.

In [1], the authors have proposed a distribution-free game framework to study incomplete-information games. In their proposed game, *robust game*, players are assumed to be robust to the uncertainty. Formally, a robust game can be defined as a tuple,

$$\mathcal{G} := (\mathcal{I}, \mathcal{S}), \tag{1}$$

where $\mathcal{S}$ is the state of nature. Every state of nature $s \in \mathcal{S}$ is a vector

$$s = (\mathcal{I}, (\mathcal{A}_i)_{i \in \mathcal{I}}, (c_i)_{i \in \mathcal{I}}),$$

where $\mathcal{A}_i$ is a nonempty finite set of actions of Player i. $c_i : \mathcal{A}_i \times \Xi \to \mathbb{R}$ the cost function of Player i where $\mathcal{A} = \bigtimes_{i \in \mathcal{I}} \mathcal{A}_i$. In this work, we assume that the players do not have private information and this allows us to transform the I-game $\mathcal{G}$ to a C-game.

For every $i \in \mathcal{I}$, let $\mathbf{x}_i$ be the mixed strategy of Player i, which is defined to be a probability over the action space, i.e., $\mathbf{x}_i = (x_i(a_i))_{a_i \in \mathcal{A}_i} \in \Delta_i := \Delta(\mathcal{A}_i)$ and $\Delta(\cdot)$ is the simplex of a finite set. For the ease of notation, define the expected cost induced by the mixed strategy profile $\mathbf{x} = (\mathbf{x}_i, \mathbf{x}_{-i})$ as the following

$$c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi) = \sum_{a_i \in \mathcal{A}_i} \sum_{a_{-i} \in \mathcal{A}_{-i}} C_i(a_i, a_{-i}; \xi) x_i(a_i) \prod_{j \neq i, j \in \mathcal{I}} x_j(a_j),$$

where $\mathcal{A}_{-i} = \bigtimes_{j \neq i, j \in \mathcal{I}} \mathcal{A}_i$. The equilibrium concept used in robust game $\mathcal{G}$ is given by the following:

Definition 2. *A mixed strategy profile $\mathbf{x}^* = (\mathbf{x}_i^*, \mathbf{x}_{-i}^*)$ is robust-optimization equilibrium solution in $\mathcal{G}$ if for $i \in \mathcal{I}$,*

$$\max_{\xi \in \Xi} \; c_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*; \xi) \leqslant \max_{\xi \in \Xi} \; c_i(\mathbf{x}_i, \mathbf{x}_{-i}^*; \xi), \tag{2}$$

where $\mathbf{x}_{-i} \in \Delta_{-i} := \bigtimes_{j \neq i, j \in \mathcal{I}} \Delta(\mathcal{A}_j)$.

The following theorem guarantees the existence of the robust-optimization equilibrium in $\mathcal{G}$.

Theorem 2. *(Existence of Equilibria in Robust Finite Games) [1] In the game defined by $\mathcal{G}$, if $C_i(a_i, a_{-i}; \xi)$ is bounded for all $i \in \mathcal{I}$, $(a_i, a_{-i}) \in \mathcal{A}$ and $\xi \in \Xi$, then there exists an ex post equilibrium.*

4 Data-Driven Empirical Game

In a data-driven empirical game (EG), we assume that the uncertainty parameter is a random variable, and is selected according to some unknown distribution by a chance move at the beginning of the game. Let $\mathbb{P}$ be the measure induced by the random variable ξ. The players can observe N such games played independently and the realizations of the uncertainty parameter. Then every player makes the estimation from the same set

$$\hat{\Xi}^{(N)} = \left\{\hat{\xi}^{(1)}, \hat{\xi}^{(2)}, \ldots, \hat{\xi}^{(N)}\right\} \subseteq \Xi^N, \tag{3}$$

which consists of N independent realizations of the random variable ξ. We call $\hat{\Xi}^{(N)}$ *dataset*, and each element in it *data point*. In [10], the author formalizes a framework which enables the players to learn as *statisticians*. Formally, define the learning rule as a mapping from the dataset (3) to the belief space:

$$\ell \; : \; \Xi^N \; \to \; \Delta(\Xi).$$

In particular, we are interested in *empirical players* in this work, who estimate $\mathbb{P}$ using an empirical approach as follows

$$\hat{\mathbb{Q}}^{(N)} := \ell\left(\hat{\Xi}^{(N)}\right) = \frac{1}{N}\sum_{n=1}^{N} \delta_{\hat{\xi}^{(n)}},$$

where δ is the Dirac delta function. We term $\hat{\mathbb{Q}}^{(N)}$ as *common empirical prior* in this work. The empirical learning rule not only is appealing for its neat and simple form, but also enjoys the following property:

Lemma 1. *Let the dataset $\hat{\Xi}^{(N)}$ be defined as (3) which contains N independent realizations of ξ. When the number of realizations (data points) N goes to infinity,*

$$\ell\left(\hat{\Xi}^{(N)}\right) \; \to \; \mathbb{P}, \quad a.s.$$

Proof. The proof of this lemma is an immediate result of law of large numbers. □

The lemma above says that as the number of samples goes to infinity, the empirical players can learn the *true distribution* of ξ, $\mathbb{P}$.

The empirical players can benefit from the information obtained from the dataset. Indeed, it is not hard to show that given $\mathbf{x}_{-i} \in \Delta_{-i}$, for every possible empirical measure $\hat{\mathbb{Q}}^{(N)}$ induced by the dataset $\hat{\Xi}^{(N)}$

$$\mathbf{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right] \leqslant \max_{\xi \in \Xi} c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi).$$

By letting

$$\zeta_i^*(\mathbf{x}_{-i}) \in \arg\min_{\mathbf{x}_i \in \Delta_i} \mathbf{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right],$$

we have that

$$\mathbf{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\zeta_i^*(\mathbf{x}_{-i}),\mathbf{x}_{-i};\xi)\right] \leqslant \min_{\mathbf{x}_i\in\Delta_i}\max_{\xi\in\Xi} c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi).$$

This inequality says that by leveraging the stochastic information from the dataset, the players behave *less conservatively.*

It is worth noting that as all the players make use of the same dataset to estimate in the same empirical fashion, they share the same empirical distribution. Thus, this distribution is also common knowledge. We further assume that the fact that all the players being empirical is common knowledge. We proceed by defining the data-driven EG, which falls into the category of I-game. A data-driven EG is given by a tuple

$$\mathcal{G}^{(N)} := \left(\mathcal{I},\mathcal{S},\mathbb{P},\hat{\Xi}^{(N)}\right),$$

where $\mathbb{P}$ is the true measure of ξ. Now, we are ready to show that data-driven EBG as an I-game is equivalent to a C-game. As mentioned earlier, the players acknowledge that all of them are empirical and they share the empirical distribution, the data-driven EG is equivalent to a Nash game by replacing the cost matrix $C_i(a_i,a_i;\xi)$ with

$$\tilde{C}_i(a_i,a_{-i}) := \mathbb{E}_{\hat{\mathbb{Q}}^{(N)}}\left[C_i(a_i,a_i;\xi)\right],$$

where the expectation is taken over ξ with respect to $\hat{\mathbb{Q}}^{(N)}$. Thus, data-driven EG is also equivalent to a C-game. Moreover, as it is equivalent to a Nash game, the existence of the Nash equilibrium is also guaranteed.

5 Data-Driven Distributionally Robust Game

In this section, we propose a new class of games which is termed data-driven Distributionally Robust Game (DRG) in which we use Wasserstein distance as the distribution metric. To motive this framework, we first answer a few essential questions.

5.1 Motivation

Why Distributionally Robust Formulation? The direct application of empirical distribution as estimated distribution suffers from *optimizer's curse* [18]. It is well known that the empirical estimator $\hat{\mathbb{Q}}^{(N)}$ is be an unbiased estimator of $\mathbb{P}$, i.e.,

$$\mathbb{E}_{\mathbb{Q}^{(N)}}\left[\hat{\mathbb{Q}}^{(N)}\right] = \mathbb{P}.$$

where $\mathbb{Q}^{(N)}$ is the measure induced by the N data points. With fixed $\mathbf{x}_{-i}$,

$$\mathbb{E}_{\mathbb{Q}^{(N)}}\left[\mathbb{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right]\right] = \mathbb{E}_{\mathbb{P}}\left[c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right].$$

By Jensen's inequality,

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}^{(N)}}\left[\min_{\mathbf{x}_i\in\Delta_i}\mathbb{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right]\right] &\leqslant \min_{\mathbf{x}_i\in\Delta_i}\mathbb{E}_{\mathbb{Q}^{(N)}}\left[\mathbb{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right]\right]\\ &= \min_{\mathbf{x}_i\in\Delta_i}\mathbb{E}_{\mathbb{P}}\left[\ c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right].\end{aligned}$$

Let

$$\zeta_i^*(\mathbf{x}_{-i})\ \in\ \arg\min_{\mathbf{x}_i\in\Delta_i}\mathbb{E}_{\hat{\mathbb{Q}}^{(N)}}\left[c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi)\right].$$

Then, for every $\mathbf{x}_{-i}\in\Delta_{-i}$,

$$\mathbb{E}_{\mathbb{P}}\left[\ c_i(\zeta_i^*(\mathbf{x}_{-i}),\mathbf{x}_{-i};\xi)\right]\geqslant\min_{\mathbf{x}_i\in\Delta_i}\mathbb{E}_{\mathbb{P}}\left[\ c_i(\mathbf{x}_i,\mathbf{x}_{-i};\xi_i)\right].$$

As shown above, given the other players' strategies, a player inclines to be *overoptimistic* due to the optimizer's curse. Therefore, it is reasonable to employ some "robustness" to deal with this overoptimism. In this work, given a tuple of his counterparts' strategies, we suppose that each player formulates the best response as the solution of a distributionally robust optimization problem.

Note that, in our framework, we assume that a player's opponents are outside the scope of the player's viewpoint. That is, the player takes the distributionally robust view only of the uncertainties of his cost function, with a tuple of the other players' strategies given. From this perspective, each player does not take a distributionally robust approach to his uncertainty with respect to this tuple itself. Moreover, we assume that each player's distributionally robust view of the game is common knowledge, which allows the players to predict each other's best-response correspondences. Thus, the players in the game defined by (4) can reach consistent predictions of what each other will play.

We interpret the distributionally robust game in a security setting. Suppose the players (defender and attacker) have the access to the same open-source dataset. On one hand, the players aim to obtain useful information from this dataset to achieve better defend/attack results. On the other hand, the dataset may not be reliable. It is natural for the players to be robust to the inference of the dataset. Hence, the distributionally robust formulation is a reasonable choice in a security problem in order to balance the optimism and conservativeness. However, one may have the concern over the reason why the players are assumed to know the same dataset. Indeed, in real world, the defender and the attacker oftentimes have different information (knowledge) due to different financial capabilities, backgrounds, identities, etc. In such cases, one may need to resort to Bayesian game framework. The information-asymmetric case is beyond the scope of this work and we leave it to future work.

Why Wasserstein Distance? In this work, we assume that each player adopts *Wasserstein Distance* as the metric measuring the difference between two distributions. Formally, a distributionally robust game using Wasserstein distance is defined by the following vector

$$\mathcal{G}_\epsilon^{(N)}=\left(\mathcal{I},\tilde{\mathcal{S}},\mathbb{P},\hat{\Xi}^{(N)}\right),\tag{4}$$

where $\tilde{s} := (\mathcal{I}, (\mathcal{A}_i)_{i\in\mathcal{I}}, (c_i)_{i\in\mathcal{I}}, \epsilon)$ and $\tilde{s} \in \tilde{\mathcal{S}}$. The parameter ϵ is the radius of the Wasserstein ball, which stands for the *robustness* of the players. It is determined by the nature and assumed to be common knowledge and the same for all the players. The key feature distinguishing Wasserstein distance as a distribution metric from other distribution metrics is that the worst-case distribution can be supported outside the dataset. In a game setting, the utilization of Wasserstein distance can be interpreted as the following: the knowledge of the support set the types are common knowledge shared between the players. Using Wasserstein distance as distribution metric enables the players to utilize this support information. Moreover, this allows the players to be robust against perturbations of the data points [4]. It also makes sense in a security scenario: both of the defender and the attacker want to use every bit of information available to improve their performance while maintaining certain level of robustness.

5.2 Equilibrium Concept

With the empirical distribution $\hat{\mathbb{Q}}^{(N)}$ being centered, we construct a Wasserstein ball as follows:

$$\mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right) = \left\{\mathbb{Q} \in \mathcal{M}(\Xi) : d(\mathbb{Q}, \hat{\mathbb{Q}}^{(N)}) \leqslant \epsilon\right\},$$

which contains all the possible probability measures, whose Wasserstein distance with the empirical distribution is less than ϵ. Here, $\mathcal{M}(\Xi)$ is the set of all the possible distributions whose support is Ξ.

Definition 3. *A mixed strategy profile* $\mathbf{x} = (\mathbf{x}_i^*, \mathbf{x}_{-i}^*)$ *is an* ***distributionally robust equilibrium (DRE)*** *solution if no player can decrease their interim expected cost by unilaterally changing their strategy: for i* $\in \mathcal{I}$ *and every mixed strategy* $\mathbf{x}_i \in \Delta_i$*,*

$$\sup_{\mathbb{Q}\in\mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}}\left[c_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*; \xi)\right] \leqslant \sup_{\mathbb{Q}\in\mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}^*, ; \xi)\right]. \tag{5}$$

Remark 1. By definition, DRE is a relaxation of *Knightian equilibrium*. In a homogeneous game where each player has the same objective function and action set, DRE falls in the category of *Knightian equilibrium*. The DRE exhibits several advantageous features:

1. The proposed DRE can be used as a solution concept in mechanism design and characterize the incentive compatibility such that each player has the incentive to truthfully reveal his private information in DRE. The players' uncertainty about their cost functions provides a potential opportunity for the mechanism designer to strategically design the ambiguity set as an additional rule of encounter to achieve the designer's social goal.
2. Suppose that the ambiguity set is given and not a part of the design. When the ambiguity set is different, one will need to solve the design problem all over again. The ambiguity set in DRE being induced by a dataset enables one to design a data-driven mechanism sequentially, as the only difference in ambiguity sets is the center of the Wasserstein ball.

When the robustness parameter ϵ goes to 0, then the Wasserstein ball collapses inward to $\{\hat{\mathbb{Q}}^{(N)}\}$. Consequently, a data-driven DRG becomes a data-driven EG, i.e.,

$$\lim_{\epsilon \to 0} \mathcal{G}_\epsilon^{(N)} = \mathcal{G}^{(N)}.$$

On the other hand, when ϵ goes to infinity, Data-Driven DRG becomes a classical robust game, as all the probability mass will concentrate on the worst-case support, i.e.,

$$\lim_{\epsilon \to \infty} \mathcal{G}_\epsilon^{(N)} = \mathcal{G}.$$

If we see ϵ as a tuning parameter, then $\mathcal{G}_\epsilon^{(N)}$ can be regarded as a generalization which bridges the robust game and data-driven EG.

5.3 Existence of DRE

In this section, we give the theoretical guarantee of the existence of DRE, which largely follows from Theorem 1 in [1]. In order to prove the existence of DRE in the game defined by $\mathcal{G}^{(N)}$, we first define the mapping $\rho_{i,\epsilon}^{(N)} : \Delta \times \Xi^N \to \Delta_i$ as the following

$$\rho_{i,\epsilon}^{(N)}(\mathbf{x}_i, \mathbf{x}_{-i}, \hat{\Xi}^{(N)}) = \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right]. \tag{6}$$

Moreover, we define the following "point-to-set" mapping for game $\mathcal{G}^{(N)}$,

$$\Phi_\epsilon^{(N)} : \Delta \times \Xi^N \to \Delta.$$

Specially, we choose $\Phi_\epsilon^{(N)}$ to be the following

$$\Phi_\epsilon^{(N)}(\mathbf{x}, \hat{\Xi}^{(N)}) = \left\{ \tilde{\mathbf{x}} = (\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_{-i}) \;\middle|\; \tilde{\mathbf{x}}_i \in \arg\min_{\mathbf{u}_i \in \Delta_i} \rho_{i,\epsilon}^{(N)}(\mathbf{u}_i, \mathbf{x}_{-i}, \hat{\Xi}^{(N)}), i \in \mathcal{I} \right\}, \tag{7}$$

which is the set of all the best response strategies given the strategy profile $\mathbf{x}$.

Theorem 3. *Let Ξ be finite, and $C_i(a_i, a_{-i}; \xi)$ be bounded for all $\xi \in \Xi$. There exists at least one DRE in the game defined by $\mathcal{G}_\epsilon^{(N)}$.*

Sketch of Proof. We start the proof by proving that $\rho_{i,\epsilon}^{(N)}(\mathbf{x}_i, \mathbf{x}_{-i}, \hat{\Xi}^{(N)})$ is continuous on Δ, and that for each $i \in \mathcal{I}$, $\rho_{i,\epsilon}^{(N)}(\mathbf{x}_i, \mathbf{x}_{-i}, \hat{\Xi}^{(N)})$ is convex in $\mathbf{x}_i$. Then, the mapping $\Phi_\epsilon^{(N)}$ can be shown to be non-empty, convex and upper semicontinuous. Applying Kakutani's fixed-point theorem immediately gives us the theorem. □

5.4 Asymptotic Consistency

We must notice that there may exist more than one equilibrium, i.e., the equilibrium may not be unique. Then, from now on, it will be reasonable to work on the set of equilibriums, which is given by

$$\mathcal{E}_{\epsilon}^{(N)}(\hat{\Xi}^{(N)}) = \left\{ \mathbf{x} \mid \mathbf{x} \in \Phi_{\epsilon}^{(N)}(\mathbf{x}, \hat{\Xi}^{(N)}) \right\}.$$

This set is non-empty due to Theorem 3.

When the true distribution of ξ is known to all the players, the problem boils down to a standard Nash game. This Nash game can be represented by a tuple $\mathcal{G}_{\text{Nash}} = (\mathcal{I}, \mathcal{S}, \mathbb{P})$. Similar to (6) and (7), define

$$\rho_i(\mathbf{x}_i, \mathbf{x}_{-i}) = \mathbb{E}_{\mathbb{P}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right],$$

and

$$\Phi(\mathbf{x}) = \left\{ \tilde{\mathbf{x}} = (\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_{-i}) \;\middle|\; \tilde{\mathbf{x}}_i \in \arg\min_{\mathbf{u}_i \in \Delta_i} \rho_i(\mathbf{u}_i, \mathbf{x}_{-i}),\ i \in \mathcal{I} \right\},$$

respectively. Characterized by fixed points, the set of equilibria in $\mathcal{G}_{\text{Nash}}$ is given by

$$\mathcal{E} = \left\{ \mathbf{x} \mid \mathbf{x} \in \Phi(\mathbf{x}) \right\}.$$

It is not hard to see that $\mathcal{E}$ is non-empty.

Proposition 1. *Define a sequence of Wasserstein ball radius $\{\epsilon_N\}_{N=1}^{\infty}$ with the following property*

$$\lim_{N \to \infty} \epsilon_N = 0.$$

Then,

$$\lim_{N \to \infty} \mathcal{E}_{\epsilon_N}^{(N)}(\hat{\Xi}^{(N)}) \;=\; \mathcal{E}, \quad a.\ s.$$

Proof. When N goes to infinity, by using Lemma 3.7 from [4], we obtain that

$$\mathbb{Q}^{\infty}\left[\lim_{N \to \infty} d(\mathbb{P}, \hat{\mathbb{Q}}^{(N)}) = 0\right] = 1.$$

Hence,

$$\begin{aligned}
\lim_{N \to \infty} \rho_{i,\epsilon}^{(N)}(\mathbf{x}_i, \mathbf{x}_{-i}, \hat{\Xi}^{(N)}) &= \lim_{N \to \infty} \sup_{\mathbb{Q} \in \mathcal{B}_{\epsilon}\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right] \\
&= \mathbb{E}_{\mathbb{P}}\left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}; \xi)\right] \\
&= \rho_i(\mathbf{x}_i, \mathbf{x}_{-i}), \quad \text{a. s.}
\end{aligned}$$

Then, it is clear that

$$\lim_{N \to \infty} \Phi_{\epsilon}^{(N)}(\mathbf{x}, \hat{\Xi}^{(N)}) \;=\; \Phi(\mathbf{x}), \quad \text{a. s.}$$

The argument in the proposition follows.

□

Remark 2. Proposition 1 exhibits the convergence result concerning the equilibrium set. As the number of data points goes to infinity, the distributionally robust game $\mathcal{G}^{(N)}$ is equivalent to the standard Nash game $\mathcal{G}$ in terms of the equilibria.

5.5 Tractable Formulations

In this section, we derive a tractable formulation using which one can solve for the DRE (as defined in (5)) in $\mathcal{G}_\epsilon^{(N)}$. Without loss of generality, we study a two-player game, i.e., $\mathcal{I} = \{1, 2\}$. Denote the cost matrix of the i-th player by

$$\mathbf{C}_i(\xi) = \left[C_i(a_1, a_2; \xi)\right]_{a_1 \in \mathcal{A}_1, a_2 \in \mathcal{A}_2}, \quad i \in \mathcal{I}.$$

Recall that the i-th player faces the following optimization problem:

$$\min_{\mathbf{x}_i \in \Delta_i} \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}} \left[\mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2\right]. \tag{8}$$

We drop the outer minimization for the clarity of notations. By the definition of Wasserstein ball, (8) can be rewritten as

$$\begin{aligned} \sup_{\mathbb{Q}} \quad & \sum_{\xi \in \Xi} \mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2 \ \mathbb{Q}(\xi) \\ \text{s.t.} \quad & d(\hat{\mathbb{Q}}^{(N)}, \mathbb{Q}) \leqslant \epsilon. \end{aligned} \tag{9}$$

By the definition of Wasserstein distance,

$$\begin{aligned} \sup_{\mathbb{Q}} \quad & \sum_{\xi \in \Xi} \mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2 \ \mathbb{Q}(\xi) \\ \text{s.t.} \quad & \min_{\Pi} \sum_{\xi; \xi' \in \Xi} |\xi - \xi'| \Pi(\xi; \xi') \leqslant \epsilon \\ & \sum_{\xi \in \Xi} \Pi(\xi; \xi') = \hat{\mathbb{Q}}^{(N)}(\xi') \\ & \sum_{\xi' \in \Xi} \Pi(\xi; \xi') = \mathbb{Q}(\xi). \end{aligned} \tag{10}$$

By eliminating the variable $\mathbb{Q}$, we reduce (9) equivalently to

$$\begin{aligned} \sup_{\Pi} \quad & \sum_{\xi; \xi' \in \Xi} \mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2 \ \Pi(\xi; \xi') \\ \text{s.t.} \quad & \sum_{\xi; \xi' \in \Xi} |\xi - \xi'| \Pi(\xi; \xi') \leqslant \epsilon \\ & \sum_{\xi \in \Xi} \Pi(\xi; \xi') = \hat{\mathbb{Q}}^{(N)}(\xi'), \quad \forall\, \xi' \in \Xi. \end{aligned} \tag{11}$$

The dual optimization of (11) is given by

$$\begin{aligned} \min_{\lambda \geqslant 0} \quad & \lambda \epsilon + \sum_{\xi' \in \Xi} \hat{\mathbb{Q}}^{(N)}(\xi') s(\xi') \\ \text{s.t.} \quad & s(\xi') + \lambda |\xi - \xi'| \geqslant \mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2, \quad \forall\, \xi; \xi' \in \Xi. \end{aligned} \tag{12}$$

It is worth noting that there is no duality gap between (11) and (12) as (11) is essentially a linear programming. We can also write (12) as

$$\begin{aligned} \min_{\lambda \geqslant 0} \quad & \lambda \epsilon + \sum_{\xi' \in \Xi} \hat{\mathbb{Q}}^{(N)}(\xi') s(\xi') \\ \text{s.t.} \quad & s(\xi') \geqslant \max_{\xi \in \Xi} \left[\mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2 - \lambda |\xi - \xi'| \right], \quad \forall\, \xi' \in \Xi \end{aligned}$$

So far, we have reduced the robust formulation using Wasserstein to a simpler form.

5.6 Mathematical Programming for DRE

By expanding the constraint that $\mathbf{x}_i \in \Delta_i$ and writing down (12) in the epigraph form, we have that for Player i,

$$\begin{aligned} \min_{\mathbf{x}_i, \lambda_i \geqslant 0, \eta_i, \{s(\xi')\}_{\xi' \in \Xi}} \quad & \eta_i \\ \text{s.t.} \quad & \lambda_i \epsilon + \sum_{\xi' \in \Xi} \hat{\mathbb{Q}}^{(N)}(\xi') s_i(\xi') \leqslant \eta_i \\ & s_i(\xi') + \lambda_i |\xi - \xi'| \geqslant \sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} C_i(a_1, a_2; \xi) x_1(a_1) x_2(a_2), \\ & \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \forall\, \xi; \xi' \in \Xi \\ & x_i(a_i) \geqslant 0, \quad \forall\, a_i \in \mathcal{A}_i \\ & \sum_{a_i \in \mathcal{A}_i} x_i(a_i) = 1. \end{aligned} \tag{13}$$

The Lagrange function of (13) is given by

$$\begin{aligned} & \mathcal{L}_i(\mathbf{x}_i, \lambda_i, \eta_i, \{s(\xi')\}_{\xi' \in \Xi}, \{\omega_i(\xi, \xi')\}_{\xi, \xi' \in \Xi}, \theta_i, \sigma_i) \\ =& \eta_i + \sum_{\xi; \xi' \in \Xi} \omega_i(\xi, \xi') \left(\mathbf{x}_1^{\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2 - s_i(\xi') - \lambda_i |\xi - \xi'| \right) \\ & \qquad + \theta_i \left(\lambda_i \epsilon + \sum_{\xi' \in \Xi} \hat{\mathbb{Q}}^{(N)}(\xi') s(\xi') - \eta_i \right) + \sigma_i \left(1 - \sum_{a_i \in \mathcal{A}_i} x_i(a_i) \right) \\ =& (1 - \theta_i) \eta_i + \lambda_i \left(\theta_i \epsilon - \sum_{\xi; \xi' \in \Xi} \omega_i(\xi, \xi') |\xi - \xi'| \right) + \sigma_i \\ & \qquad + \sum_{a_i \in \mathcal{A}_i} \left(\sum_{a_{-i} \in \mathcal{A}_{-i}} \sum_{\xi; \xi' \in \Xi} \omega_i(\xi, \xi') C_i(a_1, a_2; \xi) x_{-i}(a_{-i}) - \sigma_i \right) x_i(a_i) \\ & \qquad \qquad + \sum_{\xi' \in \Xi} \left(\theta_i \hat{\mathbb{Q}}^{(N)}(\xi') - \sum_{\xi \in \Xi} \omega_i(\xi, \xi') \right) s_i(\xi'). \end{aligned}$$

Here, $x_i(a_i) \geqslant 0, \lambda_i \geqslant 0$, and η_i and $\{s_i(\xi')\}_{\xi' \in \Xi}$ are free variables. Thus, we need

$$\begin{aligned} 1 - \theta_i &= 0, \\ \theta_i \epsilon - \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi')|\xi - \xi'| &\geqslant 0, \\ \theta_i \hat{\mathbb{Q}}^{(N)}(\xi') - \sum_{\xi \in \Xi} \omega_i(\xi, \xi') &= 0, \\ \sum_{a_{-i} \in \mathcal{A}_{-i}} \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi') C_i(a_1, a_2; \xi) x_{-i}(a_{-i}) &\geqslant \sigma_i. \end{aligned}$$

After some algebraic operations, we can write the dual problem to (13) as

$$\begin{aligned} \max_{\{\omega_i(\xi,\xi') \geqslant 0\}_{\xi,\xi' \in \Xi}, \sigma_i} \quad & \sigma_i \\ \text{s.t.} \quad & \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi')|\xi - \xi'| \leqslant \epsilon, \\ & \hat{\mathbb{Q}}^{(N)}(\xi') = \sum_{\xi \in \Xi} \omega_i(\xi, \xi'), \quad \forall\ \xi \in \Xi, \\ & \sum_{a_{-i} \in \mathcal{A}_{-i}} \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi') C_i(a_1, a_2; \xi) x_{-i}(a_{-i}) \geqslant \sigma_i, \quad \forall\ a_i \in \mathcal{A}_i. \end{aligned}$$

The mathematical problem used to solve for DRE is given by the following,

$$\begin{aligned} \max_{\kappa} \quad & \sum_{i \in \mathcal{I}} (\sigma_i - \eta_i) \\ \text{s.t.} \quad & \lambda_i \epsilon + \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi') s_i(\xi') \leqslant \eta_i, \quad \forall\ i \in \mathcal{I}, \\ & s_i(\xi') + \lambda_i |\xi - \xi'| \geqslant \sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} C_i(a_1, a_2; \xi) x_1(a_1) x_2(a_2), \\ & \qquad \forall\ \xi, \xi' \in \Xi \\ & \sum_{a_i \in \mathcal{A}_i} x_i(a_i) = 1, \quad \forall\ i \in \mathcal{I}, \\ & \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi')|\xi - \xi'| \leqslant \epsilon, \quad \forall\ i \in \mathcal{I}, \\ & \\ & \sum_{a_{-i} \in \mathcal{A}_{-i}} \sum_{\xi;\xi' \in \Xi} \omega_i(\xi, \xi') C_i(a_1, a_2; \xi) x_{-i}(a_{-i}) \geqslant \sigma_i, \quad \forall\ a_i \in \mathcal{A}_i, i \in \mathcal{I}, \\ & \hat{\mathbb{Q}}^{(N)}(\xi') = \sum_{\xi \in \Xi} \omega_i(\xi, \xi'), \quad \forall\ \xi' \in \Xi, i \in \mathcal{I}, \end{aligned} \tag{14}$$

where

$$\kappa = \{x_i(a_i) \geqslant 0, \lambda_i \geqslant 0, \eta_i, \{s_i(\xi')\}_{\xi' \in \Xi}, \{\omega_i(\xi, \xi') \geqslant 0\}_{\xi, \xi' \in \Xi}, \sigma_i\}_{i \in \mathcal{I}}$$

is the collection of decision variables.

Theorem 4. *Solving the mathematical programming above is equivalent to finding the DRE (as defined in (5)) in $\mathcal{G}_\epsilon^{(N)}$.*

Proof. "$\Leftarrow$" Let $(\mathbf{x}_1^*, \mathbf{x}_2^*)$ be an DRE. Then, for $i \in \mathcal{I}$, $\mathbf{x}_i^*$ is the best response to $\mathbf{x}_{-i}^*$. As there is no duality gap between dual and primal, $\sum_{i\in\mathcal{I}} (\sigma_i^* - \eta_i^*) = 0$. Show κ^* is global maximum. We first notice that

$$\sum_{a_1\in\mathcal{A}_1, a_2\in\mathcal{A}_2} \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') C_i(a_1,a_2;\xi) x_2(a_2) x_1(a_1) \geqslant \sigma_i.$$

By the second, the fourth and the fifth constraints in (14),

$$\begin{aligned}\eta_i &= \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') s_i(\xi') + \lambda_i \epsilon \\ &\geqslant \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') s_i(\xi') + \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') \lambda_i |\xi - \xi'| \\ &\geqslant \sum_{a_1\in\mathcal{A}_1} \sum_{a_2\in\mathcal{A}_2} \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') C_i(a_1,a_2;\xi) x_1(a_1) x_2(a_2).\end{aligned}$$

Thus,

$$\eta_i \geqslant \sigma_i.$$

"$\Rightarrow$" Let κ^* be the maximizer of (14). Then, we show that

$$\sigma_i^* \;=\; \eta_i^*, \quad i \in \mathcal{I}. \tag{15}$$

From the first, the second and the fourth constraints, we have

$$\begin{aligned}\sigma_i^* &= \eta_i^* \\ &\geqslant \lambda_i^* \epsilon + \sum_{\xi;\xi'\in\Xi} \omega_i^*(\xi;\xi') \left[\mathbf{x}_1^{*\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2^* - \lambda_i^* |\xi - \xi'| \right] \\ &\geqslant \lambda_i^* \sum_{\xi;\xi'\in\Xi} \omega_i(\xi,\xi') |\xi - \xi'| + \sum_{\xi;\xi'\in\Xi} \omega_i^*(\xi;\xi') \left[\mathbf{x}_1^{*\mathrm{T}} \mathbf{C}_i(\xi) \mathbf{x}_2^* - \lambda_i^* |\xi - \xi'| \right], \\ s^*(\xi') &\geqslant \max_{\xi} \sum_{a_1\in\mathcal{A}_1} \sum_{a_2\in\mathcal{A}_2} C_i(a_1,a_2;\xi) x_1^*(a_1) x_2^*(a_2) - \lambda_i^* |\xi - \xi'|,\end{aligned}$$

and

$$\begin{aligned}\eta_i^* &\geqslant \lambda_i^* \epsilon + \sum_{\xi;\xi'\in\Xi} \omega_i^*(\xi;\xi') s_i^*(\xi') \\ &\geqslant \lambda_i^* \epsilon + \sum_{\xi;\xi'\in\Xi} \omega_i^*(\xi;\xi') \max_{\xi} \sum_{a_1\in\mathcal{A}_1} \sum_{a_2\in\mathcal{A}_2} C_i(a_1,a_2;\xi) x_1^*(a_1) x_2^*(a_2) - \lambda_i^* |\xi - \xi'|.\end{aligned}$$

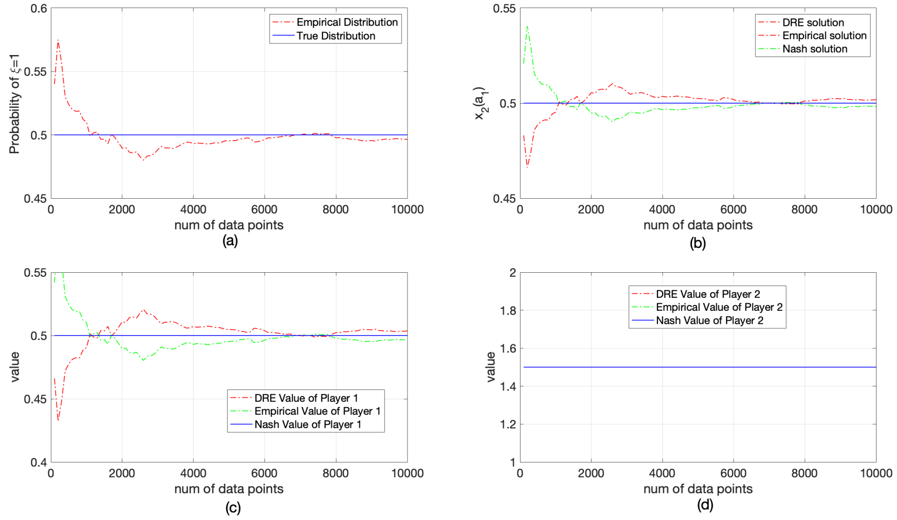

Fig. 2. The comparison of data-driven EG, Nash Game, and data-driven DRG.

From the last constraint,

$$
\begin{aligned}
&\sigma_i^* \\
&\leqslant \sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} \sum_{\xi;\xi' \in \Xi} \omega_i^*(\xi;\xi') C_i(a_1, a_2; \xi) x_1(a_1) x_2^*(a_2) \\
&= \sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} \sum_{\xi' \in \Xi} \mathbb{Q}^{*(N)}(\xi') C_i(a_1, a_2; \xi) x_1(a_1) x_2^*(a_2) \\
&\leqslant \lambda_i^* \epsilon + \sum_{\xi;\xi' \in \Xi} \omega_i^*(\xi;\xi') \left(\sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} C_i(a_1, a_2; \xi) x_1(a_1) x_2^*(a_2) - \lambda_i^* |\xi - \xi'| \right) \\
&\leqslant \lambda_i \epsilon + \sum_{\xi;\xi' \in \Xi} \omega_i^*(\xi;\xi') \max_{\xi} \left(\sum_{a_1 \in \mathcal{A}_1} \sum_{a_2 \in \mathcal{A}_2} C_i(a_1, a_2; \xi) x_1(a_1) x_2^*(a_2) - \lambda_i^* |\xi - \xi'| \right)
\end{aligned}
$$

Thus, we have

$$
\sup_{\mathbb{Q} \in \mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}} \left[c_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*; \xi) \right] \leqslant \sigma_i^* = \eta_i^* \leqslant \sup_{\mathbb{Q} \in \mathcal{B}_\epsilon\left(\hat{\mathbb{Q}}^{(N)}\right)} \mathbb{E}_{\mathbb{Q}} \left[c_i(\mathbf{x}_i, \mathbf{x}_{-i}^*; \xi) \right].
$$

Therefore, $(\mathbf{x}_1^*, \mathbf{x}_2^*)$ is a DRE.

□

6 Numerical Example

Consider a security game which is captured by a nonzero-sum game. The uncertainty parameter ξ represents the security environment, which influences the payoff. Assume that the payoff matrices are given by the following:

$$\begin{bmatrix} (1+\xi,3) & (0,2) \\ (2,0) & (-1,1) \end{bmatrix},$$

where $\xi \in \Xi = \{-1,1\}$. The true distribution of ξ is $\mathbb{P}[\xi = 1] = 1/2$, and $\mathbb{P}[\xi = 1] = 1/2$. When there is perfect distribution information (both players know the true distribution), the Nash equilibrium is $(1/2, 1/2)$, and the expected value of the game is $(1/2, 3/2)$. For the distributionally robust case, let the radius of Wasserstein ball be $\epsilon_N = 1/N$. As illustrated in Fig. 2, the value and equilibrium of both data-driven EG and data-driven DRG converge to the ones in Nash Game. We notice that the strategy of Player 1 and the value of Player 2 stay unaltered. By the indifferent principle [13], the bimatrix game considered here is fully mixed. In this case, the strategy of Player 1 only depends the payoff matrix of Player 2, and the value of Player 2 only depends on the strategy of Player 1.

7 Conclusions and Future Work

7.1 Conclusions

In this paper, we have proposed a new type of data-driven game model in which the players are capable of exploiting the information in the dataset. We have adopted the distributionally robust formulation to address the issue arising from the curse of the optimizer. We have used Wasserstein ball as the ambiguity set with the empirical distribution centered. By tuning the radius of the Wasserstein ball, we have demonstrated the relations between the proposed game and the existing games. We have also given the mathematical programming whose solutions are a subset of data-driven DRE.

7.2 Future Work

1. **Data-Driven Distributionally Robust Bayesian Game** In this work, we did not consider the case where there exists private information. As Harsanyi pointed out, the incomplete information is quite involving as there is *belief hierarchy*. We can use the information from the dataset to form the player's beliefs. As the belief are not generated from the common prior, the players are suspicious about the data-based belief. Therefore, it is reasonable to introduce robustness.
2. **Data-Driven Dynamic Game** It is also possible to extend the data-driven dynamic game. In a dynamic system, the agents do not have perfect nor complete knowledge of the system. While making decisions, they observe the outcomes of the system and update their knowledge. And with the updated knowledge, the agents are able to make *better* decisions.
3. **One-Sided Data-Driven Game** Consider a two-player game. One player has the access to the dataset $\hat{\Xi}$ and the other player has the access to the dataset $\tilde{\Xi}$. If $\tilde{\Xi} \subseteq \hat{\Xi}$, this becomes a one-sided information game, in which one player has more information than the other [7].

References

1. Aghassi, M., Bertsimas, D.: Robust game theory. Math. Program. **107**(1–2), 231–273 (2006). https://doi.org/10.1007/s10107-005-0686-0
2. Aumann, R.J.: Correlated equilibrium as an expression of Bayesian rationality. Econometrica J. Econometric Soc. **55**, 1–18 (1987)
3. Bauso, D., Gao, J., Tembine, H.: Distributionally robust games: f-divergence and learning. In: Proceedings of the 11th EAI International Conference on Performance Evaluation Methodologies and Tools, pp. 148–155 (2017)
4. Esfahani, P.M., Kuhn, D.: Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. Math. Program. **171**(1–2), 115–166 (2018). https://doi.org/10.1007/s10107-017-1172-1
5. Farooq, M.J., Zhu, Q.: On the secure and reconfigurable multi-layer network design for critical information dissemination in the internet of battlefield things (IoBT). IEEE Trans. Wirel. Commun. **17**(4), 2618–2632 (2018)
6. Harsanyi, J.C.: Games with incomplete information played by "Bayesian" players part ii. Bayesian equilibrium points. Manag. Sci. **14**(5), 320–334 (1968)
7. Horák, K., Bošanskỳ, B., Péchouček, M.: Heuristic search value iteration for one-sided partially observable stochastic games. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, pp. 558–564 (2017)
8. Kakutani, S., et al.: A generalization of Brouwer's fixed point theorem. Duke Math. J. **8**(3), 457–459 (1941)
9. Koçyiğit, Ç., Iyengar, G., Kuhn, D., Wiesemann, W.: Distributionally robust mechanism design. Manage. Sci. **66**(1), 159–189 (2020)
10. Liang, A.: Games of incomplete information played by statisticians. arXiv preprint arXiv:1910.07018 (2019)
11. Liu, S., Zhu, Q.: Robust and stochastic optimization with a hybrid coherent risk measure with an application to supervised learning. IEEE Control Syst. Lett. **5**(3), 965–970 (2020)
12. Loizou, N.: Distributionally robust game theory. arXiv preprint arXiv:1512.03253 (2015)
13. Maschler, M., Solan, E., Zamir, S.: Game Theory (translated from the Hebrew by Ziv Hellman and edited by Mike Borns), pp. xxvi 979, 4. Cambridge University Press, Cambridge (2013)
14. Peng, G., Zhu, Q.: Game-theoretic analysis of optimal control and sampling for linear stochastic systems. In: 2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 647–654. IEEE (2019)
15. Peng, G., Zhu, Q.: Sequential hypothesis testing game. In: 2020 54th Annual Conference on Information Sciences and Systems (CISS), pp. 1–6. IEEE (2020)
16. Provost, F., Fawcett, T.: Data science and its relationship to big data and data-driven decision making. Big data **1**(1), 51–59 (2013)
17. Singh, V.V., Jouini, O., Lisser, A.: Distributionally robust chance-constrained games: existence and characterization of Nash equilibrium. Optim. Lett. **11**(7), 1385–1405 (2017). https://doi.org/10.1007/s11590-016-1077-6
18. Smith, J.E., Winkler, R.L.: The optimizer's curse: skepticism and postdecision surprise in decision analysis. Manag. Sci. **52**(3), 311–322 (2006)
19. Villani, C.: Optimal Transport: Old and New, vol. 338. Springer, Heidelberg (2008)
20. Zhang, Q., Zhu, Q., Boutaba, R.: Dynamic resource allocation for spot markets in cloud computing environments. In: 2011 Fourth IEEE International Conference on Utility and Cloud Computing, pp. 178–185. IEEE (2011)