Efficiently characterizing games consistent with perturbed equilibrium observations

Juba ZIANI^{*} Venkat CHANDRASEKARAN[†] Katrina LIGETT[‡]

November 10, 2017

Abstract

We study the problem of characterizing the set of games that are consistent with observed equilibrium play. Our contribution is to develop and analyze a new methodology based on convex optimization to address this problem for many classes of games and observation models of interest. Our approach provides a *sharp, computationally efficient* characterization of the extent to which a particular set of observations constrains the space of games that could have generated them. This allows us to solve a number of variants of this problem as well as to quantify the power of games from particular classes (e.g., zero-sum, potential, linearly parameterized) to explain player behavior. We illustrate our approach with numerical simulations.

1 Introduction

This paper considers inference in game theoretic models of complete information. More precisely, we study the problem of recovering properties and characterizing parameter values of the games that are consistent with observed equilibrium play, and provide a simple procedure based on convex optimization to recover both the region of consistent games and properties of said region, in a computationally efficient manner. Further, our approach has the power to compute the size of the region of consistent games, and hence to determine when approximate point identification of the true payoff matrices (or parameter values) is possible.

Our approach differs from most related work in that we depart from the usual distributional assumptions on the observer's knowledge of the unobserved variables and payoff shifters; instead, we adopt the weaker assumption that the unobserved variables belong to a known set. This increases the robustness of our approach, reflecting ideas from the robust optimization literature—see [BTEGN09, BBC11, BB12]. Furthermore, our approach is formally computationally efficient and, when implemented, is able to handle games of much larger size than those considered in previous work.

Our approach may be viewed as complementary to a model-driven approach, in that the tools we provide here may be used to objectively evaluate the quality of fit one achieves under certain modeling assumptions. Our approach also allows us to explore a variety of assumptions about the information that might available to an observer of game play, and the effects that these assumptions would have on constraining the space of consistent games.

1.1 Summary of results

We consider a setting where, at each step, a finite-action, 2-player game is played, and an observer observes a correlated equilibrium (a more permissive concept than Nash equilibrium) of the game.¹ We assume that the games played on each step are closely related, in that each reflects a small perturbation in the payoffs of some underlying game; such perturbations are often referred to as *payoff shifters*.

^{*}California Institute of Technology, jziani@caltech.edu

 $^{^{\}dagger}\mbox{California Institute of Technology, venkatc@caltech.edu}$

[‡]Hebrew University of Jerusalem and California Institute of Technology, katrina.ligett@mail.huji.ac.il

 $^{^{1}}$ Our framework extends to multi-player games with succinct representations; for clarity, we focus here on the two-player case. See Section 5.1 for a discussion of succinct multi-player games.

In a departure from previous work, we do not make distributional assumptions on the payoff shifters, nor do we make any assumption on how the players decide which equilibrium to play, when multiple equilibria are present. Instead, we assume that the observer knows nothing of the equilibrium selection rule, and that the information the observer has on the unobserved payoff shifters is simply that the unobserved payoff shifters belong to a known set (see Section 2 for more details); this is a significantly weaker assumption than knowing exactly what distribution the payoff shifters are taken from. For example, imagine an analyst observes a routing game every day; the shifts in payoffs may come from a combination of several events such as changes in road conditions, traffic accidents, and work zones, whose potential effects on the costs of paths in a routing game may be difficult to predict and quantify precisely as a probability distribution.

In this setting, we give a computationally efficient characterization of the set of games that are consistent with the observations (Section 3.1); this set is "sharp", in the sense that it does not contain any game that is not consistent with the observations. One of our main new contributions is computational efficiency itself: the pioneering work of Beresteanu, Molchanov, and Molinari [BMM11] only checks membership of a game to the set of consistent games, and does so in a manner that is tractable in small games but intractable for larger games—see Section 1.2 for a more in-depth discussion. We also show that our framework accommodates an alternate model wherein the observer learns the expected payoff of each player at each equilibrium he sees; in our routing game example, think of an observer who sees the expected time each player spends in traffic. We refer to this setting as "partial payoff information," and discuss it in Sections 3 and 4.

Our second main contribution is our ability to quantify the size of the set of consistent games. We give an efficient algorithm (see Section 3.4, Algorithm 1) that takes a set of observations as input and computes the diameter of the sharp region of consistent games. The diameter of the consistent set is of interest to an observer, because it gives him a measure of how sharp the conclusions he can draw from the observations are (the larger the diameter, the less sharp the conclusions), and in particular the diameter quantifies the level of approximate point identification that is achievable in a particular setting. Additionally, in Section 4, Lemmas 2 and 3, we give structural conditions on the sets of observations that allow for accurate recovery. We also exhibit examples in which said conditions do not hold, and therefore accurate recovery is not possible.

We show we can extend our framework (Section 3.3) to find the set of consistent games when restricted to games with certain linear properties, e.g., zero-sum games, potential games, and games whose utilities can be parametrized by linear functions; this allows us to determine to what extent the observed behavior is consistent with such assumptions on the underlying game.

In Section 5, we show we can extend our framework to finite games with a large number of players, provided the game has a succinct representation. We further show our framework's potential to deal with games with infinite action spaces, using Cournot competition as an example.

Finally, in Section 6, we illustrate our approach with simulations, in both a simple entry game (Section 6.1), and in large Cournot competition games (Section 6.2).

1.2 Related work

One important modeling issue is whether and why one would ever observe multiple, differing behaviors of a single agent. A natural, well-established approach models different observations found in the data as stemming from random perturbations to the agents' utilities, as in [BV84, BR91, Tam03, Sei06, AL10, AL11, AL12, BHKN10, BHHR11]. In dynamic panel models, one observes equilibria across several markets sharing common underlying parameters, and in particular [HM10] considers a setting in which a unique, fixed equilibrium is played within each market. We adopt a similar approach here, and assume that we have access to several markets or locations that play perturbed versions of the same game, and that a single (mixed) equilibrium is played in each market.

In the branch of the econometrics literature that aims to infer parameters in game theoretic model, it is often the case that one requires that the game be small or that the utilities of the players can be written as simple functions of a restricted number of parameters. For example, 2-player entry games with entry payoffs parametrized as linear functions of a small number of variables, as seen in [Tam03] and subsequent work, are among the most-studied in the literature. One drawback of this literature is that when the space of parameters is high-dimensional or when multiple equilibria exist, identification of the true parameters of the game often becomes impossible, since the observations do not correspond to a unique consistent explanatory game. A number of recent papers [ABJ04, CT09, BMM11, NST15] consider instead the problem

of constructing *regions* of parameters that contain the true value of the parameters they aim to recover from equilibrium observations of games. For example, Nekipelov, Syrgkanis, and Tardos [NST15] study a dynamic sponsored search auction game, and provide a characterization of the *rationalizable set*, consisting of the set of private parameters that are consistent with the observations, under the relaxed assumption that players need not follow equilibrium play, but rather use some form of no-regret learning. Relatedly, Andrews, Berry, and Jia [ABJ04] and Ciliberto and Tamer [CT09] compute confidence regions for the value of the true parameter, but their regions are not "sharp," in the sense that they may contain parameter values that are not consistent with some of the implications of their models.

Perhaps closest to the present work, Beresteanu, Molchanov, and Molinari [BMM11] combine random set theory and convex optimization to give a representation of the *sharp identification region* as the set of values for which the solution to a convex optimization program with a random objective function is almost surely (in the payoff shifters) equal to 0.² Hence, verifying membership of a parameter value to the *sharp identification region* can be done efficiently in simple settings such as entry-games with linearly parametrized payoffs. This is an exciting advance, especially when considering games with few players and small action sets; however, for computational reasons, the approach becomes impractical in large games, such as 2-player games with many actions per player:

- The Beresteanu, Molchanov, and Molinari [BMM11] framework can verify that a vector of parameter values belongs to the sharp identification set, but does not provide an efficient, searchable representation of the sharp identification set itself, nor an algorithm to efficiently find a point in the set.
- One can verify that a parameter vector belongs to the sharp identification set by checking that a particular condition holds for almost all possible realizations of the payoff shifters. Beresteanu, Molchanov, and Molinari [BMM11] further show that one can cluster payoff shifters into groups such that all perturbed games in the same group have the same set of Nash equilibria; one then must check the condition only once per group. In particular, in their entry-game example, the number of such groups is small, and thus this is a computationally tractable task. However, in more complex games, or for more general equilibrium concepts, the number of such groups can become intractably large in the number of actions available to each player.
- Finally, the BMM framework [BMM11] relies on being able to compute *all* equilibria of each of the perturbed games, which may be impractical. (In general, in the worst case, it is PPAD-hard to find *any* Nash equilibrium, even for 2-player games [CDT07].)

The goal of the present paper is similar to the goal of [BMM11], in the sense that we wish to sharply understand the set of games that are consistent with a set of observations (for us, correlated equilibria of perturbed games). We also use the setting of their simulations as the jumping off point for our own experimental section. However, our approach differs from that of [BMM11] in two main ways. First, we make weaker assumptions on the information on the unobserved payoff shifters available to the observer; our approach to modeling the perturbations is inspired by the concept of uncertainty sets in robust optimization (see [BTEGN09, BBC11, BB12]). Second, our framework provides a computationally efficient characterization of the consistent set, both in theory and in practice, on games of large size, and also gives efficient and practical algorithms to find points in the consistent set, compute its diameter, and test whether it contains games with certain properties.

A handful of papers in the computer science literature have looked at slightly different but related questions, arising when observing equilibria of games whose payoffs are unknown. In particular, Bhaskar et al. [BLSS14] and Rogers et al. [RRUW15] study a network routing setting in which equilibrium behavior can be observed but edge costs are unknown, and study the query complexity of devising a variant of the game to induce desired target flows as equilibria. Barman et al. [BBEW13] adopt a model in which the observer observes what joint strategies are played when restricting the actions of the players in a complete information game with no perturbations, and show that data with certain properties can be rationalized by games with low complexity. Kuleshov and Schrijvers [KS15] assume the observer sees equilibria from a collection of succinct games that share common parameters, and give a convex and computationally efficient

 $^{^{2}}$ Our notion of the *consistent set* is closely analogous to the sharp identification region of [BMM11]. We use different terminology to highlight that they are derived under somewhat different settings.

characterization of the set of consistent parameters, for a wide class of games with known succinct structure. Their work makes the interesting and important observation that correlated equilibria of succinct games define a tractable number of linear constraints on the games whose equilibria are observed; we make use of said observation in our work. We note however that our framework is different in scope, as i) our framework accommodates more general relationships between the games whose equilibria are observed (specifically, we do not assume common parameters and known structure; we only make the assumption that the games are close to each other according to the metrics we define in Section 2); and ii) our framework can compute metrics and recover properties of the set of underlying games that are consistent with the observations, such as the diameter of the set of consistent games and whether the observations can be explained by games with specific properties.

2 Model and setting

2.1 Players' behavior

Consider a finite two-player game G; we will refer to it as the *true* or *underlying* game. Let $\mathcal{A}_1, \mathcal{A}_2$ be the finite sets of actions available to players 1 and 2, respectively, and let $m_1 = |\mathcal{A}_1|$ and $m_2 = |\mathcal{A}_2|$ be the number of actions available to them. For every $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$, we denote by $G_p(i, j)$ the payoff of player p when player 1 chooses action i and player 2 chooses action j. $G_p \in \mathbb{R}^{m_1 \times m_2}$ is the vector representation of the utility of player p, and we often abuse notation and write $G = (G_1, G_2)$. The *strategies* available to player p are simply the distributions over \mathcal{A}_p . A *strategy profile* is a pair of strategies (distributions over actions), one for each player. A *joint strategy profile* is a distribution over pairs of actions (one for each player); it is not required to be a product distribution. We refer to strategies as *pure* when they place their entire probability mass on a single action, and *mixed* otherwise.

We consider l perturbed versions of the game G, indexed by $k \in [l]$ so that the k^{th} perturbed game is denoted G^k ; one can for instance imagine each G^k as a version of the game G played in a different location or market k. The same notation as for G applies to the G^k 's.

Throughout the paper, we assume that for each k, the players' strategies are given by a correlated equilibrium of the complete information game G^k . In the presence of several such equilibria, no assumption is made on the selection rule the players use to pick which equilibrium to play (though we assume they both play according to the same equilibrium). Correlated equilibria are defined as follows:

Definition 1. A probability distribution e is a correlated equilibrium of game $G = (G_1, G_2)$ if and only if

$$\sum_{j=1}^{m_2} G_1(i,j) e_{ij} \ge \sum_{j=1}^{m_2} G_1(i',j) e_{ij} \ \forall i,i' \in \mathcal{A}_1$$
$$\sum_{i=1}^{m_1} G_2(i,j) e_{ij} \ge \sum_{i=1}^{m_1} G_2(i,j') e_{ij} \ \forall j,j' \in \mathcal{A}_2$$

The notion of correlated equilibrium extends the classical notion of Nash equilibrium by allowing players to act jointly; as every Nash equilibrium of a game is a correlated equilibrium of the same game, many of our results also have implications for Nash equilibria.

2.2 Observation model

Our observer does not have access to the payoffs of the underlying game G nor of the perturbed games G^k , for any k in [l]. We model an observer as observing, for each perturbed game G^k , the entire correlated equilibrium distribution $e^k \in \mathbb{R}^{m_1 \times m_2}$, where $e^k(i, j)$ denotes the joint probability in the kth perturbed game of player 1 playing action $i \in \mathcal{A}_1$ while player 2 plays action $j \in \mathcal{A}_2$. Note that as e^k represents a probability distribution, we require $e^k(i, j) \ge 0 \forall (i, j)$ and $\sum_{i,j} e^k_{ij} = 1$.

In practice, an observer may not have access to the full distribution of each equilibrium. Our approach can also accommodate a less-demanding assumption, that the observer receives samples from the equilibrium distribution, perhaps from multiple plays of the same (possibly mixed) equilibrium. If one additionally imposes the mild assumption that the payoffs of the underlying game G are bounded, one can use the empirical distribution on the sampled observations to approximate the set of consistent games, by adding a slack parameter to the inequalities of Definition 1; this can easily be incorporated into our framework. Upper bounds on how to pick the parameter can be obtained through classical concentration inequalities, or through the techniques of [GJN17, CHT07].³ For ease of exposition, in this paper, we focus on the case where the full distribution is observed.

In this paper, we consider two variants of the model of observations we just described:

- In the partial payoff information setting, the observer has access to equilibrium observations e^1, \ldots, e^l , and additionally to the expected payoff of equilibrium e^k on perturbed games G^k , for each player p and for all $k \in [l]$; we denote said payoff v_p^k and note that $v_p^k = e^k G_p^k$.
- In the payoff shifter information setting, at each step k, a payoff shifter $\beta^k = (\beta_1^k, \beta_2^k) \in \mathcal{R}^{m_1 \times m_2} \times \mathcal{R}^{m_1 \times m_2}$ is added to game $G = (G_1, G_2)$, and the perturbed games G^k result from the further addition of small perturbations to the $G + \beta^k$'s. The observer knows β^1, \ldots, β^l and observes e^1, \ldots, e^l of perturbed games G^1, \ldots, G^l . This setting represents a situation in which changes in the behavior of agents are observed as a function of changes in observable economic parameters (taxes, etc.).

While the payoff shifter information setting is the model of perturbations that is commonly used in the literature, the partial payoff information setting has not been used in previous work, to the best of our knowledge. We introduce it to model the following types of situation:

Example 1. Two firms are competing for customers in Los Angeles, and an observer follows what actions the two L.A. firms take over the course of each quarter. The observer also learns the quarterly revenue of each firm.

Example 2. Several agents are playing a routing game, and the observer sees not just the routes the players choose, but also the amount of time players spend in traffic.

2.3 Observer's knowledge about the perturbations

Our paper aims to characterize the games that explain equilibrium observations under the partial payoff and payoff shifter information settings when the perturbations are known to be "small" and the perturbed games are thus "close" to the underlying game. The next few definitions formalize our notion of closeness, and Assumption 1 formalizes the information the observer has about the perturbations added to the underlying game G.

Definition 2. A game G is δ -close to games G^1, \ldots, G^l with respect to metric d for $\delta > 0$ if and only if $d(G^1, \ldots, G^l | G) \leq \delta$.

We think of d as distances and therefore convex functions of the perturbations $G - G^k$ for all k. For the above definitions to make sense in the context of this paper, we need a metric whose value on a set of games G, G^1, \ldots, G^l is small when G, G^1, \ldots, G^l are close in terms of payoffs. We consider the following metrics:

Definition 3. The sum-of-squares distance between games G and G^1, \ldots, G^l is given by

$$d_2(G^1,\ldots,G^l|G) = \sum_{k=1}^l (G_1 - G_1^k)'(G_1 - G_1^k) + \sum_{k=1}^l (G_2 - G_2^k)'(G_2 - G_2^k).$$

³More precisely, [GJN17] provide upper bounds on the slack parameter that ensure the approximated set covers the true parameters with high probability; therefore, the consistent sets of Section 3.1 and Section 3.3 can be generalized to approximate sets for the sample setting. Further, [GJN17] also provide bounds on how much the optimal value of a convex optimization problem changes when replacing the true constraints by their empirical counterparts with an appropriately chosen slack parameter. In this manner, our results in Section 3.4 and Algorithm 1 can be extended to obtain high probability bounds on the diameter of the true consistent set, when the observer only has access to samples. [CHT07] propose a framework that ensures the approximated set is on the one hand a consistent estimator of the sharp set of consistent games, and on the other hand covers said set with high probability as the number of samples grows large.

The maximum distance between games G and G^1, \ldots, G^l is defined as

$$d_{\infty}(G^{1},\ldots,G^{l}|G) = \max_{p \in \{1,2\}, \ k \in [l]} \|G_{p} - G_{p}^{k}\|_{\infty},$$

where $\|.\|_{\infty}$ denotes the usual infinity norm.

Both distances are useful, in different situations. The sum-of-squares distance is small when the variance of the perturbations added to G is known to be small, but allows for worst-case perturbations to be large. An example is when the G^{k} 's are randomly sampled from a distribution with mean G, unbounded support, and small covariance matrix, in which case some of the perturbations may deviate significantly from the mean but with low probability, while the average squared perturbation remains small. If the distribution of perturbations is i.i.d Gaussian, the sum-of-squares norm replicates the log-likelihood of the estimations and follows a Chi-square distribution. The maximum distance, in contrast, is small when it is known that all perturbations are small and bounded; one example is when the perturbations are uniform in a small interval $[-\delta, \delta]$.

Throughout the paper, we make the following assumption on the information about the perturbations that is available to the observer:

Assumption 1. Let G be the underlying game and G^1, \ldots, G^l be the perturbed games that generated observations e^1, \ldots, e^l .

- In the partial payoff information settings, the observer knows that G is δ -close to games G^1, \ldots, G^l with respect to some metric d and magnitude $\delta \geq 0$.
- In the payoff shifter information setting with observed shifters β^1, \ldots, β^l , the observer knows that G is δ -close to the unshifted games $(G^1 \beta^1), \ldots, (G^l \beta^l)$ with respect to some metric d and magnitude $\delta \geq 0$.

Assumption 1 defines a convex set the observer knows the perturbations must belong to, much like the uncertainty sets given in [BB12, BTEGN09, BBC11]. We note that the d_2 and d_{∞} distances we focus on define respectively an ellipsoidal and a polyhedral uncertainty set (as seen in [BBC11]).

Remark 1 (On the shape of the set of perturbations). While we make Assumption 1 for convenience and simplicity of exposition, our framework is able to handle more general sets of perturbations. In particular, the results of Section 3 can easily be extended to any convex set of perturbations that has an efficient, easy-to-optimize-over representation. This includes classes of sets defined by a tractable number of linear or convex quadratic constraints, which in turn encompasses many of the uncertainty sets considered in [BB12], such as the central limit theorem or correlation information sets, and most of the typical sets presented therein.

Remark 2 (On the observer's knowledge). In some settings, one may have reasonable upper bounds on the value of δ : for example, in a routing game, the observer may know that delays in traffic cannot exceed a few hours. However, in some settings it may be unreasonable to assume the observer knows a reasonably tight upper bound δ . We remark that in such a case, our framework still gives the observer the power to experiment with different values of δ and to see how they affect the set of consistent games and its diameter. Along these lines, we plot heatmaps of the consistent set as a function of δ in Section 6.

It is usually the case in the econometrics literature that the distribution of perturbations is assumed to be known. In such a case, there is a direct mapping from distributions over perturbations to values of δ : one can pick δ such that $d(G^1, \ldots, G^l|G) \leq \delta$ with high probability. For example, if the perturbations are known to be Gaussian with a known mean and standard deviation, $d_2(G^1, \ldots, G^l|G)$ follows a Chi-square distribution whose quantiles are known; one can then pick δ such that $d(G^1, \ldots, G^l|G) \leq \delta$ with probability $p(\delta)$. This immediately guarantees that the recovered set of consistent games contains the true parameters (G, G^1, \ldots, G^l) with probability at least $p(\delta)$. Such a bound does not use all of the information available to the observer about the perturbations, hence the recovered set is not sharp; however, in our simulations in Section 6 we observed that under reasonable distributional assumptions, the set identified by our framework remained reasonably small.

2.4 Consistent games

In this paper, as in [BMM11], we adopt an observation-driven view that describes the class of games that are consistent with the observed behavior. Given a set of observations, we define the set of consistent games as follows:

Definition 4 (δ -consistency). We say a game \tilde{G} is δ -consistent with the observations when there exists a set of games $(\tilde{G}, \tilde{G}^1, \ldots, \tilde{G}^l)$ such that for all k, e^k is an equilibrium of \tilde{G}^k , and:

- If in the partial payoff information model of observations, $\sum_{p} \tilde{G}_{p}^{k} e^{k} = v_{p}^{k}$ for all players p and $d(\tilde{G}^{1}, \ldots, \tilde{G}^{l}|\tilde{G}) < \delta$.
- If in the payoff shifter information model, $d(\tilde{G}^1 \beta^1, \dots, \tilde{G}^l \beta^l | \tilde{G}) \leq \delta$.

The set of all δ -consistent games with respect to metric d is denoted $S_d(\delta)$.

Given the specifications of our model, it is often the case that, given a set of observations with no additional assumption on the distribution of perturbations nor on a the rule used to select among multiple equilibria, it is not possible to recover an approximation to a unique game that generated these observations (no matter what recovery framework is used). That is, the diameter of the consistent set can sometimes be too large for approximate point identification to be possible, which is highlighted in the following example:

Example 3. Take any set of observations e^1, \ldots, e^l under the no payoff information observation model, and let \hat{G} be the all-constant game, i.e., $\hat{G}_1(i,j) = \hat{G}_2(i,j) = c$ for some $c \in \mathbb{R}$ and for all $(i,j) \in \mathcal{A}_1 \times \mathcal{A}_2$. Let $\hat{G}^1 = \ldots = \hat{G}^l = \hat{G}$. Then for all $k \in [l]$, e^k is an equilibrium of \hat{G}^k , and $d_2(G^1, \ldots, G^l|G) = d_{\infty}(G^1, l \operatorname{dots}, G^l|G) = 0$. That is, \hat{G} is a trivial game, and it is consistent with all possible observations. Even when e^1, \ldots, e^l are generated by a non-trivial G, without any additional observations, an observer cannot determine whether G or \hat{G} is the underlying game. In fact, both games are consistent with all implications of our model. We note that this issue arises regardless of how inferences will be drawn about the observations, so long as the approach does not discard consistent games.

It may thus be of interest to an observer to compute the diameter of the consistent set, either to determine whether point identification is possible, or simply to understand how tightly the observations constrain the space of consistent games. We define it as follows:

Definition 5. The diameter $D(S_d(\delta))$ of consistent set $S_d(\delta)$ is given by

$$D(S_d(\delta)) = \sup\{\|\hat{G} - \tilde{G}\|_{+\infty} \text{ s.t. } \tilde{G}, \hat{G} \in S_d(\delta)\}$$

When the diameter is small, then every game in the consistent set is close to the true underlying game, and approximate point identification is achievable. When the diameter $D(S_d(\delta))$ grows large, point identification is impossible independently of what framework is used for recovery, as there exist two games that are $D(S_d(\delta))$ -far apart in terms of payoff, yet either could have generated all observations.

3 A convex optimization framework

In this section, we show how techniques from convex optimization can be used to recover the perturbationminimizing explanation for a set of observations, determine the extent to which observations are consistent with certain assumptions on the underlying game, and determine whether a set of observations tightly constrains the set of games that could explain it well. The results in this section are not tied to a specific observation model.

3.1 Efficient characterization of the set of consistent games

We will show that for every δ , and $d \in \{d_2, d_{+\infty}\}$, the set of consistent games $S_d(\delta)$ has an efficient, convex representation.

Claim 1. If in the "partial payoff information" model of observations:

$$S_{d}(\delta) = \begin{cases} \exists (G^{1}, \dots, G^{l}) \text{ with } d(G^{1}, \dots, G^{l}|G) \leq \delta \text{ s.t.} \\ \sum \limits_{j=1}^{m_{2}} G_{1}^{k}(i, j) e_{ij}^{k} \geq \sum \limits_{j=1}^{m_{2}} G_{1}^{k}(i', j) e_{ij}^{k} \forall i, i' \in \mathcal{A}_{1}, \forall k \in [l], \\ \sum \limits_{i=1}^{m_{1}} G_{2}^{k}(i, j) e_{ij}^{k} \geq \sum \limits_{i=1}^{m_{1}} G_{2}^{k}(i, j') e_{ij}^{k} \forall j, j' \in \mathcal{A}_{2}, \forall k \in [l] \\ \sum \limits_{p} \tilde{G}_{p}^{k'} e^{k} = v_{p}^{k} \end{cases}$$

If in the "payoff shifter information" model:

$$S_{d}(\delta) = \begin{cases} \exists (G^{1}, \dots, G^{l}) \text{ with } d(G^{1} - \beta^{1}, \dots, G^{l} - \beta^{l} | G) \leq \delta \text{ s.t.} \\ \sum_{j=1}^{m_{2}} G_{1}^{k}(i, j) e_{ij}^{k} \geq \sum_{j=1}^{m_{2}} G_{1}^{k}(i', j) e_{ij}^{k} \forall i, i' \in \mathcal{A}_{1}, \forall k \in [l], \\ \sum_{i=1}^{m_{1}} G_{2}^{k}(i, j) e_{ij}^{k} \geq \sum_{i=1}^{m_{1}} G_{2}^{k}(i, j') e_{ij}^{k} \forall j, j' \in \mathcal{A}_{2}, \forall k \in [l] \end{cases}$$

Proof. Follows from the definition of δ -consistency (Definition 4)

ns bolong

We remark that as in [BMM11], our sets are *sharp*: any game that explains the observations belongs to this set, and any game that belongs to this set is consistent with our assumptions and observations. Indeed, if G is in the consistent set, there must exist perturbations of valid magnitude (given by the corresponding G^1, \ldots, G^l) and an equilibrium e^k of each perturbed game that together would lead to our observations, by the definition of the consistent set.

These consistent sets have efficient convex representations, for two reasons. First, all constraints are always linear except those of the form $d(G_p^1, \ldots, G_p^l|G) \leq \delta$. When $d = d_2$, $d(G_p^1, \ldots, G_p^l|G) \leq \delta$ is a simple convex quadratic constraint, while when $d = d_{\infty}$, $d(G_p^1, \ldots, G_p^l|G) \leq \delta$ is equivalent to the following collection of linear constraints:

$$\begin{aligned} -\delta &\leq G_p^k - G_p \; \forall p \in \{1, 2\}, \forall k \in [l] \\ G_p^k - G_p &\leq \delta \; \forall p \in \{1, 2\}, \forall k \in [l]. \end{aligned}$$

Second, the number of constraints describing each set is quadratic in the number of player actions m_1 and m_2 .

As mentioned in Section 2, in all observation models, the assumption that $d(G_p^1, \ldots, G_p^l|G) \leq \delta$ can easily be replaced by an assumption of the perturbations $G^k - G$ being in any tractable convex set. In particular, many of the sets considered in [BB12] fit this requirement, and they describe robust information that an observer without distributional knowledge of the perturbations could realistically have on said perturbations: for example, an observer could know that the sum or average of the perturbations satisfies certain lowerand upper-bounds.

3.2 Recovering the perturbation-minimizing consistent game

Here, we consider the problem of recovering a game that best explains a given set of observations from perturbed games, according to the desired distance metric d. One reason to do so is that it enables an observer to test whether there exists any game in $S_d(\delta)$ that is consistent with specific properties and to give a measure of how much of $S_d(\delta)$ has said properties—see Section 3.3. Or, it could be that the observer is simply interested in recovering the "best" game according to any simple convex metric of interest. For any metric d and any observation model, this can be done simply by solving:

$$\begin{array}{ll}
\min_{G,\delta} & \delta \\
\text{s.t.} & G \in S_d(\delta)
\end{array}$$
(1)

It is easy to see that this program returns the game G and the minimum value of δ such that $d(G^1, \ldots, G^l|G) \leq \delta$ (resp. $d(G^1 - S^1, \ldots, G^l - S^l|G) \leq \delta$) in the partial payoff information setting (resp. the payoff shifter

information setting) where the G^{k} 's satisfy all equilibrium constraints, hence Program (1) returns the perturbation-minimizing G that is consistent with all observations. When $d = d_{\infty}$, this is a linear program; when $d = d_2$, this is a second-order cone program, using the same reasoning as in Section 3.1 (this holds even with δ as a variable). Both types of programs can be solved efficiently, as seen in [BV04].

3.3 Can observations be explained by linear properties?

This convex optimization-based approach can further be used to determine whether there exists a game that is compatible with the observations and that also has certain additional properties, as long as the properties of interest can be expressed as a tractable number of linear equalities and inequalities. One can then solve program (1) with said linear equalities and inequalities as additional constraints (the program remains a SOCP or LP with a tractable number of constraints), then check whether the optimal value is greater than or less than δ . If the optimal value is greater than δ , then there exists no game with those properties that belongs to the δ -consistent set; if the optimal value is smaller than δ , then the recovered game displays the additional properties and belongs to the δ -consistent set. In what follows, we present a few examples of interesting properties that fit this framework.

Zero-sum games A zero-sum game is a game in which for each pure strategy (i, j), the sum of the payoff of player 1 and the payoff of player 2 for (i, j) is always 0. One can restrict the set of games we look for to be zero-sum games, at the cost of separability of Program (1), by adding constraints $G_1(i, j) = -G_2(i, j) \forall (i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$.

Exact potential games A 2-player game G is an exact potential game if and only if it admits an exact potential function, i.e., a function Φ that satisfies:

$$\Phi(i,j) - \Phi(i',j) = G_1(i,j) - G_1(i',j) \ \forall i,i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2$$

$$\tag{2}$$

$$\Phi(i,j) - \Phi(i,j') = G_2(i,j) - G_2(i,j') \ \forall i \in \mathcal{A}_1, \forall j,j' \in \mathcal{A}_2 \tag{3}$$

In order to restrict the set of games we are searching over to the set of potential games, one can introduce m_1m_2 variables $\Phi(i, j)$ and constraints (2), (3) in Program (1).

Games generated through linear parameter fitting It is common in the literature to recover a game with the help of a parametrized function whose parameters are calibrated using the observations. In many applications, linear functions of some parameters are considered—entry games are one example. Our framework allows one to determine whether there exist parameters for such a linear function that provide good explanation for the observations. When such parameters exist, one can use the mathematical program to find a set of parameters that describe a game which is consistent with the observations. Take two functions $f_1(\theta)$ and $f_2(\theta)$ that are linear in the vector of parameters θ and output a vector in $\mathbb{R}^{m_1 \times m_2}$. It suffices to add the optimization variable θ and the linear constraints $G_1 = f_1(\theta)$ and $G_2 = f_2(\theta)$ to Program (1) to restrict the set of games we look for to games linearly parametrized by f_1, f_2 .

3.4 Computing the diameter of the consistent set

In this section, we provide an algorithm (Algorithm 1) for computing the diameter of $S_d(\delta)$, for any given value of δ . Because the diameter is a property of the consistent set and not of the framework used to recover an element from said set, this tells an observer whether approximate point identification is possible *independently* of what framework is used for recovery. In particular, when the diameter is small, our framework approximately recovers the true underlying game (see Section 3.2). When the diameter is large, *no* framework can achieve approximate point identification of a true, underlying game.

Algorithm 1 is computationally efficient for the considered metrics d_2 and d_{∞} : it solves $2m_1m_2$ linear programs for d_{∞} , and $2m_1m_2$ second-order cone programs (SOCP) for $d = d_2$ with a tractable number of constraints. The algorithm computes the diameter of the consistent set:

Lemma 1. The output $\mathcal{A}(\delta)$ of Algorithm 1 run with input δ satisfies $\mathcal{A}(\delta) = D(S_d(\delta))$.

Proof. See Appendix A.

ALGORITHM 1: Computing the diameter of the consistent set

Input: Observations e^1, \ldots, e^l , magnitude of perturbations δ , metric dOutput: Real number $\mathcal{A}(\delta)$ (may be infinite) for $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$, player $p \in 1, 2$ do $P_{\delta,p}(i, j) = \sup_{\substack{\tilde{G}, \tilde{G}, \gamma \\ \text{s.t.}}} \gamma$ s.t. $\tilde{G} \in S_d(\delta)$ $\hat{G} \in S_d(\delta)$ $\tilde{G}_1(i, j) - \hat{G}_1(i, j) \ge \gamma$ end $\mathcal{A}(\delta) = \max_{(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2} \max_{p \in \{1, 2\}} P_{\delta, p}(i, j)$

4 Consistent games with partial payoff information: when is recovery possible?

This section considers the *partial payoff information* variant of the observation model described in Section 2. We ask the following question: when is it possible to approximate the underlying game, in the presence of partial payoff information? We answer this question by giving bounds on the diameter of the consistent set $S_d(\delta)$ as a function of δ and the observations e^1, \dots, e^l , for both metrics d_2 and d_{∞} . Recall that in this setting, for an equilibrium e^k observed from perturbed game G^k , the observer learns

Recall that in this setting, for an equilibrium e^k observed from perturbed game G^k , the observer learns not only e^k , but also the expected payoff v_p^k of player p in said equilibrium strategy on game G^k . Similar to the previous sections, we are interested in computing a game \hat{G} that is close to some perturbed games $\hat{G}^1, \dots, \hat{G}^l$ that (respectively) have equilibria e^1, \dots, e^l with payoffs v^1, \dots, v^l . For simplicity of presentation, we recall that the optimization program that the observer solves is separable and note that he can thus solve the following convex optimization problem for player 1, and a similar optimization problem for player 2:

$$P(\epsilon) = \min_{G_{1}^{k}, G_{1}} d(G_{1}^{1}, ..., G_{1}^{k} | G_{1})$$

s.t.
$$\sum_{j=1}^{d} G_{1}^{k}(i, j) e_{ij}^{k} \ge \sum_{j=1}^{d} G_{1}^{k}(i', j) e_{ij}^{k} \,\forall i, i' \in \mathcal{A}_{1}, \forall j \in \mathcal{A}_{2}, \forall k \in [l]$$

$$e^{k'} G_{1}^{k} = v_{1}^{k} \,\forall k \in [l]$$
(4)

We take $l \ge m_1 m_2$ and make the following assumption for the remainder of this subsection, unless otherwise specified:

Assumption 2. There exists a subset $E \subset \{e^1, ..., e^l\}$ of size m_1m_2 such that the vectors in E are linearly independent.

We abuse notation and denote by E the $m_1m_2 \times m_1m_2$ matrix in which row i is given by the i^{th} element of set E, for all $i \in [m_1m_2]$; also, we write $d(G^1, ..., G^l|G) = \sum_p d(G_p^1, ..., G_p^l|G_p)$, i.e., $d(G_p^1, ..., G_p^l|G_p)$ is the part of $d(G^1, ..., G^l|G)$ that corresponds to player p. For every $p \in \mathbb{N} \cup \{+\infty\}$, let $\|.\|_p$ be the p-norm. We can define the corresponding induced matrix norm $\|.\|_p$ that satisfies $\|M\|_p = \sup_{x \neq 0} \frac{\|Mx\|_p}{\|x\|_p}$ for any matrix

 $M \in \mathbb{R}^{m_1 m_2 \times m_1 m_2}.$

Lemmas 2 and 3 highlight that if one has linearly independent observations such that E is well-conditioned, any optimal solution of Program (4) necessarily recovers a game whose payoffs are close to the payoffs of the underlying game, as long as the perturbations are small. The statements are given for both metrics introduced in Section 2. Example 4 shows that E being well-conditioned is, in fact, necessary.

Lemma 2. Let G be the underlying game, and $G^1, ..., G^l$ be the games generating observations $e^1, ..., e^l$, where $l = m_1 m_2$. Suppose that for player p, $d_2(G_p^1, ..., G_p^l|G_p) \leq \delta$. Let $(\hat{G}_p, \hat{G}_p^1, ..., \hat{G}_p^l)$ be an optimal solution of Program (4) for player p with distance function d_2 . Then

$$||G_p - \hat{G}_p||_2 \le \sqrt{2||E^{-1}||_2 \cdot \delta}.$$

Proof. For simplicity of notation, we drop the p indices. We first remark that $(G, G^1, ..., G^l)$ is feasible for Program (4); as $(\hat{G}, \hat{G}^1, ..., \hat{G}^l)$ is optimal, it is necessarily the case that

$$\sum_{k=1}^{l} \|\hat{G} - \hat{G}^k\|_2^2 \le \sum_{k=1}^{l} \|G - G^k\|_2^2 \le \delta.$$

Let us write $\Delta G = G - \hat{G}$. We know that for all k, $e^{k'}G^k = e^{k'}\hat{G}^k = v^k$, and thus $e^{k'}(G^k - \hat{G}^k) = 0$. We can write

$$\begin{split} E\Delta G &= (e_1'(G - \hat{G}) \dots e_l'(G - \hat{G}))' \\ &= (e_1'(G - G^1 + G^1 - \hat{G}^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + G^l - \hat{G}^l + \hat{G}^l - \hat{G}))' \\ &= (e_1'(G - G^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + \hat{G}^l - \hat{G}))'. \end{split}$$

Let $x_k = G - G^k + \hat{G}^k - \hat{G}$. We then have $||E\Delta G||_2^2 = \sum_{k=1}^l x'_k e_k e'_k x_k$, as $e_k e'_k$ is a symmetric, positive semi-definite, stochastic matrix, all its eigenvalues are between 0 and 1 and

$$||E\Delta G||_2^2 \le \sum_{k=1}^l x'_k x_k = \sum_{k=1}^l ||x_k||_2^2 \le 2\delta.$$

It immediately follows that $\|\Delta G\|_2 \leq \sqrt{2\|E^{-1}\|_2 \cdot \delta}$.

Lemma 3. Let G be the underlying game, and $G^1, ..., G^l$ be the games generating observations $e^1, ..., e^l$, where $l = m_1m_2$. Suppose that for player p, $d_{\infty}(G_p^1, ..., G_p^l|G_p) \leq \delta$. Let $(\hat{G}_p, \hat{G}_p^1, ..., \hat{G}_p^l)$ be an optimal solution of Program (4) for player p with distance function d_{∞} . Then

$$\|G_p - \hat{G}_p\|_{\infty} \le 2\|E^{-1}\|_{\infty} \cdot \delta.$$

Proof. See Appendix B.

When E is far from being singular, as long as the perturbations are small, we can accurately recover the payoff matrix of each player. This has a simple interpretation: the further E is from being singular, the more diverse the observations are. More diverse observations means more information for the observer, and allows for more accurate recovery. On the other hand, if the matrix E is close to being singular, it means that one sees the same or similar observations over and over again, and does not gain much information about the underlying game—there are more payoffs to recover than different observations, and the system is underdetermined.

An extreme example arises when we take E to be the identity matrix, in which case we observe every single pure strategy of the game and an approximation of the payoff of each of these strategies, allowing us to approximately reconstruct the game. It is also the case that there are examples in which $||E^{-1}||_{\infty}$ is large and there exist two games that are far from one another, yet both explain the observations, making our bound essentially tight:

Example 4. Consider the square matrix $E \in \mathbb{R}^{4\times 4}$ with probability $0.25 + \epsilon$ on the diagonal and $\frac{0.75-\epsilon}{3}$ off the diagonal, i.e., we observe four equilibria, each placing probability slightly higher than 0.25 on a different action profile; the first equilibrium has a higher probability on action profile (1,1), the second on (1,2), the third on (2,1) and the last one on (2,2). Suppose the vector of observed payoffs is $v = (\delta, -\delta, \delta, -\delta)$, where v(i) is the payoff for the i^{th} equilibrium. Note that there exists a constant C such that for all $\epsilon > 0$ small enough, $||E^{-1}||_{+\infty} \leq \frac{C}{\epsilon}$.

In the rest of the example, we fix the payoff matrix of player 2 for all considered games to be all-zero so that it is consistent with every equilibrium observation, and describe a game through the payoff matrix of player 1. Let G be the all-zero game, $G^1 = G^3$ be the game with payoff $\frac{\delta}{0.5+2\epsilon/3}$ on actions (1,1) and (1,2) and 0 everywhere else, and $G^2 = G^4$ be the game with payoff $-\frac{\delta}{0.5+2\epsilon/3}$ on actions (2,1) and (2,2) and 0 everywhere else. The G^i 's are consistent with the payoff observations as the payoffs are constant across rows

	_	-	

on the same column, making no deviation profitable, and the payoff of equilibria e^1 and e^3 on G^1 and G^3 is indeed δ , and $-\delta$ for e^2 and e^4 on G^2 and G^4 . We have

$$d_{\infty}(G^1, G^2, G^3, G^4 | G) = \frac{\delta}{0.5 + 2\epsilon/3} \le 2\delta$$

and

$$\lim_{\epsilon \to 0} d_{\infty}(G^1, G^2, G^3, G^4 | G) = 2\delta.$$

Now, take \hat{G} to be the game that has payoff δ/ϵ for action profiles (1,1) and (1,2), and $-\delta/\epsilon$ for (2,1) and (2,2). Take $\hat{G}^1 = \hat{G}^3$ to be the game with payoffs $\frac{\delta}{\epsilon}$ in the first column, and $-\frac{\delta}{\epsilon}\frac{3-2\epsilon}{3-4\epsilon}$ in the second column; similarly, take $\hat{G}^2 = \hat{G}^4$ to be the game with payoffs $\frac{\delta}{\epsilon}\frac{3-2\epsilon}{3-4\epsilon}$ in the first column and $-\frac{\delta}{\epsilon}$ in the second column. The observations are equilibria of the \hat{G}^i 's and yield payoff δ . Now, note that for $\epsilon < 3/4$,

$$d_{\infty}(G^1, G^2, G^3, G^4 | G) = \frac{\delta}{\epsilon} \left| 1 - \frac{3 - 2\epsilon}{3 - 4\epsilon} \right| = \frac{2}{3 - 4\epsilon} \delta.$$

Therefore, both G and \hat{G} are good explanations of the equilibrium observations, in the sense that for $\epsilon \leq 1/4$, G is δ -close to $G^1, ..., G^l$ and \hat{G} is δ -close to $\hat{G}^1, ..., \hat{G}^l$ that have $e^1, ..., e^l$ as equilibria, respectively. However,

$$\|G - \hat{G}\|_{\infty} = \frac{\delta}{\epsilon} - \frac{\delta}{0.5 + 2\epsilon/3} \ge \delta\left(\frac{1}{\epsilon} - 2\right),$$

which immediately implies

$$\|G - \hat{G}\|_{\infty} = \Omega_{\epsilon \to 0} \left(\frac{\delta}{\epsilon}\right) = \Omega_{\epsilon \to 0} \left(\|E^{-1}\|_{\infty}\delta\right).$$

Remark 3. In the case of sparse games, in which some action profiles are never profitable to the players, and are therefore never played, one can reduce the number of linearly independent, well-conditioned observations needed for accurate recovery. Under the assumption that the action profiles that are never played with positive probability have payoffs strictly worse than the lowest payoff of any action profile played with non-zero probability, one can solve the optimization problem on the restricted set of action profiles that are observed in at least one equilibrium, and set the payoffs of the remaining action profiles to be lower than the lowest payoff of the recovered subgame, without affecting the equilibrium structure of the game. While the recovered game may not be the unique good explanation of the observations when looking at the full payoff matrix, it is unique with respect to the subgame of non-trivial actions when one has access to sufficiently many linearly independent, well-conditioned equilibrium observations.

5 Extensions

In this section, we show extensions of our framework—first to succinct games with many players, and second to some games with infinite action spaces.

5.1 Linear succinct games (as per [KS15])

In general, computational tractability cannot be achieved as the number of players increases. A reason for this is that in the general case, an intractable, exponential number (in the number of players) of variables need be used to represent the game and its equilibria: in a game with n players and m actions per player, there are m^n pure action profiles, hence m^n variables are needed simply to represent the payoff matrices and the equilibria of the recovered games.

However, if the game and the observed equilibria have a compact representation, the equilibrium constraints can be written down using a tractable number of variables, and our framework provides efficient algorithms to find an element in the consistent set, compute its diameter, and test for linear properties. Kuleshov and Schrijvers [KS15] consider linear succinct games and make the interesting observation that if the structure of the succinct game is known and if we observe an equilibrium such that the "equilibrium summation property holds" (roughly, the exact expected utility of the players can be computed efficiently), then a game is consistent with the equilibrium observations if and only if a polynomial number of tractable, linear constraints are satisfied. Such constraints can easily be incorporated into our framework. (See Property 1 and Lemma 1 of [KS15] for more details.)

5.2 Cournot competition and infinite action space

In the general case, our framework cannot directly deal with games with infinite action spaces in a tractable way: to write down an equilibrium constraint, one needs a constraint for each of the infinite number of possible deviations. In this section, we show that, nevertheless, for some games with infinite action spaces, only a finite, tractable number of constraints is needed to characterize the equilibria; we show how to adapt our framework to such games. For illustration, we focus on the Cournot competition game with continuous spaces of production levels.

Consider a Cournot competition with n players selling the same good. Each player i chooses a production level $q_i \ge 0$, and sells all produced goods at price $P(q_1, ..., q_n)$ common to all players, and each player i incurs a production cost $c_i(q_i)$ to produce q_i units of the good; we write G=(P,c) where $c = (c_1, ..., c_n)$. We assume that P is concave in each q_i .

We assume the observer knows the function P and wants to recover the costs c_i of the players, where the underlying costs change slightly with each observation. Formally, consider that we have l perturbed games such that in every pertubed game k, the players play a Cournot competition with the same, commonly known price function P but perturbed cost functions c_i^k for each player i, known to be convex. We obtain equilibrium observations q^1, \ldots, q^l , where q_i^k is the equilibrium production level of player i in perturbed game k and $q^k = (q_1^k, \ldots, q_n^k)$, that is, $G^k = (P, c^k)$.

Suppose the following hold:

• The observer knows the costs belong to the space of polynomials of any chosen fixed degree $d \ge 1$; i.e., the observer parameterizes the underlying and perturbed cost functions in the following way:

$$c_i(q_i) = \sum_{ex=1}^d a_i(ex)q_i^{ex} \tag{5}$$

where the $a_i(ex)$'s are now the variables the observer want to recover.

- $d(c^1, ..., c^l | c) \leq \delta$ can be written as a tractable number of semidefinite constraints on the a_i^k 's and a_i 's (this includes, but is not limited to, the d_2 and d_{∞} distances).
- P(q) and $\frac{\partial P}{\partial q_i}$ can be computed efficiently for *i*, given *q*.

Then $S_d(\delta)$ has an efficiently computable and tractable representation (as a function of n, l and d) as the intersection of SDP constraints. This means in particular that optimizing a linear function over $S_d(\delta)$ can be cast as a tractable semidefinite program, for which efficient solvers are known – see [BV04]; one such solver, that we use in the simulations of Section 6, is CVX [GB14]. This enables us to efficiently recover a game in the consistent set, efficiently compute its diameter, and efficiently test for linear properties (by simply adding linear and thus SDP constraints, as needed).

To obtain such a tractable characterization, we only need to note that i) the equilibrium constraints can be rewritten as a tractable number of tractable linear constraints, and ii) convexity constraints on polynomials can be classically cast as tractable SDP constraints. This is the object of Appendix C.1 and C.2.

6 Simulations for entry games and Cournot competition

In this section, we run simulations for two concrete settings to illustrate the power of our approach. We first (Section 6.1) illustrate how our framework performs on a simple entry game. We then (Section 6.2) show that it is able to handle much larger games.

6.1 2-player entry game

We first consider an *entry game*, in which each of two players (think of them as companies deciding whether to open a store in a new location) has two actions available to him (enter the market; don't enter the market). Entry games are common in the literature, as seen in [ABJ04, CT09, BMM11], and, because of their simplicity, allow us to cleanly visualize the consistent region.

Each player p has two actions: $\mathcal{A}_p = \{0, 1\}$; $a_p = 0$ if player p does not enter the market, $a_p = 1$ if he does. The utility of a player is given by $G_p(a_p, a_{-p}) = a_p((1 - a_{-p})\gamma_p + a_{-p}\theta_p)$ for some parameters $\gamma_p \ge 0$ and $\theta_p \le \gamma_p$, similarly to [Tam03]: if player p does not enter the market, his utility is zero; if he enters the game but the other player does not, p has a monopoly on the market and gets non-negative utility; finally, if both players enter the game, they compete with each other and get less utility than if they had a monopoly.

In our simulations, we fix values for the parameters (γ_p, θ_p) and generate the perturbed games as follows:

- In the partial payoff information setting, we add independent Gaussian noise with mean 0 and standard deviation σ to $G_p(a_p = 1, a_{-p})$ (we vary the value of σ) to obtain the perturbed games $G^1, ..., G^l$.
- In the payoff shifter information setting, we sample the payoff shifters $\beta^1, ..., \beta^l$ such that for all $k \in [l]$, for all players p, $\beta_p^k(a_p = 1, a_{-p})$ follows a normal distribution of mean 0 and standard deviation σ_s . We then add Gaussian noise with mean 0 and standard deviation σ to $G_p(a_p = 1, a_{-p})$ to obtain the perturbed games $G^1, ..., G^l$.

In both observation models, paralleling the setting of [BMM11], no observed payoff shifter nor unknown noise is added to the payoff of action $a_p = 0$ for player p; action $a_p = 0$ is always assumed to yield payoff 0 for player p, independently of a_{-p} . In order to generate the equilibrium observations, once the perturbed games are generated, we find the set of equilibria of each of the G^k , and sample a point e^k in said set. In the payoff information case, we also compute $v^k = e^k G^k$.

In order to parallel the setting of Beresteanu et al. [BMM11], we assume the observer knows the form of the utility function, i.e., that $G_p(0,0) = 0$ and $G_p(a_p = 0, a_{-p} = 1) = 0$, and that he aims to recover the values of γ_p and θ_p . Thus, we add linear constraints $G_p(0,0) = 0$ and $G_p(a_p = 0, a_{-p} = 1) = 0$ in the optimization programs that we solve (see Program (1)) in the payoff shifter information and partial payoff information settings. Furthermore, we assume as in [BMM11] that the observer knows that perturbations are only added to γ and θ , and therefore we add linear constraints $G_p^k(0,0) = 0$ and $G_p^k(a_p = 0, a_{-p} = 1) = 0$ for all $k \in [l]$ to the optimization problems for player p in each of the observation models. All optimization problems are solved in Matlab, using CVX (see [GB14]).

Our model for entry-games is similar to the ones presented in [Tam03] and used in simulations in [BMM11], so as to facilitate informal comparisons of the simulation results of both papers; in particular, the parametrization of the utility functions of the players in our simulations is inspired by [BMM11], and noise is generated and added in a similar fashion. However, while we attempt to parallel the simulations run by Beresteanu et al. [BMM11], it is important to note that this is not an apples-to-apples comparison, because of key differences in the setting. In particular, our observation models (seeing full equilibria) and the information available to the observer (no distributional assumptions) are different from those in [BMM11].

6.1.1 Consistent regions for Player 1

We fix l = 500, $\gamma = 5$, $\theta = -10$ in all simulations, and vary the values of σ and σ_s . Because the observations are generated by adding i.i.d Gaussian noise with mean 0 and variance σ^2 to the two payoffs for entry of each player, if G is the underlying game and $G^1, ..., G^l$ are its perturbations,

$$\frac{1}{\sigma^2} d_2(G^1, ..., G^l | G) \text{ (resp. } \frac{1}{\sigma^2} d_2(G^1 - \beta^1, ..., G^l - \beta^l | G))$$

follows a Chi-square distribution with 4l degrees of freedom in the partial payoff information case (resp. in the payoff shifter case). We choose δ such that $P(d_2(G^1, ..., G^l | G) \leq \delta) \approx 0.99$, and suppose the observer sees said value of δ . While the observer may not have access to the distribution of perturbations, it is extremely likely he will experience a magnitude of perturbations equal to or less than δ , and we can use δ as a high-probability upper bound on the information on the perturbations accessible to the observer. It may also be the case that the observer knows the distribution of perturbations, in which case the observer could pick such a value of δ to guarantee that with 0.99 probability, the recovered set covers the true, underlying game.



Figure 1: Plots of the consistent region for different values of σ , σ_s in the payoff shifter information observation model

In all plots, the colored region in the plots is the projection over the space (γ_1, θ_1) for player 1 of the set of parameters $(\gamma_1, \theta_1, \gamma_2, \theta_2)$ that are in the δ -consistent region. The darker the region, the smaller the objective value of the best explanation for the corresponding values of γ and θ . The black, center of the region represents the value of (γ, θ) that minimizes $d_2(G^1, ..., G^l|G)$. We note that darker colors define consistent regions with smaller values of δ ; therefore the heatmap allows the observer to visualize how varying δ affects the set of consistent games.

Figure 1 shows the evolution of the consistent region when varying σ and σ_s in the payoff shifter information setting. The smaller the standard deviation σ of the unknown noise, the tighter the consistent region. On the other hand, reasonably increasing the value of σ_s can be beneficial, at least when it comes to centering the consistent region on the true values of the parameters: this comes from the fact that when the game is sufficiently perturbed, new equilibria arise and new, informative behavior is observed, while not adding significant uncertainty to the payoffs of the game.

Figure 2 shows the evolution of the consistent region when varying σ . The larger the value of σ , the larger the consistent region, and the further away its center is from the underlying, true value of the parameters.



Figure 2: Plots of the consistent region for different values of σ in the partial payoff information observation model

6.1.2 Testing for linear properties

We also illustrate via simulation how our framework can test the ability of linear properties to explain observed behavior. In particular, here we test whether a set of observations is likely to be explained by a zero-sum game. We consider entry games as defined in the previous section, and assume the observer wants to test whether observations were generated by a game that is approximately zero-sum, without any information on the parametric form of the game (the observer does not know the game is an entry game).

Formally, we say a game $G = (G_1, G_2)$ is ε -zero-sum with respect to the *p*-norm if and only if $||G_1 + G_2||_p \le \varepsilon$. Note that a game being ε -zero-sum is a linear property and therefore can be included in our framework.



Figure 3: Testing for zero-sum with respect to the 1-norm

The smaller the value of ε , the more stringent the condition is and the closer G must be to a zero-sum game. We use l = 500, $\sigma = 0.5$, $\sigma_s = 10$ in all simulations.

As before, we pessimistically assume the observer sees δ such that

 $P(d_2(G^1, ..., G^l | G) \leq \delta) \approx 0.99$, that is, $\delta = 537.5$ for l = 500, $\sigma = 0.5$. Figure 3 shows for which values of ε one can recover a ε -zero-sum game with objective value less than 537.5 that explains the observations for different values of γ and θ . Values of ε to the left of the intersection between the red and the blue line are impossible, while values to the right of this intersection indicate there is a ε -zero-sum game that explain the observations. In both cases, we see that no zero-sum game or game close to being zerosum is a good explanation for the observations; in the payoff information setting, no game less than 21-zerosum explains the observations, while in the payoff shifter setting, no game less than 15-zerosum explains the observations.

6.2 Multiplayer Cournot competition

In this section, we run simulations on a Cournot competition with varying number of players. See Section 5.2 for a discussion of how our framework can be modified to accommodate Cournot games with many players and an action set of infinite size for each player. All simulations were performed on a laptop with an Intel Core i7-4700MQ at 2.40GHz and 16 GB RAM.

6.2.1 Generating the games

Let *n* be the number of players, and q_i the production level of player *i*. We fix a parameter $\alpha = 0.05$, and set the price function to be given by $P(q_1, ..., q_n) = 1 - \alpha \sum_{i=1}^n q_i$; the price function is known to the observer. We fix the form of the cost function to be linear; that is, the cost of producing q_i of goods incurred by player *i* is given by $c_i(q_i) = a_i q_i + b_i$. Without loss of generality, we set $b_i = 0$: b_i does not affect the maximization problem nor the first order condition solved by player *i*, and hence does not impact the chosen production level of the players.

We generate $n_g = 10$ underlying Cournot games with heterogeneous, linear cost functions $c_i(x) = a_i x$ as follows:

- We first set $\hat{a}_i = 0.01$ for every player *i*.
- We generate each of the n_g games by adding i.i.d. truncated Gaussian noise X_i with mean 0 and standard deviation 0.01 to the \hat{a}_i 's. I.e., $a_i = 0.01 + X_i$ where X_i can be written as $X_i = max(Z_i, -0.01)$



Figure 4: Average recovery time as a function of the degree with 10 players, 10 equilibrium observations

and Z_i is a non-truncated Gaussian with mean 0, standard deviation 0.01. This ensures the a_i 's are always non-negative, hence the production costs are always non-decreasing.

Note that the same n_g games are used in all plots and simulations.

For each of the n_g games, we then generate l perturbed games by adding truncated Gaussian noise with standard deviation $\sigma = 0.001$ to each of the a_i 's. As before, the noise is truncated to ensure non-negativity of the perturbed a_i 's. We then solve the first-order condition to obtain equilibrium observations and note that all obtained q_i 's are positive with extremely high probability.

6.2.2 Observer's problem and simulation results

We assume the observer wants to recover a cost that is polynomial of chosen degree $d \ge 1$. I.e., the observer parameterizes the cost functions in the following way: $c_i(x) = a_i(d)x^d + a_i(d-1)x^{d-1} + ... + a_i(1)x$; we always assume that for every player *i*, every perturbed game k, $a_i(0) = 0$ and $a_i^k(0) = 0$ for simplicity (not producing anything costs the players nothing); we note that this without loss of generality, as a constant shift does not change the utility-maximizing strategies of the players. The observer also knows the perturbed cost function c_i^k 's of the perturbed games are convex. The program then solved by the observer is derived from the results of Section 5.2. Throughout this subsection, all results are averaged over the n_g games originally generated: for each game, we measure the diameter of the consistent set, the time taken to recover a game within the set, and the time taken to compute the diameter, then average it over the n_g games we consider.

Figure 4 shows the time it takes the observer to recover a game within the consistent set as a function of the degree of the polynomial cost function considered, fixing the number of players to 10 and the number of perturbed games/equilibrium observations to 50 per underlying game. We see the recovery time is less than a minute, even when considering polynomials of degree 10 or 12, and that said time increases roughly linearly in the degree of the polynomial used for recovery. This allows for recovery within minutes for high-degree polynomials, even when using the minimal computing power of a personal laptop.

Figure 5 shows the value of the diameter of the consistent set when the observer assumes the cost function is linear. The diameter is plotted as a function of the number of observations, in the presence of a fixed number of players (10). The figure shows the diameter decreases quickly as the number of equilibrium observations. When only one equilibrium observation is available, the diameter is 3.3×10^{-3} , which is 35%of the expected true cost $\hat{a}_i = 0.01$; at 10 equilibrium observations, it is 3.3×10^{-4} , i.e. only 3.5% of \hat{a}_i , and at 100 equilibrium observations, it is 6.6×10^{-5} , i.e. 0.66% of \hat{a}_i . Hence, very few equilibrium observations are necessary to recover the underlying game accurately.



Figure 5: Average diameter as a function of the number equilibrium observations



Figure 6: Average recovery/diameter time as a function of the number of observations/number of players, when recovering linear costs

Figures 6a and 6b show the time it takes the observer to recover a game within the consistent set and to compute its diameter as a function of the number of players, fixing the degree of the polynomial to 1 and equilibrium observations to 50 per underlying game. Similarly to before, we see that the recovery time increases roughly linearly as a function of the number of players, and it takes less than 1.5 seconds on average to recover a consistent game, even with 50 players. The diameter time increases roughly quadratically: this was expected, as both the number of programs to solve and the size of each program increase linearly in the number of players. Hence, while computing the diameter scales superlinearly, it remains computable within minutes, even with larger numbers of players, on a personal laptop. To the best of our knowledge, no previous approach to understanding the games consistent with observed behavior offers comparable scalability.

Finally, Figures 6c and 6d show the time it takes the observer to recover a game within the consistent set and to compute its diameter as a function of the number of equilibria, fixing the degree of the polynomial to 1 and the number of players to 10 per underlying game. Both the recovery time and the diameter time scale linearly with the number of equilibrium observations. Unlike before, while the size of each subprogram to solve to compute the diameter increases, the number of such programs is independent of the number of observations, allowing for extremely good scalability of our framework as function of the number of equilibrium observations. In practice, the number of players is fixed but the observer may see more observations over time, and our framework is able to deal with such an increasing number of observations.

Acknowledgments

We thank Federico Echenique, Denis Nekipelov, Matt Shum, and Vasilis Syrgkanis for extremely helpful comments and suggestions. Ziani's research was funded in part by NSF grant CNS-1254169. Chandrasekaran's research was funded in part by NSF awards CCF-1350590 and CCF-1637598, Air Force Office of Scientific Research grants FA9550-14-1-0098 and FA9550-16-1-0210, and the Sloan research fellowship. Ligett's research was funded in part by ISF grant 1044/16, NSF grants CNS-1254169 and CNS-1518941, a subcontract under the DARPA Brandeis Project, and the Hebrew University Cybersecurity Research Center in conjunction with the Israel National Cyber Bureau in the Prime Minister's Office. Ligett's work was done in part while the author was visiting the Simons Institute for the Theory of Computing at Berkeley.

Bibliography

- [ABJ04] Donald W. K. Andrews, Steven Berry, and Panle Jia. Confidence regions for parameters in discrete games with multiple equilibria, with an application to discount chain store location. Unpublished manuscript, Yale University, 2004.
- [AL10] Andres Aradillas-Lopez. Semiparametric estimation of a simultaneous game with incomplete information. Journal of Econometrics, 157(2):409 431, 2010.
- [AL11] Andres Aradillas-Lopez. Nonparametric probability bounds for nash equilibrium actions in a simultaneous discrete game. *Quantitative Economics*, 2(2):135–171, 2011.
- [AL12] Andres Aradillas-Lopez. Pairwise-difference estimation of incomplete information games. Journal of Econometrics, 168(1):120 – 140, 2012. The Econometrics of Auctions and Games.
- [BB12] Chaithanya Bandi and Dimitris Bertsimas. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming*, 134(1):23–70, 2012.
- [BBC11] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.
- [BBEW13] Siddharth Barman, Umang Bhaskar, Federico Echenique, and Adam Wierman. The empirical implications of rank in bimatrix games. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 55–72, New York, NY, USA, 2013. ACM.
- [BHHR11] Patrick Bajari, Jinyong Hahn, Han Hong, and Geert Ridder. A note on semiparametric estimation of finite mixtures of discrete choice models with application to game theoretic models. *International Economic Review*, 52(3):807–824, 08 2011.
- [BHKN10] Patrick Bajari, Han Hong, John Krainer, and Denis Nekipelov. Estimating static models of strategic interactions. Journal of Business & Economic Statistics, 28(4):469–482, 2010.
- [BLSS14] Umang Bhaskar, Katrina Ligett, Leonard J. Schulman, and Chaitanya Swamy. Achieving target equilibria in network routing games without knowing the latency functions. *CoRR*, abs/1408.1429, 2014.
- [BMM11] Arie Beresteanu, Ilya Molchanov, and Francesca Molinari. Sharp identification regions in models with convex moment predictions. *Econometrica*, 79(6):1785–1821, 2011.
 - [BR91] Timothy Bresnahan and Peter C. Reiss. Empirical models of discrete games. Journal of Econometrics, 48(1-2):57–81, 1991.
- [BTEGN09] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Princeton University Press, 2009.
 - [BV84] Paul A. Bjorn and Quang H. Vuong. Simultaneous equations models for dummy endogenous variables: a game theoretic formulation with an application to labor force participation. Working Papers 537, California Institute of Technology, Division of the Humanities and Social Sciences, 1984.
 - [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

- [CDT07] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. CoRR, abs/0704.1678, 2007.
- [CHT07] Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007.
- [CT09] Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. Econometrica, 77(6):1791–1828, 2009.
- [GB14] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, March 2014.
- [GJN17] Vincent Guigues, Anatoli Juditsky, and Arkadi Nemirovski. Non-asymptotic confidence bounds for the optimal value of a stochastic program. Optimization Methods and Software, 32(5):1033– 1058, 2017.
- [HM10] Jinyong Hahn and Hyungsik Roger Moon. Panel data models with finite number of multiple equilibria, 2010.
- [KS15] Volodymyr Kuleshov and Okke Schrijvers. Inverse game theory: Learning utilities in succinct games. In Proceedings of the 11th International Conference on Web and Internet Economics - Volume 9470, WINE 2015, pages 413–427, New York, NY, USA, 2015. Springer-Verlag New York, Inc.
- [NST15] Denis Nekipelov, Vasilis Syrgkanis, and Éva Tardos. Econometrics for learning agents. CoRR, abs/1505.00720, 2015.
- [Par04] Pablo A Parrilo. Sum of squares programs and polynomial inequalities. In SIAG/OPT Viewsand-News: A Forum for the SIAM Activity Group on Optimization, volume 15, pages 7–15, 2004.
- [RRUW15] Ryan Rogers, Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Inducing approximately optimal flow using truthful mediators. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 471–488, New York, NY, USA, 2015. ACM.
 - [Sei06] Katja Seim. An empirical model of firm entry with endogenous product-type choices. RAND Journal of Economics, 37(3):619–640, 2006.
 - [Tam03] Elie Tamer. Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies*, 70(1):147–165, 2003.

A Proof of performance of the algorithm

Consider the following optimization program:

$$P_{\delta} = \sup_{\substack{\tilde{G}, \hat{G}, \gamma \\ \text{s.t.}}} \gamma$$

s.t.
$$\begin{array}{c} \tilde{G} \in S_d(\delta) \\ \tilde{G} \in S_d(\delta) \\ \max(\|\tilde{G}_1 - \hat{G}_1\|_{\infty}, \|\tilde{G}_2 - \hat{G}_2\|_{\infty}) \ge \gamma \end{array}$$

Clearly, $P_{\delta} = D(S_d(\delta))$, simply by noting that the program is a rewriting of Definition 5. Now if $P_{\delta} = \max_{(i,j)\in m_1m_2} \max(P_{\delta,1}(i,j), P_{\delta,2}(i,j))$, then we have shown that $\mathcal{A}(\delta) = D(S_d(\delta))$. This holds because:

- For every player p and action profile (i, j), if $\tilde{G}_p(i, j) \hat{G}_p(i, j) \ge \gamma$ then $\max(\|\tilde{G}_1 \hat{G}_1\|_{\infty}, \|\tilde{G}_2 \hat{G}_2\|_{\infty}) \ge \gamma$. Hence $P_{\delta} \ge \max_{p, (i, j)} P_{\delta, p}(i, j)$.
- If $\max(\|\tilde{G}_1 \hat{G}_1\|_{\infty}, \|\tilde{G}_2 \hat{G}_2\|_{\infty}) \ge \gamma$, then there exists a player p and a set of actions (i, j) such that $\tilde{G}_p(i, j) \hat{G}_p(i, j) \ge \gamma$ w.l.o.g. (remember \tilde{G} and \hat{G} play symmetric roles) and so one of the $P_{\delta,p}(i, j)$ must have objective value at least γ . This means that $P_{\delta} \le \max_{p,(i,j)} P_{\delta,p}(i, j)$.

B Proof of recovery lemma under infinite norm and payoff information

Proof. For simplicity of notation, we drop the indices p. We first remark that (G, G^1, \ldots, G^l) is feasible for Program (4); as $(\hat{G}, \hat{G}^1, \ldots, \hat{G}^l)$ is optimal, it is necessarily the case that

$$\max_{k} \|\hat{G} - \hat{G}^{k}\|_{\infty} \le \max_{k} \|G - G^{k}\|_{\infty} \le \delta.$$

Let us write $\Delta G = G - \hat{G}$. We know that for all k, $e^{k'}G^k = e^{k'}\hat{G}^k = v^k$, and thus $e^{k'}(G^k - \hat{G}^k) = 0$. We can write

$$\begin{split} E\Delta G &= (e_1'(G-\hat{G}) \ \dots \ e_l'(G-\hat{G}))' \\ &= (e_1'(G-G^1+G^1-\hat{G}^1+\hat{G}^1-\hat{G}) \\ \qquad \dots \ e_l'(G-G^l+G^l-\hat{G}^l+\hat{G}^l-\hat{G}))' \\ &= (e_1'(G-G^1+\hat{G}^1-\hat{G}) \ \dots \ e_l'(G-G^l+\hat{G}^l-\hat{G}))'. \end{split}$$

Let $x_k = G - G^k + \hat{G}^k - \hat{G}$. We then have $||E\Delta G||_{\infty} \leq \max_k ||x_k||_{\infty}$ as e^k has only elements between 0 and 1. Therefore, by the triangle inequality,

$$||E\Delta G||_{\infty} \le 2\delta.$$

It immediately follows that $\|\Delta G\|_{\infty} \leq 2\|E^{-1}\|_{\infty} \cdot \delta$.

C Writing Cournot constraints efficiently

C.1 Casting the equilibrium constraints as linear constraints

When action profile (q_1, \ldots, q_n) is chosen, player *i* gets utility $u_i(q_i, q_{-i}) = q_i P(q_1, \ldots, q_n) - c_i(q_i)$ where q_{-i} denotes the production levels of all players but *i*. A pure action profile $q^* = (q_1^*, \ldots, q_n^*)$ is a Nash Equilibrium if and only if for all players *i*, q_i^* maximizes $u_i(q_i, q_{-i}^*)$; as *P* is concave and c_i is convex, u_i is convex in q_i and the equilibrium condition is equivalent to the first order condition

$$q_i \frac{\partial P}{\partial q_i}(q_1, \dots, q_n) + P(q_1, \dots, q_n) = c'_i(q_i), \ \forall i$$
(6)

Then, combining Equations 6 and 5, the equilibrium constraints become

$$q_i \frac{\partial P}{\partial q_i}(q_1, \dots, q_n) + P(q_1, \dots, q_n) = \sum_{k=1}^d k a_i(k) q_i^{k-1}$$

$$\tag{7}$$

which are linear and tractable in the variables $(a_i(0), a_i(1), \ldots, a_i(d)))$, as long as $P(q_1, \ldots, q_n)$, $\frac{\partial P}{\partial q_i}(q_1, \ldots, q_n)$ can be efficiently computed given observations q_1, \ldots, q_n . Such equilibrium constraints can be incorporated into our framework.

C.2 Casting the convexity constraints as SDP constraints

We need to be able to deal with " c_i is convex polynomial of degree d" constraints for all i, in a computationally efficient manner. This constraint can be rewritten as " c''_i is a non-negative polynomial of degree d." Fortunately, this is a classic constraint in the realm of convex optimization, and can be dealt with in the following ways:

- If d = 1, then $c_i(q_i) = a_i(1)q_i + a_i(0)$ $(c''_i = 0)$ is always convex. In this case, no constraint need be added.
- If $d \ge 2$, then we need to ensure that c''_i is non-negative in all points. It is known that a univariate polynomial is non-negative if and only if it can be written as a sum-of-squares; such constraints can efficiently be transformed into tractable semidefinite constraints. (For more details on the SDP formulation of sum-of-squares constraints, see [Par04].)