


Data-Driven Behaviour Estimation in Parametric Games

Conference Paper**Author(s):**

Maddux, Anna M.; Pagan, Nicolò; Belgioioso, Giuseppe ; Dörfler, Florian 

Publication date:

2023-11-22

Permanent link:

<https://doi.org/10.3929/ethz-b-000660774>

Rights / license:

[Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International](#)

Originally published in:

IFAC-PapersOnLine 56(2), <https://doi.org/10.1016/j.ifacol.2023.10.220>

Funding acknowledgement:

180545 - NCCR Automation (phase I) (SNF)

Data-Driven Behaviour Estimation in Parametric Games[★]

Anna M. Maddux^{*} Nicolò Pagan^{**} Giuseppe Belgioioso^{***}
Florian Dörfler^{***}

^{*} *Sycamore, EPF Lausanne, Switzerland, (e-mail: anna.maddux@epfl.ch).*

^{**} *Social Computing Group, UZH, Switzerland, (e-mail: nicolo.pagan@uzh.ch)*

^{***} *Automatic Control Laboratory, ETH Zürich, Switzerland, (e-mail: {gbelgioioso,doerfler}@ethz.ch)*

Abstract: A central question in multi-agent strategic games deals with learning the underlying utilities driving the agents' behaviour. Motivated by the increasing availability of large datasets, we develop an unifying data-driven technique to estimate agents' utility functions from their observed behaviour, irrespective of whether the observations correspond to equilibrium configurations or to temporal sequences of action profiles. Under standard assumptions on the parametrization of the utilities, the proposed inference method is computationally efficient and finds all the parameters that rationalize the observed behaviour best. We numerically validate our theoretical findings on the market share estimation problem under advertising competition, using historical data from the Coca-Cola Company and Pepsi Inc. duopoly.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Game theory; Multi-agent systems; Data-Driven Decision Making

1. INTRODUCTION

The study of multi-agent systems (MAS) has been of interest to various research communities, such as economics, e.g. Bertrand–Nash competition (Narahari et al., 2009), computational social science, e.g. social networks formation (Pagan and Dörfler, 2019), and automatic control, e.g. power systems (Belgioioso et al., 2022). The decision-making processes in such MAS can be modelled as a strategic game, where, typically, the objective of each agent is to select an action that maximizes its own utility.

A pivotal problem in game theory deals with predicting the behaviour of agents from their utilities, namely, characterizing the game equilibria. In many relevant cases, though, the underlying utilities are not directly observable. With the emergence of the internet of things, however, increasing amounts of data regarding the behaviour of each agent are available and can be used to infer the underlying utilities that rationally led to that behaviour. This line of research, termed *inverse game theory*, has many useful applications: Inferred utilities can be used to predict the agents' behaviour in other similarly structured games, where the behaviour is unobservable (Vaughan et al., 2011), or to cross-validate results from classical game theory, namely, to confirm whether a specified game admits any utilities that are consistent with the observed behaviour (Kuleshov and Schrijvers, 2015).

Existing data-driven methods for inverse game theory can be broadly divided into two sub-classes: methods based on observing time-independent snapshots of a MAS, and

methods based on action trajectories observations. Among works of the former class, (Pagan and Dörfler, 2019) and (Kuleshov and Schrijvers, 2015) propose an inverse-optimization approach to determine the utilities from an observed Nash and correlated equilibrium, respectively. (Bertsimas et al., 2015) combine ideas from inverse optimization and variational inequality (VI) theory to build a (non-) parametric estimation method. (Li et al., 2019) and (Trivedi et al., 2020) propose a data-driven framework specific to matchings and network formation, respectively.

The latter sub-class of methods (inverse-game-theoretic problems with action trajectories observations) has mostly been studied by the automatic control community. (Tsai et al., 2016) and (Molloy et al., 2017) propose an inference method based on the minimum principle of optimal control for estimating the unknown parameters of agents' cost functions from observed open-loop Nash equilibrium in two-player zero-sum and N-player dynamic games, respectively. (Inga et al., 2019) apply the principle of maximum entropy in N-player inverse dynamic games to infer the cost functions from both observed open-loop and feedback Nash equilibria. Alternatively, (Nekipelov et al., 2015) develop a regret minimization method for auction games to infer players' private valuations (parameters).

From the aforementioned literature, several critical issues emerge. First, the objective is to find a parameter tuple which rationalizes the observed behaviour best¹. However, the inverse game-theoretic problem is, in general, ill-posed, in the sense that possibly multiple parameter tuples exist which simultaneously are the most rational estimates of

[★] The authors gratefully acknowledge financial support from ETHZ and the SNSF under the NCCR Automation grant 180545.

¹ Under certain circumstances (Bertsimas et al., 2015; Pagan and Dörfler, 2019) find sets (not singletons) of parameters.

the observed behaviour. Moreover, all previously presented works are based on the assumption that the available observations are either snapshots or trajectories of action profiles, but always equilibrium configurations with the exception of (Nekipelov et al., 2015). In reality, however, interactions between agents may evolve over time, sometimes generating persistent fluctuations and equilibrium behaviour may not be observed.

Motivated by these issues, in this paper, we develop a novel unifying data-driven behavior estimation method to learn parametric utilities in multi-agent games. The contributions are summarized next.

- (i) Our inference method is unifying, in the sense that it handles observations of both snapshots and of trajectories of action profiles in a multi-agent game. It builds on the assumptions that any observed snapshot corresponds to an approximate Nash equilibrium and that any observed action trajectory is compatible with an approximate better-response dynamics, i.e., in which each agent aims at improving its utility with respect to its previous action.
- (ii) Our method leverages a computationally tractable linear programming reformulation of the original inference problem to produce polyhedral sets containing all the parameters that rationalize the observed game behaviour best.
- (iii) We numerically validate our method on the market share evolution estimation problem under advertising competition (Gasmi et al., 1988), using historical data of advertising expenditures trajectories from the Coca-Cola Company and Pepsi Inc. duopoly.

2. MATHEMATICAL SETUP

We consider a set of players (or agents) $\mathcal{I} := \{1, \dots, N\}$, where each agent $i \in \mathcal{I}$ chooses its action (or strategy) a_i from its action set \mathcal{A}_i . Let $a := (a_1, \dots, a_N) \in \mathcal{A}$ be the action profile, where $\mathcal{A} := \times_{i \in \mathcal{I}} \mathcal{A}_i$ is the global action set. Each agent $i \in \mathcal{I}$ aims at maximizing its own utility $U_i(a_i, a_{-i})$, which depends on its action $a_i \in \mathcal{A}_i$ and on the actions of the other agents $a_{-i} := (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_N) \in \mathcal{A}_{-i}$, where $\mathcal{A}_{-i} := \times_{j \in \mathcal{I} \setminus \{i\}} \mathcal{A}_j$. The collection of players, strategy sets, and utility functions, i.e., $\mathcal{G} := (\mathcal{I}, (\mathcal{A}_i)_{i \in \mathcal{I}}, (U_i)_{i \in \mathcal{I}})$, constitutes a game.

We assume that agents are rational decision makers. Namely, if an agent i is allowed to revise its current action \bar{a}_i , it will choose an action $a_i \in \mathcal{A}_i$ that leads to a larger payoff, i.e., a so-called *better response*.

Definition 1. For all $i \in \mathcal{I}$, an action $a_i \in \mathcal{A}_i$ is an ϵ -*better response* to $\bar{a} = (\bar{a}_i, \bar{a}_{-i}) \in \mathcal{A}_i \times \mathcal{A}_{-i}$, with $\epsilon \geq 0$, if

$$U_i(a_i, \bar{a}_{-i}) \geq U_i(\bar{a}_i, \bar{a}_{-i}) - \epsilon. \quad (1)$$

Accordingly, the ϵ -better response mapping is defined as

$$\mathcal{R}_i^\epsilon(\bar{a}) := \{a_i \in \mathcal{A}_i \mid (1) \text{ holds}\}.$$

If (1) holds with $\epsilon = 0$, a_i is known as better response. \square

An action a_i is a *best response* to \bar{a}_{-i} , if it is a better response compared to any alternative admissible action $\bar{a}_i \in \mathcal{A}_i$.

Definition 2. For all $i \in \mathcal{I}$, an action $a_i \in \mathcal{A}_i$ is an ϵ -*best response* to $\bar{a}_{-i} \in \mathcal{A}_{-i}$, with $\epsilon \geq 0$, if

$$U_i(a_i, \bar{a}_{-i}) \geq U_i(\bar{a}_i, \bar{a}_{-i}) - \epsilon, \quad \forall \bar{a}_i \in \mathcal{A}_i. \quad (2)$$

Accordingly, the ϵ -best response mapping is defined as

$$\mathcal{B}_i^\epsilon(\bar{a}_{-i}) := \{a_i \in \mathcal{A}_i \mid (2) \text{ holds}\}.$$

If (2) holds with $\epsilon = 0$, a_i is known as *best response*. \square

If all agents play an ϵ -best response action simultaneously, this action profile is an ϵ -Nash equilibrium.

Definition 3. An action profile $a^* = (a_1^*, \dots, a_N^*)$ is an ϵ -Nash equilibrium (ϵ -NE), with $\epsilon \geq 0$, if

$$\forall i \in \mathcal{I}: \quad a_i^* \in \mathcal{B}_i^\epsilon(a_{-i}^*). \quad (3)$$

If (3) holds with $\epsilon = 0$, a^* is a Nash equilibrium (NE). \square

A NE is strategically-stable in the sense that no agent has an interest in unilaterally deviating. Next, we assume that the local utility functions admit a finite parametrization.

Assumption 1. For player $i \in \mathcal{I}$, the utility function $U_i(a_i, a_{-i}; \theta_i)$ is parameterized by $\theta_i \in \Theta_i$, where $\Theta_i \neq \emptyset$ is the parameter space. Additionally, Definitions 1-3 hold with respect to a specified parameterization. \square

In this paper, we assume that the utilities are unknown but observations of the agents' actions are available. Thus, we focus on the problem of inferring the underlying parameters θ_i that led the system to the observed actions. Concretely, we consider two types of observations:

- (i) Dynamic observations: a temporal sequence of $m \in \mathbb{N}$ action profiles, $(a(j))_{j \in [m]}$, where $a(j) \in \mathcal{A}$ is the action profile at time step j , and $[m] := \{1, \dots, m\}$.
- (ii) Static observations: an equilibrium profile $a^* \in \mathcal{A}$.

We illustrate the difference with an example: in economics, competing firms repeatedly change their prices in response to other firms. In this case, a sequence of action profiles, i.e., prices, is observed, and the firms' utility functions can be inferred by assuming that pairwise consecutive elements of the sequence are approximately compatible with better responses, i.e., $a_i(j+1) \in \mathcal{R}_i^{\epsilon_i}(a(j))$, for some $\epsilon_i \geq 0$. On the other hand, if the firms have reached an equilibrium, their utility functions can still be inferred by assuming that the observed action profile a^* is approximately a NE, namely, $a_i^* \in \mathcal{B}_i^{\epsilon_i}(a_{-i}^*)$ for all $i \in \mathcal{I}$. The term "approximately" stems from the fact that the agents may not act fully rational. Namely, the observed action trajectories and equilibrium profiles may violate some of the perfect better-response (1) and NE conditions (3), respectively.

3. DATA-DRIVEN BEHAVIOR ESTIMATION

Before proceeding, this next assumption helps to unify the notation for the static and dynamic cases.

Assumption 2. For all $i \in \mathcal{I}$, we construct a set of m data points $\mathcal{D}_i := \{(a_i^j, \bar{a}^j)\}_{j \in [m]}$, where $a_i^j \in \mathcal{A}_i$ and $\bar{a}^j = (\bar{a}_i^j, \bar{a}_{-i}^j) \in \mathcal{A}$.

- (i) If the observations are dynamic, the data points are of the form $(a_i^j, \bar{a}^j) = (a_i(j+1), a(j))$, for all time steps $j \in [m]$, and $a_i(j+1) \in \mathcal{R}_i^{\epsilon_i}(a(j))$, for all $j \in [m]$, w.r.t. the true parameterization θ_i of $U_i(\cdot, \cdot; \theta_i)$ for some $\epsilon_i \geq 0$.

(ii) If the observations are static, the data points are of the form $(a_i^j, \bar{a}^j) = (a_i^*, (\bar{a}_i^j, a_{-i}^*))$, where $\bar{a}_i^j \in \mathcal{A}_i$, for all $j \in [m]$, are m alternative admissible actions of player i . Furthermore, $a_i^* \in \mathcal{B}_i^{\epsilon_i}(a_{-i}^*)$ w.r.t. to the true parameterization θ_i of $U_i(\cdot, \cdot; \theta_i)$ for some $\epsilon_i \geq 0$. \square

This assumption states that any observed action trajectory $(a_i(j))_{j \in [m]}$ is an ϵ -better response dynamics and any observed action profile a^* is an ϵ -NE for the true parameterization of the utility function. Then, our objective is to find all admissible preference vectors θ_i that fully rationalize the observed behaviour, namely, the action trajectories or the NE. Alternatively, if no such vector exists, we aim to find all the vectors θ_i that violate the better response or NE conditions the least. To measure the irrationality, we introduce the error functions $e_i : \mathcal{A}_i \times \mathcal{A} \times \Theta_i \rightarrow \mathbb{R}$, defined as in (Pagan and Dörfler, 2019)

$$e_i(a_i, \bar{a}; \theta_i) := U_i(\bar{a}_i, \bar{a}_{-i}; \theta_i) - U_i(a_i, \bar{a}_{-i}; \theta_i), \quad (4)$$

which takes on positive values whenever agent i 's action a_i yields a lower utility than \bar{a}_i , given that the other agents' actions \bar{a}_{-i} remains fixed. For dynamic observations, θ_i is a rational preference if and only if $a_i(j+1)$ leads to higher utility compared to $a_i(j)$, when $a_{-i}(j)$ remains fixed. For static observations, instead, we expect that a_i^* leads to higher payoffs compared to any other admissible alternative $\bar{a}_i \in \mathcal{A}_i$. Regardless of whether the observations are evolving or static, if $e_i(a_i, \bar{a}; \theta_i) \leq 0$ then θ_i is a rational behavior. Vice-versa, the irrationality of θ_i comes from the positive contributions of the error function in (4): $e_i^+(a_i, \bar{a}; \theta_i) := \max\{0, e_i(a_i, \bar{a}; \theta_i)\}$.

To measure an agent's maximal violation of the rationality assumption, we introduce a function,

$$d_i(\theta_i; \mathcal{D}_i) := \max_{j \in [m]} e_i^+(a_i^j, \bar{a}^j; \theta_i), \quad (5)$$

termed the irrationality loss function. Any preference vector $\hat{\theta}_i$ which minimizes $d_i(\cdot; \mathcal{D}_i)$ ensures that the worst violation of the rationality assumption over all m observations \mathcal{D}_i is minimal and $\hat{\theta}_i$ rationalizes the observed individual behaviour best. Hence, our inference problem can be cast as N decoupled optimization problems:

Problem 1. Let assumptions 1-2 hold. For all $i \in \mathcal{I}$, find all the vectors of preferences $\hat{\theta}_i$ that solve

$$\begin{cases} \min_{\theta_i \in \mathbb{R}^{p_i}} d_i(\theta_i; \mathcal{D}_i) \\ \text{s.t. } \theta_i \in \Theta_i. \end{cases} \quad (6)$$

Next, we show that the observed trajectory of action profiles $(a_i(j))_{j \in [m]}$ can be interpreted as an ϵ -better response dynamics for each agent i and, if the available data points \mathcal{D}_i in Assumption 2 (ii) are exhaustive, i.e., \mathcal{D}_i contains all feasible actions of agent i , the observed static profile a^* can be interpreted as an ϵ -NE³, where ϵ is the maximum value among the solutions of the N problems in (6).

Proposition 1. Let assumptions 1-2 hold true.

² The equilibrium profile (a_i^*, a_{-i}^*) is observed while $\bar{a}_i^j \in \mathcal{A}_i$, $j \in [m]$, are fictitious admissible actions agent i could have played.

³ Note that we make no assumptions about the utility functions w.r.t. the actions. Thus, if the game is ill-posed, i.e., a NE does not exist, then our method finds a preference vector that violates the NE conditions the least.

- (i) For $i \in \mathcal{I}$, the observed action trajectory, $(a_i(j))_{j \in [m]}$ is an $\bar{\epsilon}$ -better response dynamics w.r.t. $U_i(\cdot, \cdot; \hat{\theta}_i)$, where $\bar{\epsilon} := \max_{i \in \mathcal{I}} d_i(\hat{\theta}_i; \mathcal{D}_i)$, and $\hat{\theta}_i$ as in (6).
(ii) Assume that $\bigcup_{i=1}^m \bar{a}_i^j = \mathcal{A}_i$, for all $i \in \mathcal{I}$. Then, the observed action profile a^* is an $\bar{\epsilon}$ -NE w.r.t. $U_i(\cdot, \cdot; \hat{\theta}_i)$, where $\bar{\epsilon} := \max_{i \in \mathcal{I}} d_i(\hat{\theta}_i; \mathcal{D}_i)$, and $\hat{\theta}_i$ as in (6) for all $i \in \mathcal{I}$. \square

For the proof we refer the reader to Appendix A in (Maddux et al., 2022).

3.1 Behavior estimation in contextual games

In this section, we extend the formulation of Problem 1 to the case of multiple data-sets, i.e., multiple trajectories or equilibria observations. In many relevant scenarios, (the outcome of) multi-agent games depend on exogenous but observable factors (e.g. GDP, shocks, seasonality in Bertrand–Nash games). More formally, we denote such factors as the context $\xi \in \Xi$ of the game, where $\Xi \subseteq \mathbb{R}^c$ is the set of possible contexts (Sessa et al., 2020). Hereafter, we assume that each agent's parametric utility depends also on a known context ξ , as formalized next.

Assumption 3. For all $i \in \mathcal{I}$, the local utility function $U_i(a_i, a_{-i}; \theta_i, \xi)$ depends also on an observable parameter (context) $\xi \in \Xi$. \square

To generalize the inference problem in Section 2, we assume to have n observations, i.e., action profile trajectories $(a^k(j))_{j \in [m^k]}$ or equilibrium profiles $a^{*,k}$, each connected to a different game context ξ^k , with $k \in [n]$. As in Section 2, we use the following assumption to construct a data-set that helps to unify the notation for the static and dynamic cases.

Assumption 4. Let $\{\xi^k\}_{k \in [n]}$, with $\xi^k \in \Xi$, be a set of contexts of a multi-agent game $\mathcal{G} = (\mathcal{I}, (\mathcal{A}_i)_{i \in \mathcal{I}}, (U_i)_{i \in \mathcal{I}}, \Xi)$. For each context ξ^k , we construct n sets of data points $\{\mathcal{D}_i^k\}_{i \in \mathcal{I}}$ as in Assumption 2, where $\mathcal{D}_i^k := \{(a_i^{j,k}, \bar{a}^{j,k})\}_{j \in [m^k]}$ and m^k is the number of data points relative to the k -th context. \square

The multi data-sets extension of the original behaviour estimation problem (Problem 1) can be recast as N decoupled optimization problems, as formalized next.

Problem 2. Let assumptions 1-4 hold true. For all $i \in \mathcal{I}$, find all vectors of preferences $\hat{\theta}_i$ that solve

$$\begin{cases} \min_{\theta_i \in \mathbb{R}^{p_i}} \max_{k \in [n]} d_i(\theta_i; \mathcal{D}_i^k, \xi^k) \\ \text{s.t. } \theta_i \in \Theta_i, \end{cases} \quad (7)$$

where the distance function⁴ d_i in (7) is now parametric also with respect to the game context ξ^k , i.e.,

$$d_i(\theta; \mathcal{D}_i^k, \xi^k) := \max_{j \in [m^k]} e_i^+(a_i^{j,k}, \bar{a}^{j,k}; \theta_i, \xi^k). \quad (8)$$

3.2 A linear programming approach

Before proceeding, we assume that the utilities are linearly parametrized, as in (Waugh et al., 2011; Kuleshov and Schrijvers, 2015; Pagan and Dörfler, 2019).

⁴ With some abuse of notation, we keep the same notation for the functions e_i , e_i^+ and d_i in contextual games.

Assumption 5. For all $i \in \mathcal{I}$, the function $U_i(a_i, a_{-i}; \theta_i, \xi)$ is linear in its parameters $\theta_i \in \Theta_i$. \square

It follows that d_i in (8) is piece-wise linear, thus convex, in its parameters θ_i . Therefore, the minimax optimization in Problem 2 can be equivalently recast as a linear program (LP), see e.g. § 1 in (Ahuja, 1985), which is always feasible.

Proposition 2. Let the preconditions of Problem 2 hold. For all $i \in \mathcal{I}$, vector $\hat{\theta}_i$ solves (7) if and only if there exists a constant $\hat{\epsilon}_i$ such that the pair $(\hat{\theta}_i, \hat{\epsilon}_i)$ solves the LP

$$\begin{cases} \min_{\epsilon_i, \theta_i} \epsilon_i \\ \text{s.t. } e_i(a_i^{j,k}, \bar{a}^{j,k}; \theta_i, \xi^k) \leq \epsilon_i, \forall j \in [m^k], \forall k \in [n], \\ \epsilon_i \geq 0, \\ \theta_i \in \Theta_i. \end{cases} \quad (9)$$

Lemma 1. For all $i \in \mathcal{I}$, the solution set of the LP in (9) is nonempty. \square

For the proof we refer the reader to Appendix B in (Maddux et al., 2022).

Finally, by leveraging the LP reformulation in Proposition 2, we can characterize the solution set of the original multi data-set behavior estimation problem (Problem 2).

Proposition 3. Let assumptions 1-4 hold true. For all $i \in \mathcal{I}$, the solution set $\hat{\Theta}_i$ of (7) reads as

$$\hat{\Theta}_i := \left\{ \theta_i \in \Theta_i \mid e_i(a_i^{j,k}, \bar{a}^{j,k}; \xi^k, \theta_i) \leq \hat{\epsilon}_i, \forall j \in [m^k], \forall k \in [n] \right\}, \quad (10)$$

where $\hat{\epsilon}_i$ is the optimal value of the LP in (9). \square

Note that $\hat{\Theta}_i$ is a polyhedron which contains all parameters θ_i that rationalize best the observed behaviour of agent i over all contexts ξ^k . If $\hat{\epsilon}_i = 0$, then $\max_{k \in [n]} d_i(\hat{\theta}_i; \mathcal{D}_i^k, \xi^k) = 0$ for all $\hat{\theta}_i \in \hat{\Theta}_i$, which implies that agent i 's observed behaviour is perfectly rationalizable.

4. NUMERICAL AND EMPIRICAL ANALYSIS

In this section, we numerically validate our inference method on (a) the demand estimation problem under Bertrand–Nash competition, using static observations of NE prices, and on (b) the market share evolution estimation problem under advertising competition (Gasmi et al., 1988), using dynamic observations constituted by historical data of advertising expenditures.

4.1 Demand estimation in Bertrand–Nash competition

Consider two firms competing to maximize their revenues by setting prices p_1 and p_2 , respectively. The demand of each firm $i \in \{1, 2\}$ is modelled by the parametric function

$$D_i(p_1, p_2; \theta_i, \xi) = \theta_{i,0} + \theta_{i,1}p_1 + \theta_{i,2}p_2 + \theta_{i,3}\xi, \quad (11)$$

where ξ represents time and region specific economic indicators (e.g. shocks, seasonality). We assume $\theta_{i,i} \leq 0$ and $\theta_{i,-i} \geq 0$ such that, for firm i , an increase of price p_i has a negative effect, and an increase of price p_{-i} has a positive effect on its demand ($\theta_{i,-i} = \theta_{1,2}$, if $i = 1$ and,




Fig. 1. The true marginal revenue function (blue line), the inverse-VI fit (Bertsimas et al., 2015) (black line, square markers), and the fit obtained from the polyhedral solution set which reconcile all the data exactly (grey region) for both firms.

$\theta_{i,-i} = \theta_{2,1}$, if $i = 2$). Then, given context ξ , each firm sets its prices to maximize its revenue, i.e., $U_i(p_1, p_2; \theta_i, \xi) = p_i \cdot D_i(p_1, p_2; \theta_i, \xi)$.

Given a set of observed contexts $\{\xi^k\}_{k \in [n]}$, we seek to estimate the most rational normalized⁵ parameters $\tilde{\theta}_i := (\frac{\theta_{i,0}}{\theta_{i,-i}}, \frac{\theta_{i,i}}{\theta_{i,-i}}, \frac{\theta_{i,3}}{\theta_{i,-i}})^\top$ from a set of observed price tuples $\{(p_1^{*,k}, p_2^{*,k})\}_{k \in [n]}$, each connected to a different context ξ^k and assumed to be a NE. Thus, we construct n batches of data-sets $D_i^k := \{(p_1^{*,k}, \tilde{p}_i^{j,k}, p_{-i}^{*,k})\}_{j \in [m^k]}$, where $\{(\tilde{p}_1^{j,k}, \tilde{p}_2^{j,k})\}_{j \in [m^k]}$ is a set of alternative feasible price profiles. This demand estimation problem satisfies the assumptions of our framework. Thus, its solutions can be found by exploiting Proposition 3.

For our numerical study, we assume the true parameters $\theta_1^* = (1, -1.2, 0.5, 1)^\top$ and $\theta_2^* = (1, 0.3, -1, 1)^\top$ of the demand function (11) to be unknown. Drawing inspiration from § 8.1 in (Bertsimas et al., 2015), we model contexts $\xi^k \in \{\xi^k\}_{k \in [50]}$ as i.i.d. normally distributed random variables with mean 5 and standard deviation 1.5. and solve for the NE prices $(p_1^{*,k}, p_2^{*,k})$ connected to context ξ^k . For each agent i and each context ξ^k , we estimate $\tilde{\theta}_i$ from data-set⁶ $D_i^k := \{(p_i^{*,k}, \tilde{p}_i^{j,k}, p_{-i}^{*,k})\}_{j \in [2^7+1]}$. To further validate our method, we compare its results with those obtained via the VI-based method proposed in § 3.1 of (Bertsimas et al., 2015). The polyhedral solution set $\hat{\Theta}_1$ of firm 1 contains θ_1^* and θ_1^{VI} . Since for each context ξ^k , the price pair $(p_1^{*,k}, p_2^{*,k})$ is a NE, both methods unsurprisingly find parameter vectors that perfectly rationalize the observed prices. However, the solution of the inverse VI problem θ_1^{VI} does not coincide with the true parameter vector θ_1^* . On the other hand, solving Problem 2 gives all the parameter vectors that perfectly rationalize the generated prices, including the true parameter vector θ_1^* .

We plot the fitted marginal revenue function $MR_1 = \frac{\partial}{\partial p_1} p_1 \cdot D_1(p_1, p_2; \theta_i, \xi)$ of firm 1 obtained from the vertices of the polyhedral solution sets $\hat{\Theta}_1$, see Figure 1. The grey area displays all potential marginal revenue functions in the parametric family specified in (11) which exactly reconcile the simulated NE prices. Note that it contains the true

⁵ Normalization ensures that the everywhere-zero revenue function does not let all prices appear rational (Ng et al., 2000).

⁶ $\tilde{p}_i^{j,k}$ is the j -th point in a regular price grid $[0, p_{\max}]$ of $2^7 + 1$ points, which can be seen as a discretization of the price sets \mathcal{A}_i

Table 1. Rationalization error ϵ_i and estimated beliefs $k_{1,i}, k_{2,i}$ obtained via our method (L_∞), the L_2 -based variant, and OLS regression.

	Coke			Pepsi		
	$\hat{\epsilon}_1$	$\hat{k}_{\{1,1\}}$	$\hat{k}_{\{2,1\}}$	$\hat{\epsilon}_2$	$\hat{k}_{\{1,2\}}$	$\hat{k}_{\{2,2\}}$
L_∞	0.018	0.066	0.057	0.011	0.036	0.022
L_2	0.021	0.078	0.052	0.013	0.021	0.016
OLS	0.059	0.011	0.007	0.015	0.012	0.011

marginal revenue function and the fitted function obtained from (Bertsimas et al., 2015) since both reconcile the data.

4.2 Estimating market shares in advertising competition

In this subsection, we formulate the problem of estimating the beliefs of two firms in an advertising competition on the market share evolution. Under the assumption that the trajectory of observed advertising expenditure profiles $(a_1(j), a_2(t))_{t \in [T]}$ is compatible with an approximate better-response dynamics (Assumption 2), we can solve the belief estimation problem using Proposition 3.

We model the market share evolution via the parametric Lanchester model (LM) of advertising competition in (Chintagunta and Vilcassim, 1992). Its discrete-time specification, taken from Eqn. (2) in (Chintagunta and Vilcassim, 1992), is given by

$$M(t+1) = M(t) + k_1 \sqrt{a_1(t+1)}(1 - M(t)) - k_2 \sqrt{a_2(t+1)}M(t), \quad (12)$$

where M is the market share of firm 1, $(1 - M)$ is the share of firm 2, and $k_1, k_2 \geq 0$ are measures of the advertising effectiveness of firm 1 and firm 2, respectively. To include word-of-mouth effects on the market share evolution, (Sorger, 1989) suggests a modification of the LM, as shown in Eqn. (4) of (Chintagunta and Jain, 1995)

$$M(t+1) = M(t) + k_1 \sqrt{a_1(t+1)}\sqrt{1 - M(t)} - k_2 \sqrt{a_2(t+1)}\sqrt{M(t)}. \quad (13)$$

In a duopolistic market share competition, both firms update their advertising expenditure to gain market share. Moreover, each firm's choice of advertising expenditure depends on its beliefs on its own and the competitor's advertising effectiveness. Let us denote by $\theta_1 = (k_{1,1}, k_{2,1})$ and $\theta_2 = (k_{1,2}, k_{2,2})$ firm 1's and firm 2's beliefs on their true advertising effectiveness k_1 and k_2 , respectively. We assume that firm 1 sets its new advertising expenditure, $a_1(t+1)$, with the objective of increasing its current market share, i.e., $M(a_1(t+1), a_2(t); \theta_1) \geq M(a_1(t), a_2(t); \theta_1)$. Similarly, we assume that firm 2 sets $a_2(t+1)$ such that $1 - M(a_1(t), a_2(t+1); \theta_2) \geq 1 - M(a_1(t), a_2(t); \theta_2)$. Thus, $\mathcal{D}_i := \{(a_i(t+1), a_i(t), a_{-i}(t))\}_{t \in [T]}$, constitutes the data-set of agent i , from which we estimate $\hat{\theta}_i$ such that the observed advertising expenditure trajectory is compatible with a better-response dynamics.

Empirical analysis In this section, we show the applicability of our method to a real-world data-set. The Coca-Cola Company (firm 1) and Pepsi Inc. (firm 2) are amongst the most famous examples of duopolistic rivalries over market share, while other soft-drinks companies own a negligible fraction of market shares. We assume Coke and




Fig. 2. Coke's and Pepsi's irrationality, expressed by violations of the better-response conditions in (1), over the time horizon. Resulting irrationality from the OLS regression fit (black line), the L_2 -based variant (red line), and from our estimation method (grey line).

Pepsi set their advertising expenditures to gain market share, where Coke and Pepsi consider the market share to evolve according to the LM (12) and SM, respectively. Thus, using data⁷ from 1968 to 1986, which spans a period of intense advertising competition, we estimate Coke's and Pepsi's beliefs on the market share evolution or, more concretely, the advertising effectiveness parameters using Proposition 3.

We find⁸ polyhedral parameter sets of Coke and Pepsi which rationalize the observed advertising expenditures up to an error of at most $\hat{\epsilon}_i$, $i \in \{1, 2\}$, reported in Table 1. Coke's and Pepsi's estimated effectiveness parameters $\hat{\theta}_1 = (\hat{k}_{1,1}, \hat{k}_{2,1})$ and $\hat{\theta}_2 = (\hat{k}_{1,2}, \hat{k}_{2,2})$, respectively, are positive, as required, and the believed advertising effectiveness of Coke is slightly larger than that of Pepsi, which agrees with qualitative findings (Yoffie and Wang, 2006). In Figure 2 we compare our results with those obtained using a L_2 -modification of our proposed method⁹ and the ordinary least square method (OLS), which differs fundamentally from the other two as it simply performs an ordinary least squares regression on the actual market share evolution. As expected, our method minimizes the maximum deviation from rational better-responding behavior, while OLS regression leads to a much more irrational behavior (since it neglects any relationship between consecutive action pairs). On the other hand, the objective of the OLS regression is to find the parameters that lead to the best fit of the market share evolution, shown in Figure 3¹⁰. In

⁷ The description of the data-set, generously provided by Gasmı, Laffont and Vuong, can be found in (Gasmı et al., 1988).

⁸ The code is made available at <https://github.com/amaddux9/Data-Driven-Behaviour-Estimation>.

⁹ The irrationality loss function is defined as $d_i(\theta_i, \mathcal{D}_i) = (1/m \sum_{j \in [m]} e^{-i^+(a_i^j, \bar{a}^j; \theta_i)^2})^{\frac{1}{2}}$

¹⁰ We added parameter estimates that are at most 5% more irrational than $\hat{\epsilon}_i$ to the polyhedral solution set (in this case a singleton).

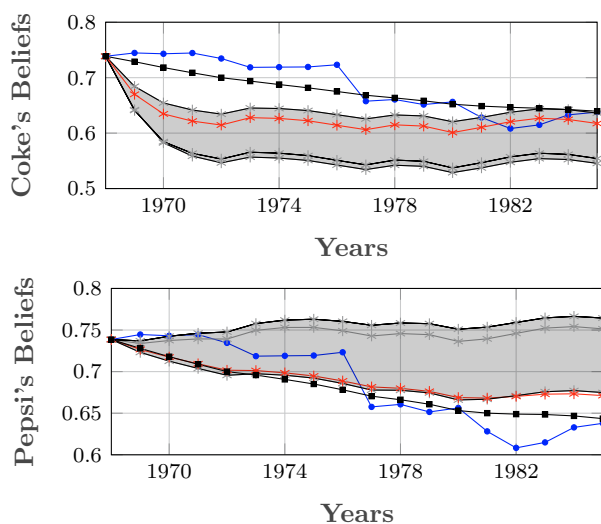


Fig. 3. Coke's and Pepsi's beliefs on market share evolution of Coke. The true market share evolution (blue line), the OLS regression fit (black line), the fit from the L_2 -based variant (red line) and the fit from the polyhedral solution set (grey region).

this case, the estimates given by our method are less in agreement with the actual market share evolution. This might be due to the fact that the market share is not only affected by the advertising expenditure but also by other economic factors and that the market share evolution is more complex than captured by the simplified Lanchester (resp. Sorger) model. Despite differences in the estimates (whose accuracy can only be computed by means of ground truth values of the parameters) and in their performances under different metrics, the above example has shown the applicability of our method to a real-world case scenario.

5. CONCLUSION

We designed a data-driven inverse-game-theoretic inference method to learn parametric utilities in multi-agent games. Our proposed technique works with both Nash equilibrium profiles or better-response trajectories observations, is computationally tractable, and produces polyhedral solution sets containing all the parameters which rationalize the observed behaviour best. The proposed dynamic inference method has been used to estimate Coke's and Pepsi's beliefs on market share evolution from their historical advertising expenditures. Unlike other standard inference methods the market share beliefs estimated with our approach minimize irrationality of the firms' observed expenditures.

REFERENCES

Ahuja, R.K. (1985). Minimax linear programming problem. *Operations Research Letters*, 4, 131–134.

Belgioioso, G., Ananduta, W., Grammatico, S., and Ocampo-Martinez, C. (2022). Operationally-safe peer-to-peer energy trading in distribution grids: A game-theoretic market-clearing mechanism. *IEEE Transactions on Smart Grid*.

Bertsimas, D., Gupta, V., and Paschalidis, I.C. (2015). Data-driven estimation in equilibrium using inverse op-

timization. *Mathematical Programming*, 153(2), 595–633.

Chintagunta, P.K. and Jain, D.C. (1995). Empirical analysis of a dynamic duopoly model of competition. *Journal of Economics & Management Strategy*, 4(1), 109–131.

Chintagunta, P.K. and Vilcassim, N.J. (1992). An empirical investigation of advertising strategies in a dynamic duopoly. *Management science*, 38(9), 1230–1244.

Gasmi, F., Vuong, Q.H., et al. (1988). An econometric analysis of some duopolistic games in prices and advertising. Technical report.

Inga, J., Bischoff, E., Köpf, F., and Hohmann, S. (2019). Inverse dynamic games based on maximum entropy inverse reinforcement learning. *arXiv preprint arXiv:1911.07503*.

Kuleshov, V. and Schrijvers, O. (2015). Inverse game theory: Learning utilities in succinct games. In *International Conference on Web and Internet Economics*, 413–427. Springer.

Li, R., Ye, X., Zhou, H., and Zha, H. (2019). Learning to match via inverse optimal transport. *Journal of machine learning research*, 20.

Maddux, A.M., Pagan, N., Belgioioso, G., and Dörfler, F. (2022). Data-driven behaviour estimation in parametric games.

Molloy, T.L., Ford, J.J., and Perez, T. (2017). Inverse noncooperative dynamic games. *IFAC-PapersOnLine*, 50(1), 11788–11793.

Narahari, Y., Garg, D., Narayanam, R., and Prakash, H. (2009). *Game theoretic problems in network economics and mechanism design solutions*. Springer Science & Business Media.

Nekipelov, D., Syrgkanis, V., and Tardos, E. (2015). Econometrics for learning agents.

Ng, A.Y., Russell, S.J., et al. (2000). Algorithms for inverse reinforcement learning. In *Icml*, volume 1, 2.

Pagan, N. and Dörfler, F. (2019). Game theoretical inference of human behavior in social networks. *Nature communications*, 10(1), 1–12.

Sessa, P.G., Bogunovic, I., Krause, A., and Kamgarpour, M. (2020). Contextual games: Multi-agent learning with side information. *Advances in Neural Information Processing Systems*, 33.

Sorger, G. (1989). Competitive dynamic advertising: A modification of the case game. *Journal of Economic Dynamics and Control*, 13(1), 55–80.

Trivedi, R., Yang, J., and Zha, H. (2020). Graphopt: Learning optimization models of graph formation. In *International Conference on Machine Learning*, 9603–9613. PMLR.

Tsai, D., Molloy, T.L., and Perez, T. (2016). Inverse two-player zero-sum dynamic games. In *2016 Australian Control Conference (AuCC)*, 192–196. IEEE.

Waugh, K., Ziebart, B.D., and Bagnell, J.A. (2011). Computational rationalization: the inverse equilibrium problem. *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 1169–1176.

Yoffie, D.B. and Wang, Y. (2006). *Cola wars continue: Coke and Pepsi in 2006*. Harvard Business School Boston, MA.