

From Behavior to Sparse Graphical Games: Efficient Recovery of Equilibria

Asish Ghoshal
Department of Computer Science
Purdue University
West Lafayette, IN - 47907
Email: aghoshal@purdue.edu

Jean Honorio
Department of Computer Science
Purdue University
West Lafayette, IN - 47907
Email: jhonorio@purdue.edu

Abstract—In this paper we study the problem of exact recovery of the pure-strategy Nash equilibria (PSNE) set of a graphical game from noisy observations of joint actions of the players alone. We consider sparse linear influence games — a parametric class of graphical games with linear payoffs, and represented by directed graphs of n nodes (players) and in-degree of at most k . We present an ℓ_1 -regularized logistic regression based algorithm for recovering the PSNE set exactly, that is both computationally efficient — i.e. runs in polynomial time — and statistically efficient — i.e. has logarithmic sample complexity. Specifically, we show that the sufficient number of samples required for exact PSNE recovery scales as $\mathcal{O}(\text{poly}(k) \log n)$. We also validate our theoretical results using synthetic experiments.

I. INTRODUCTION AND RELATED WORK

Non-cooperative game theory is widely regarded as an appropriate mathematical framework for studying *strategic* behavior in multi-agent scenarios. The core solution concept of *Nash equilibrium* describes the stable outcome of the overall behavior of self-interested agents — for instance people, companies, governments, groups or autonomous systems — interacting strategically with each other and in distributed settings.

Over the past few years, considerable progress has been made in analyzing behavioral data using game-theoretic tools, e.g. computing Nash equilibria [1], [2], [3], most influential agents [4], price of anarchy [5] and related concepts in the context of graphical games. In *political science* for instance, Irfan and Ortiz [4] identified, from congressional voting records, the most influential senators in the U.S. congress — a small set of senators whose collective behavior forces every other senator to a unique choice of vote. Irfan and Ortiz [4] also observed that the most influential senators were strikingly similar to the gang-of-six senators, formed during the national debt ceiling negotiations of 2011. Further, using graphical games, Honorio and Ortiz [6] showed that Obama’s influence on Republicans increased in the last sessions before candidacy, while McCain’s influence on Republicans decreased.

The problems in *algorithmic game theory* described above, i.e. computing the Nash equilibria, computing the price of anarchy or finding the most influential agents, require a known graphical game which is not available apriori in real-world settings. Therefore, Honorio and Ortiz [6] proposed

learning graphical games from behavioral data, using maximum likelihood estimation (MLE) and *sparsity*-promoting methods. Honorio and Ortiz [6] and Irfan and Ortiz [4] have also demonstrated the usefulness of learning *sparse* graphical games from behavioral data in real-world settings, through their analysis of the voting records of the U.S. congress as well as the U.S. supreme court.

In this paper, we analyze a particular method proposed by Honorio and Ortiz [6] for learning sparse linear influence games, namely, using logistic regression for learning the neighborhood of each player in the graphical game, independently. Honorio and Ortiz [6] showed that the method of independent logistic regression is likelihood consistent; i.e. in the infinite sample limit, the likelihood estimate converges to the true achievable likelihood. In this paper we obtain the stronger guarantee of recovering the true PSNE set exactly. Needless to say, our stronger guarantee comes with additional conditions that the true game must satisfy in order to allow for exact PSNE recovery. The most crucial among our conditions is the assumption that the minimum payoff across all players and all joint actions in the true PSNE set is strictly positive. We show, through simulation experiments, that our assumptions indeed bears out for large classes of graphical games, where we are able to exactly recover the PSNE set.

Finally, we would like to draw the attention of the reader to the fact that ℓ_1 -regularized logistic regression has been analyzed by Ravikumar et. al. [7] in the context of learning sparse Ising models. Apart from technical differences and differences in proof techniques, our analysis of ℓ_1 -penalized logistic regression for learning sparse graphical games differs from Ravikumar et. al. [7] conceptually — in the sense that we are not interested in recovering the edges of the true game graph, but only the PSNE set. Therefore, we are able to avoid some stronger and non-intuitive conditions required by Ravikumar et. al. [7], such as mutual incoherence.

The rest of the paper is organized as follows. In section II we provide a brief overview of graphical games and, specifically, linear influence games. We then formalize the problem of learning linear influence games in section III. In section IV, we present our main method and associated theoretical results. Then, in section V we provide experimental validation of our theoretical guarantees. Finally, in section VI we conclude by

discussing avenues for future work.

II. PRELIMINARIES

In this section we review key concepts behind graphical games introduced by Kearns et. al. [8] and linear influence games (LIGs) introduced by Irfan and Ortiz [4] and Honorio and Ortiz [6].

A. Graphical Games

A *normal-form game* \mathcal{G} in classical game theory is defined by the triple $\mathcal{G} = (V, \mathcal{A}, \mathcal{U})$ of players, actions and payoffs. V is the set of players, and is given by the set $V = \{1, \dots, n\}$, if there are n players. \mathcal{A} is the set of actions or *pure-strategies* and is given by the Cartesian product $\mathcal{A} \doteq \times_{i \in V} A_i$, where A_i is the set of pure-strategies of the i -th player. Finally, $\mathcal{U} \doteq \{u_i\}_{i=1}^n$, is the set of payoffs, where $u_i : A_i \times \prod_{j \in V \setminus i} A_j \rightarrow \mathbb{R}$ specifies the payoff for the i -th player given its action and the joint actions of the all the remaining players.

A key solution concept in non-cooperative game theory is the *Nash equilibrium*. For a non-cooperative game, a joint action $\mathbf{x}^* \in \mathcal{A}$ is a pure-strategy Nash equilibrium (PSNE) if, for each player i , $x_i^* \in \operatorname{argmax}_{x_i \in A_i} u_i(x_i, \mathbf{x}_{-i}^*)$, where $\mathbf{x}_{-i}^* = \{x_j^* | j \neq i\}$. In other words, \mathbf{x}^* constitutes the mutual best-response for all players and no player has any incentive to unilaterally deviate from their optimal action x_i^* given the joint actions of the remaining players \mathbf{x}_{-i}^* . The set of all *pure-strategy Nash equilibrium* (PSNE) for a game \mathcal{G} is defined as follows:

$$\mathcal{NE}(\mathcal{G}) = \left\{ \mathbf{x}^* \mid (\forall i \in V) x_i^* \in \operatorname{argmax}_{x_i \in A_i} u_i(x_i, \mathbf{x}_{-i}^*) \right\}. \quad (1)$$

Graphical games, introduced by Kearns et. al. [8], extend the formalism of *Graphical models* to games. That is, a graphical game G is defined by the *directed graph*, $G = (V, E)$, of vertices and directed edges (arcs), where vertices correspond to players and arcs encode “influence” among players i.e. the payoff of the i -th player only depends on the actions of its (incoming) neighbors.

B. Linear Influence Games

Linear influence games (LIGs), introduced by Irfan and Ortiz [4] and Honorio and Ortiz [6], are graphical games with binary actions, or pure strategies, and parametric (linear) payoff functions. We assume, without loss of generality, that the joint action space $\mathcal{A} = \{-1, +1\}^n$. A linear influence game between n players, $\mathcal{G}(n) = (\mathbf{W}, \mathbf{b})$, is characterized by (i) a matrix of weights $\mathbf{W} \in \mathbb{R}^{n \times n}$, where the entry W_{ij} indicates the amount of influence (signed) that the j -th player has on the i -th player and (ii) a bias vector $\mathbf{b} \in \mathbb{R}^n$, where b_i captures the prior preference of the i -th player for a particular action $x_i \in \{-1, +1\}$. The payoff of the i -th player given the actions of the remaining players is then given as $u_i(x_i, \mathbf{x}_{-i}) = x_i(\mathbf{w}_{-i}^T \mathbf{x}_{-i} - b_i)$, and the PSNE set is defined as follows:

$$\mathcal{NE}(\mathcal{G}(n)) = \left\{ \mathbf{x} \mid (\forall i) x_i(\mathbf{w}_{-i}^T \mathbf{x}_{-i} - b_i) \geq 0 \right\}, \quad (2)$$

where \mathbf{w}_{-i} denotes the i -th row of \mathbf{W} without the i -th entry, i.e. $\mathbf{w}_{-i} = \{w_{ij} | j \neq i\}$. Note that we have $\operatorname{diag}(\mathbf{W}) = 0$. For linear influence games $\mathcal{G}(n)$, we can define the neighborhood and signed neighborhood of the i -th vertex as $\mathcal{N}(i) = \{j \mid |w_{ij}| > 0\}$ and $\mathcal{N}_{\pm}(i) = \{\operatorname{sign}(w_{ij}) | j \in \mathcal{N}(i)\}$ respectively. It is important to note that we don’t include the outgoing arcs in our definition of the neighborhood of a vertex. Thus, for linear influence games, the weight matrix \mathbf{W} and the bias vector \mathbf{b} , completely specify the game and the PSNE set induced by the game. Finally, let $\mathcal{G}(n, k)$ denote sparse games over n players where the in-degree of any vertex is at most k , i.e. for all i , $|\mathcal{N}_i| \leq k$.

III. PROBLEM FORMULATION

Having introduced the necessary definitions, we now introduce the problem of learning sparse linear influence games from observations of joint actions only. We assume that there exists a game $\mathcal{G}^*(n, k) = (\mathbf{W}^*, \mathbf{b}^*)$ from which a “noisy” data set $\mathcal{D} = \{\mathbf{x}^{(l)}\}_{l=1}^m$ of m observations is generated, where each observation $\mathbf{x}^{(l)}$ is sampled independently and identically from the following distribution:

$$p(\mathbf{x}) = \frac{q \mathbf{1}[\mathbf{x} \in \mathcal{NE}^*]}{|\mathcal{NE}^*|} + \frac{(1-q) \mathbf{1}[\mathbf{x} \notin \mathcal{NE}^*]}{2^n - |\mathcal{NE}^*|}. \quad (3)$$

In the above distribution, q is the probability of observing a data point from the set of Nash equilibria and can be thought of as the “signal” level in the data set, while $1-q$ can be thought of as the “noise” level in the data set. We use \mathcal{NE}^* instead of $\mathcal{NE}(\mathcal{G}^*(n, k))$ to simplify notation. We further assume that the game is non-trivial¹ i.e. $|\mathcal{NE}^*| \in \{1, \dots, 2^n - 1\}$ and that $q \in (|\mathcal{NE}^*|/2^n, 1)$. The latter assumption ensures that the signal level in the data set is more than the noise level². We define the equality of two games $\mathcal{G}^*(n, k) = (\mathbf{W}^*, \mathbf{b}^*)$ and $\widehat{\mathcal{G}}(n, k) = (\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$ as follows:

$$\begin{aligned} \mathcal{G}^*(n, k) = \widehat{\mathcal{G}}(n, k) \text{ iff} \\ (\forall i) \mathcal{N}_{\pm}^*(i) = \widehat{\mathcal{N}}_{\pm}(i) \wedge \operatorname{sign}(b_i^*) = \operatorname{sign}(\widehat{b}_i) \wedge \mathcal{NE}^* = \widehat{\mathcal{NE}}. \end{aligned}$$

A natural question to ask then is that, given only the data set \mathcal{D} and no other information, is it possible to recover the weight matrix $\widehat{\mathbf{W}}$ and the bias vector $\widehat{\mathbf{b}}$ such that $\widehat{\mathcal{G}}(n, k) = \mathcal{G}^*(n, k)$? Honorio and Ortiz [6] showed that it is in general impossible to learn the true game $\mathcal{G}^*(n, k)$ from observations of joint actions only because multiple weight matrices \mathbf{W} and bias vectors \mathbf{b} can induce the same PSNE set and therefore have the same likelihood under the observation model (3) — an issue known as non-identifiability in the statistics literature. It is, however, possible to learn the equivalence class of games that induce the same PSNE set. We define the equivalence of two games $\mathcal{G}^*(n, k)$ and $\widehat{\mathcal{G}}(n, k)$ simply as :

$$\mathcal{G}^*(n, k) \equiv \widehat{\mathcal{G}}(n, k) \text{ iff } \mathcal{NE}^* = \widehat{\mathcal{NE}}.$$

Therefore, our goal in this paper is efficient and consistent recovery of the pure-strategy Nash equilibria set (PSNE) from

¹This comes from Definition 4 in [6].

²See Proposition 5 and Definition 7 in [6] for a justification of this.

observations of joint actions only; i.e. given a data set \mathcal{D} , drawn from some game $\mathcal{G}^*(n, k)$ according to (3), we infer a game $\widehat{\mathcal{G}}(n, k)$ from \mathcal{D} such that $\widehat{\mathcal{G}}(n, k) \equiv \mathcal{G}^*(n, k)$.

IV. METHOD AND RESULTS

Our main method for learning the structure of a sparse LIG, $\mathcal{G}^*(n, k)$, is based on using ℓ_1 -regularized logistic regression, to learn the parameters (\mathbf{w}_{-i}, b_i) for each player i independently. We denote by $\mathbf{v}_i(\mathbf{W}, \mathbf{b}) = (\mathbf{w}_{-i}, -b_i)$ the parameter vector for the i -th player, which characterizes its payoff; and by $\mathbf{z}_i(\mathbf{x}) = (x_i \mathbf{x}_{-i}, x_i)$ the ‘‘feature’’ vector. In the rest of the paper we use \mathbf{v}_i and \mathbf{z}_i instead of $\mathbf{v}_i(\mathbf{W}, \mathbf{b})$ and $\mathbf{z}_i(\mathbf{x})$ respectively, to simplify notation. Then, we learn the parameters for the i -th player as follows:

$$\widehat{\mathbf{v}}_i = \underset{\mathbf{v}_i}{\operatorname{argmin}} \ell(\mathbf{v}_i, \mathcal{D}) + \lambda \|\mathbf{v}_i\|_1 \quad (4)$$

$$\ell(\mathbf{v}_i, \mathcal{D}) = \frac{1}{m} \sum_{l=1}^m \log(1 + \exp(-\mathbf{v}_i^T \mathbf{z}_i^{(l)})). \quad (5)$$

We then set $\widehat{\mathbf{w}}_{-i} = [\widehat{\mathbf{v}}_i]_{1:(n-1)}$ and $\widehat{b}_i = -[\widehat{\mathbf{v}}_i]_n$, where the notation $[\cdot]_{i:j}$ denotes indices i to j of the vector. We show that, under suitable assumptions on the true game $\mathcal{G}^*(n, k) = (\mathbf{W}^*, \mathbf{b}^*)$, the parameters $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{b}}$ obtained using (5) induce the same PSNE set as the true game, i.e. $\mathcal{NE}(\mathbf{W}^*, \mathbf{b}^*) = \mathcal{NE}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$. Before presenting our main results, however, we state and discuss the assumptions on the true game under which it is possible to recover the true PSNE set of the game using our proposed method.

A. Assumptions

The success of our method hinges on certain assumptions on the structure of the underlying game. These assumptions are in addition to the assumptions imposed on the parameter q and the number of Nash equilibria of the game $|\mathcal{NE}^*|$, as required by definition of a LIG. Since the assumptions are related to the Hessian of the loss function (5), we take a moment to introduce the expressions of the gradient and the Hessian here. The gradient and Hessian of the loss function for any vector \mathbf{v} and the data set \mathcal{D} is given as follows:

$$\nabla \ell(\mathbf{v}, \mathcal{D}) = \frac{1}{m} \sum_{l=1}^m \left\{ \frac{-\mathbf{z}^{(l)}}{1 + \exp(\mathbf{v}^T \mathbf{z}^{(l)})} \right\} \quad (6)$$

$$\nabla^2 \ell(\mathbf{v}, \mathcal{D}) = \frac{1}{m} \sum_{l=1}^m \eta(\mathbf{v}^T \mathbf{z}^{(l)}) \mathbf{z}^{(l)} \mathbf{z}^{(l)T}, \quad (7)$$

where $\eta(x) = 1/(e^{x/2} + e^{-x/2})^2$. Finally, \mathbf{H}_i^m denotes the sample Hessian matrix with respect to the i -th player and the true parameter \mathbf{v}_i^* , and \mathbf{H}_i^* denotes its expected value, i.e. $\mathbf{H}_i^* \doteq \mathbb{E}_{\mathcal{D}}[\mathbf{H}_i^m] = \mathbb{E}_{\mathcal{D}}[\nabla^2 \ell(\mathbf{v}_i^*, \mathcal{D})]$. In subsequent sections we drop the notational dependence of \mathbf{H}_i^* and \mathbf{z}_i on i to simplify notation. However, it should be noted that the assumptions are with respect to each player and must hold individually for each player i . Now we are ready to state and describe the implications of our assumptions.

The following assumption ensures that the expected loss is strongly convex and smooth.

Assumption 1. Let S be the support of the vector \mathbf{v} , i.e. $S \doteq \{i \mid |v_i| > 0\}$. Then, there exists constants $C_{\min} > 0$ and $D_{\max} \leq |S|$ such that:

$$\lambda_{\min}(\mathbf{H}_{SS}^*) \geq C_{\min} \text{ and } \lambda_{\max}(\mathbb{E}_{\mathbf{x}}[\mathbf{z}_S \mathbf{z}_S^T]) \leq D_{\max},$$

for all i , where $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ denote the minimum and maximum eigenvalues respectively and $\mathbf{H}_{SS}^* = \{H_{i,j}^* \mid i, j \in S\}$.

To better understand the implications of the above assumptions first note that the distribution over joint actions in (3), after some algebraic manipulations, can be written as follows:

$$p(\mathbf{x}) = \left(\frac{q - |\mathcal{NE}^*|/2^n}{1 - |\mathcal{NE}^*|/2^n} \right) \frac{\mathbf{1}[\mathbf{x} \in \mathcal{NE}^*]}{|\mathcal{NE}^*|} + \left(\frac{1 - q}{1 - |\mathcal{NE}^*|/2^n} \right) \frac{1}{2^n}. \quad (8)$$

Thus, we have that the distribution over joint actions is a mixture of two uniform distributions, one over the number of Nash equilibria and the other being the Rademacher distribution over n variables. We denote the latter distribution by \mathfrak{R}^n , i.e. $\mathbf{x} \sim \mathfrak{R}^n \implies \mathbf{x} \in \{-1, +1\}^n \wedge p(\mathbf{x}) = 1/2^n$ for all \mathbf{x} . Therefore the expected Hessian matrix \mathbf{H}^* decomposes as a convex combination of two other Hessian matrices:

$$\mathbf{H}^* = \nu \mathbf{H}^{\mathcal{NE}^*} + (1 - \nu) \mathbf{H}^{\mathfrak{R}}, \quad (9)$$

where we have defined:

$$\begin{aligned} \mathbf{H}^{\mathcal{NE}^*} &\doteq \frac{1}{|\mathcal{NE}^*|} \sum_{\mathbf{x} \in \mathcal{NE}^*} \eta(\mathbf{v}^{*T} \mathbf{z}) \mathbf{z} \mathbf{z}^T, \\ \mathbf{H}^{\mathfrak{R}} &\doteq \mathbb{E}_{\mathbf{z} \sim \mathfrak{R}^n} \left[\eta(\mathbf{v}^{*T} \mathbf{z}) \mathbf{z} \mathbf{z}^T \right], \\ \nu &\doteq \left(\frac{q - |\mathcal{NE}^*|/2^n}{1 - |\mathcal{NE}^*|/2^n} \right). \end{aligned}$$

Now, by concavity of $\lambda_{\min}(\cdot)$ and the Jensen’s inequality we have that

$$\begin{aligned} C_{\min} &\geq \nu \lambda_{\min}(\mathbf{H}_{SS}^{\mathcal{NE}^*}) + (1 - \nu) \lambda_{\min}(\mathbf{H}_{SS}^{\mathfrak{R}}) \\ &\geq \nu \lambda_{\min}(\mathbf{H}_{SS}^{\mathcal{NE}^*}) + (1 - \nu) \eta(\|\mathbf{v}^*\|_1) \\ &\geq (1 - \nu) \eta(\|\mathbf{v}^*\|_1) > 0, \end{aligned}$$

where the last line follows from the fact that $\mathbf{H}_{SS}^{\mathcal{NE}^*}$ is positive semi-definite. In fact, $\lambda_{\min}(\mathbf{H}_{SS}^{\mathcal{NE}^*}) = 0$, if $|\mathcal{NE}^*| < n$. Thus, we have that our assumption of $q \in (|\mathcal{NE}^*|/2^n, 1)$ automatically implies that $\lambda_{\min}(\mathbf{H}_i^*) \geq C_{\min} > 0$. We can also verify that the maximum eigenvalue is bounded as follows:

$$\begin{aligned} D_{\max} &\leq \nu \lambda_{\max} \left(\frac{1}{|\mathcal{NE}^*|} \sum_{\mathbf{x} \in \mathcal{NE}^*} \mathbf{z}_S \mathbf{z}_S^T \right) + \\ &\quad (1 - \nu) \lambda_{\max}(\mathbb{E}_{\mathbf{z} \sim \mathfrak{R}^n} [\mathbf{z}_S \mathbf{z}_S^T]) \\ &\leq \nu |S| + (1 - \nu). \end{aligned}$$

The following assumption characterizes the minimum payoff in the Nash equilibria set.

Assumption 2. *The minimum payoff in the PSNE set, ρ_{\min} , is strictly positive, specifically:*

$$x_i(\mathbf{w}_{-i}^{*T} \mathbf{x}_{-i} - b_i) \geq \rho_{\min} > 5C_{\min}/D_{\max} \quad (\forall \mathbf{x} \in \mathcal{N}\mathcal{E}^*).$$

Note that as long as the minimum payoff is strictly positive, we can scale the parameters $(\mathbf{W}^*, \mathbf{b}^*)$ by the constant $5C_{\min}/D_{\max}$ to satisfy the condition: $\rho_{\min} > 5C_{\min}/D_{\max}$, without changing the PSNE set or the likelihood of the data. Indeed the assumption that the minimum payoff is strictly positive is unavoidable for exact recovery of the PSNE set in a noisy setting such as ours, because otherwise this is akin to exactly recovering the parameters \mathbf{v} for each player i . For example, if $\mathbf{x} \in \mathcal{N}\mathcal{E}^*$ is such that $\mathbf{v}^{*T} \mathbf{x} = 0$, then it can be shown that even if $\|\mathbf{v}^* - \widehat{\mathbf{v}}\|_{\infty} = \varepsilon$, for any ε arbitrarily close to 0, then $\widehat{\mathbf{v}}^T \mathbf{x} < 0$ and therefore $\mathcal{N}\mathcal{E}(\mathbf{W}^*, \mathbf{b}^*) \neq \mathcal{N}\mathcal{E}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$. Next, we present our main theoretical results for learning LIGs.

B. Theoretical Guarantees

Our main strategy for obtaining exact PSNE recovery guarantees is to first show, using results from random matrix theory, that given the assumptions on the eigenvalues of the population Hessian matrices, the assumptions hold in the finite sample case with high probability. Then, we exploit the convexity properties of the logistic loss function to show that the weight vectors learned using penalized logistic regression is “close” to the true weight vectors. By our assumption that the minimum payoff in the PSNE set is strictly greater than zero, we show that the weight vectors inferred from a finite sample of joint actions induce the same PSNE set as the true weight vectors.

1) *Minimum and Maximum Eigenvalues of Finite Sample Hessian and Scatter Matrices:* The following technical lemma shows that the assumptions on the eigenvalues of the Hessian matrices, hold with high probability in the finite sample case.

Lemma 1. *If $\lambda_{\min}(\mathbf{H}_{SS}^*) \geq C_{\min}$ and $\lambda_{\max}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) \leq D_{\max}$ then we have that*

$$\lambda_{\min}(\mathbf{H}_{SS}^m) \geq \frac{C_{\min}}{2} \quad \text{and} \quad \lambda_{\max} \left(\sum_{l=1}^m \mathbf{z}_S^{(l)} \mathbf{z}_S^{(l)T} \right) \leq 2D_{\max}$$

with probability at least

$$1 - |S| \exp \left(\frac{-mC_{\min}}{2|S|} \right) \quad \text{and} \quad 1 - |S| \exp \left(\frac{-m(1-\nu)}{4|S|} \right)$$

respectively.

Proof. Let

$$\mu_{\min} \doteq \lambda_{\min}(\mathbf{H}_{SS}^*) \quad \text{and} \quad \mu_{\max} \doteq \lambda_{\max}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]).$$

First note that for all $\mathbf{z} \in \{-1, +1\}^n$:

$$\begin{aligned} \lambda_{\max}(\eta(\mathbf{v}_S^{*T} \mathbf{z}_S) \mathbf{z}_S \mathbf{z}_S^T) &\leq \frac{|S|}{4} \doteq R \\ \lambda_{\max}(\mathbf{z}_S \mathbf{z}_S^T) &\leq |S| \doteq R'. \end{aligned}$$

Using the Matrix Chernoff bounds from Tropp [9](Theorem 1.1), we have that

$$\Pr \{ \lambda_{\min}(\mathbf{H}_{SS}^m) \leq (1-\delta)\mu_{\min} \} \leq |S| \left[\frac{e^{-\delta}}{(1-\delta)^{1-\delta}} \right]^{\frac{m\mu_{\min}}{R}}.$$

Setting $\delta = 1/2$ we get that

$$\begin{aligned} \Pr \{ \lambda_{\min}(\mathbf{H}_{SS}^m) \leq \mu_{\min}/2 \} &\leq |S| \left[\sqrt{\frac{2}{e}} \right]^{\frac{4m\mu_{\min}}{|S|}} \\ &\leq |S| \exp \left(\frac{-mC_{\min}}{2|S|} \right). \end{aligned}$$

Therefore, we have

$$\Pr \{ \lambda_{\min}(\mathbf{H}_{SS}^m) > C_{\min}/2 \} > 1 - |S| \exp \left(\frac{-mC_{\min}}{2|S|} \right).$$

Next, we have that

$$\begin{aligned} \mu_{\max} &= \lambda_{\max}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) \\ &\geq \lambda_{\min}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) \\ &\geq \nu \lambda_{\min} \left(\frac{1}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \mathbf{z}_S \mathbf{z}_S^T \right) + (1-\nu) \\ &\geq (1-\nu). \end{aligned}$$

Once again invoking Theorem 1.1 from [9] and setting $\delta = 1$ we have that

$$\begin{aligned} \Pr \{ \lambda_{\max} \geq (1+\delta)\mu_{\max} \} &\leq |S| \left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}} \right]^{(m\mu_{\max})/R'} \\ \Pr \{ \lambda_{\max} \geq 2\mu_{\max} \} &\leq |S| \left[\frac{e}{4} \right]^{(m\mu_{\max})/|S|} \\ &\leq |S| \exp \left(\frac{-m\mu_{\max}}{4|S|} \right) \\ &\leq |S| \exp \left(\frac{-m(1-\nu)}{4|S|} \right) \end{aligned}$$

Therefore, we have that

$$\Pr \{ \lambda_{\max} < 2D_{\max} \} > 1 - |S| \exp \left(\frac{-m(1-\nu)}{4|S|} \right).$$

□

2) *Recovering the Pure Strategy Nash Equilibria (PSNE) Set:* Before presenting our main result on the exact recovery of the PSNE set from noisy observations of joint actions, we first present a few technical lemmas that would be helpful in proving the main result. The following lemma bounds the gradient of the loss function (5) at the true vector \mathbf{v}^* , for all players.

Lemma 2. *With probability at least $1 - \delta$ for $\delta \in [0, 1]$, we have that*

$$\|\nabla \ell(\mathbf{v}^*, \mathcal{D})\|_{\infty} < \nu \kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}},$$

where $\kappa = 1/(1+\exp(\rho_{\min}))$ and $\rho_{\min} \geq 0$ is the minimum payoff in the PSNE set.

Proof. Let $\mathbf{u}^m \doteq \nabla \ell(\mathbf{v}^*, \mathcal{D})$ and u_j^m denote the j -th index of \mathbf{u}^m . We have that

$$\begin{aligned}
& (\forall j \in S^c \wedge j \neq n) \left| \mathbb{E} [u_j^m] \right| \\
&= \left| \frac{\nu}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} + \right. \\
&\quad \left. (1 - \nu) \mathbb{E}_{\mathbf{z} \sim \mathfrak{R}^n} \left[\frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} \right] \right| \\
&= \left| \frac{\nu}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} + \right. \\
&\quad \left. (1 - \nu) \left\{ \mathbb{E}_{\mathbf{z}_S \sim \mathfrak{R}^{|S|}} \left[\frac{1}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} \right] \mathbb{E}_{z_j \sim \mathfrak{R}} [z_j] \right\} \right| \\
&\leq \left| \frac{\nu\kappa}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} x_i x_j \right| = \nu\kappa
\end{aligned}$$

Similarly,

$$\begin{aligned}
& (\forall j \in S \wedge j \neq n) \left| \mathbb{E} [u_j^m] \right| \\
&= \left| \frac{\nu}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} + \right. \\
&\quad \left. (1 - \nu) \mathbb{E}_{\mathbf{z} \sim \mathfrak{R}^n} \left[\frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} \right] \right| \\
&\leq \left| \frac{\nu}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \frac{z_j}{1 + \exp((\mathbf{v}_S^*)^T \mathbf{z}_S)} + \right. \\
&\quad \left. (1 - \nu) \mathbb{E}_{z_j \sim \mathfrak{R}} [z_j] \right| \\
&\leq \nu\kappa.
\end{aligned}$$

Following the same procedure as above, it can be easily shown that the above bounds hold for the case $j = n$ as well. Also note that $|u_j^m| \leq 1$. Therefore, by using the Hoeffding's inequality [10] and a union bound argument we have that:

$$\begin{aligned}
& \Pr \left\{ \max_{j=1}^n |u_j^m - \mathbb{E} [u_j^m]| < t \right\} > 1 - 2ne^{-mt^2/2} \\
&\implies \Pr \{ \|\mathbf{u}^m - \mathbb{E} [\mathbf{u}^m]\|_\infty < t \} > 1 - 2ne^{-mt^2/2} \\
&\implies \Pr \{ \|\mathbf{u}^m\|_\infty - \|\mathbb{E} [\mathbf{u}^m]\|_\infty < t \} > 1 - 2ne^{-mt^2/2} \\
&\implies \Pr \{ \|\mathbf{u}^m\|_\infty < \nu\kappa + t \} > 1 - 2ne^{-mt^2/2}.
\end{aligned}$$

Setting $2n \exp(-mt^2/2) = \delta$, we prove our claim. \square

A consequence of Lemma 2 is that, even with an infinite number of samples, the gradient of the loss function at the true vector \mathbf{v}^* doesn't vanish. Therefore, we cannot hope to recover the parameters of the true game perfectly even with an infinite number of samples. In the following technical lemma we show that the optimal vector $\hat{\mathbf{v}}$ for the logistic regression problem is close to the true vector \mathbf{v}^* in the support set S of \mathbf{v}^* . Next, in Lemma 4, we bound the difference between the true vector \mathbf{v}^* and the optimal vector $\hat{\mathbf{v}}$ in the non-support set.

The lemmas together show that the optimal vector is close to the true vector.

Lemma 3. *If the regularization parameter λ satisfies the following condition:*

$$\lambda \leq \frac{5C_{\min}^2}{16|S|D_{\max}} - \nu\kappa - \sqrt{\frac{2}{m} \log \frac{2n}{\delta}},$$

then

$$\|\mathbf{v}_S^* - \hat{\mathbf{v}}_S\|_2 \leq \frac{5C_{\min}}{4\sqrt{|S|}D_{\max}},$$

with probability at least $1 - (\delta + |S| \exp((-mC_{\min})/2|S|) + |S| \exp((-m(1-\nu)/4|S|))$.

Proof. The proof of this lemma follows the general proof structure of Lemma 3 in [7]. First, we reparameterize the ℓ_1 -regularized loss function

$$f(\mathbf{v}_S) = \ell(\mathbf{v}_S) + \lambda \|\mathbf{v}_S\|_1$$

as the loss function \tilde{f} , which gives the loss at a point that is Δ_S distance away from the true parameter \mathbf{v}_S^* as follows:

$$\tilde{f}(\Delta_S) = \ell(\mathbf{v}_S^* + \Delta_S) - \ell(\mathbf{v}_S^*) + \lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1),$$

where $\Delta_S = \mathbf{v}_S - \mathbf{v}_S^*$. Also note that the loss function \tilde{f} is shifted such that the loss at the true parameter \mathbf{v}_S^* is 0, i.e. $\tilde{f}(\mathbf{0}) = 0$. Further, note that the function \tilde{f} is convex and is minimized at $\hat{\Delta}_S = \hat{\mathbf{v}}_S - \mathbf{v}_S^*$, since $\hat{\mathbf{v}}_S$ minimizes f . Therefore, clearly $\tilde{f}(\hat{\Delta}_S) \leq 0$. Thus, if we can show that the function \tilde{f} is strictly positive on the surface of a ball of radius b , then the point $\hat{\Delta}_S$ lies inside the ball i.e. $\|\hat{\mathbf{v}}_S - \mathbf{v}_S^*\|_2 \leq b$. Using the Taylor's theorem we expand the first term of \tilde{f} to get the following:

$$\begin{aligned}
\tilde{f}(\Delta_S) &= \nabla \ell(\mathbf{v}_S^*)^T \Delta_S + \Delta_S^T \nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S) \Delta_S \\
&\quad + \lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1), \tag{10}
\end{aligned}$$

for some $\theta \in [0, 1]$. Next, we lower bound each of the terms in (10). Using the Cauchy-Schwartz inequality, the first term in (10) is bounded as follows:

$$\begin{aligned}
\nabla \ell(\mathbf{v}_S^*)^T \Delta_S &\geq -\|\nabla \ell(\mathbf{v}_S^*)\|_\infty \|\Delta_S\|_1 \\
&\geq -\|\nabla \ell(\mathbf{v}_S^*)\|_\infty \sqrt{|S|} \|\Delta_S\|_2 \\
&\geq -b\sqrt{|S|} \left(\nu\kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} \right), \tag{11}
\end{aligned}$$

with probability at least $1 - \delta$ for $\delta \in [0, 1]$. It is also easy to upper bound the last term in equation 10, using the reverse triangle inequality as follows:

$$\lambda \|\mathbf{v}_S^* + \Delta_S\|_1 - \lambda \|\mathbf{v}_S^*\|_1 \leq \lambda \|\Delta_S\|_1.$$

Which then implies the following lower bound:

$$\begin{aligned}
\lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1) &\geq -\lambda \|\Delta_S\|_1 \\
&\geq -\lambda \sqrt{|S|} \|\Delta_S\|_2 \\
&= -\lambda \sqrt{|S|} b. \tag{12}
\end{aligned}$$

Now we turn our attention to computing a lower bound of the second term of (10), which is a bit more involved.

$$\begin{aligned}\Delta_S^T \nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S) \Delta_S &\geq \min_{\|\Delta_S\|_2=b} \Delta_S^T \nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S) \Delta_S \\ &= b^2 \lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S)).\end{aligned}$$

Now,

$$\begin{aligned}\lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S)) &\geq \min_{\theta \in [0,1]} \lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S)) \\ &= \min_{\theta \in [0,1]} \lambda_{\min} \left(\frac{1}{m} \sum_{l=1}^m \eta((\mathbf{v}_S^* + \theta \Delta_S)^T \mathbf{z}_S^{(l)}) \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T \right)\end{aligned}$$

Again, using the Taylor's theorem to expand the function η we get

$$\begin{aligned}\eta((\mathbf{v}_S^* + \theta \Delta_S)^T \mathbf{z}_S^{(l)}) &= \eta((\mathbf{v}_S^*)^T \mathbf{z}_S^{(l)}) + \eta'((\mathbf{v}_S^* + \bar{\theta} \Delta_S)^T \mathbf{z}_S^{(l)}) (\theta \Delta_S)^T \mathbf{z}_S^{(l)}\end{aligned}$$

, where $\bar{\theta} \in [0, \theta]$. Continuing from above and from Lemma 1 we have, with probability at least $1 - |S| \exp((-mC_{\min})/2|S|)$:

$$\begin{aligned}\lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S)) &\geq \min_{\theta \in [0,1]} \lambda_{\min} \left(\frac{1}{m} \sum_{l=1}^m \eta((\mathbf{v}_S^*)^T \mathbf{z}_S^{(l)}) \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T \right. \\ &\quad \left. + \frac{1}{m} \sum_{l=1}^m \eta'((\mathbf{v}_S^* + \bar{\theta} \Delta_S)^T \mathbf{z}_S^{(l)}) ((\theta \Delta_S)^T \mathbf{z}_S^{(l)}) \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T \right) \\ &\geq \lambda_{\min}(\mathbf{H}_{SS}^m) - \max_{\theta \in [0,1]} \|\mathbf{A}(\theta)\|_2 \\ &\geq \frac{C_{\min}}{2} - \max_{\theta \in [0,1]} \|\mathbf{A}(\theta)\|_2,\end{aligned}$$

where we have defined

$$\mathbf{A}(\theta) \doteq \frac{1}{m} \sum_{l=1}^m \eta'((\mathbf{v}_S^* + \theta \Delta_S)^T \mathbf{z}_S^{(l)}) (\theta \Delta_S)^T \mathbf{z}_S^{(l)} \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T.$$

Next, the spectral norm of $\mathbf{A}(\theta)$ can be bounded as follows:

$$\begin{aligned}\|\mathbf{A}(\theta)\|_2 &\leq \max_{\|\mathbf{y}\|_2=1} \left\{ \frac{1}{m} \sum_{l=1}^m \left| \eta'((\mathbf{v}_S^* + \theta \Delta_S)^T \mathbf{z}_S^{(l)}) \right| \left| ((\theta \Delta_S)^T \mathbf{z}_S^{(l)}) \right| \right. \\ &\quad \left. \times \mathbf{y}^T (\mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T) \mathbf{y} \right\} \\ &< \max_{\|\mathbf{y}\|_2=1} \left\{ \frac{1}{10m} \sum_{l=1}^m \|(\theta \Delta_S)\|_1 \|\mathbf{z}_S^{(l)}\|_{\infty} \mathbf{y}^T (\mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T) \mathbf{y} \right\} \\ &\leq \theta \max_{\|\mathbf{y}\|_2=1} \left\{ \frac{1}{10m} \sum_{l=1}^m \sqrt{|S|} \|\Delta_S\|_2 \mathbf{y}^T (\mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T) \mathbf{y} \right\} \\ &= \theta b \sqrt{|S|} \left\| \frac{1}{10m} \sum_{l=1}^m \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T \right\|_2 \\ &\leq \frac{(b\sqrt{|S|}D_{\max})}{5} \leq \frac{C_{\min}}{4},\end{aligned}$$

where in the second line we used the fact that $\eta'(\cdot) < 1/10$ and in the last line we assumed that $(b\sqrt{|S|}D_{\max})/5 \leq C_{\min}/4$ — an assumption that we verify momentarily. Having upper bounded the spectral norm of $\mathbf{A}(\theta)$, we have

$$\lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S)) \geq \frac{C_{\min}}{4}. \quad (13)$$

Plugging back the bounds given by (11), (12) and (13) in (10) and equating to zero we get

$$\begin{aligned}-b\sqrt{|S|} \left(\nu\kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} \right) + \frac{b^2 C_{\min}}{4} - \lambda \sqrt{|S|} b &= 0 \\ \implies b = \frac{4\sqrt{|S|}}{C_{\min}} \left(\lambda + \nu\kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} \right).\end{aligned}$$

Finally, coming back to our prior assumption we have

$$b = \frac{4\sqrt{|S|}}{C_{\min}} \left(\lambda + \nu\kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} \right) \leq \frac{5C_{\min}}{4\sqrt{|S|}D_{\max}}.$$

The above assumption holds if the regularization parameter λ is bounded as follows:

$$\lambda \leq \frac{5C_{\min}^2}{16|S|D_{\max}} - \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} - \nu\kappa. \quad \square$$

Lemma 4. *If the regularization parameter λ satisfies the following condition:*

$$\lambda \geq \nu\kappa + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}},$$

then we have that

$$\|\hat{\mathbf{v}} - \mathbf{v}^*\|_1 \leq \frac{5C_{\min}}{D_{\max}}$$

with probability at least $1 - (\delta + |S| \exp((-mC_{\min})/2|S|) + |S| \exp((-m(1-\nu)/4|S|))$.

Proof. Define $\Delta \doteq \hat{\mathbf{v}} - \mathbf{v}^*$. Also for any vector \mathbf{y} let the notation $\mathbf{y}_{\bar{S}}$ denote the vector \mathbf{y} with the entries not in the support, S , set to zero, i.e.

$$[\mathbf{y}_{\bar{S}}]_i = \begin{cases} y_i & \text{if } i \in S, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, let the notation \mathbf{y}_{S^c} denote the vector \mathbf{y} with the entries not in S^c set to zero, where S^c is the complement of S . Having introduced our notation and since, S is the support of the true vector \mathbf{v}^* , we have by definition that $\mathbf{v}^* = \mathbf{v}_{\bar{S}}^*$. We then have, using the reverse triangle inequality,

$$\begin{aligned}\|\hat{\mathbf{v}}\|_1 &= \|\mathbf{v}^* + \Delta\|_1 = \|\mathbf{v}_{\bar{S}}^* + \Delta_{\bar{S}} + \Delta_{S^c}\|_1 \\ &= \|\mathbf{v}_{\bar{S}}^* - (-\Delta_{\bar{S}})\|_1 + \|\Delta_{S^c}\|_1 \\ &\geq \|\mathbf{v}^*\|_1 - \|\Delta_{\bar{S}}\|_1 + \|\Delta_{S^c}\|_1.\end{aligned} \quad (14)$$

Also, from the optimality of $\hat{\mathbf{v}}$ for the ℓ_1 -regularized problem we have that

$$\begin{aligned}\ell(\mathbf{v}^*) + \lambda \|\mathbf{v}^*\|_1 &\geq \ell(\hat{\mathbf{v}}) + \lambda \|\hat{\mathbf{v}}\|_1 \\ \implies \lambda(\|\mathbf{v}^*\|_1 - \|\hat{\mathbf{v}}\|_1) &\geq \ell(\hat{\mathbf{v}}) - \ell(\mathbf{v}^*).\end{aligned} \quad (15)$$

Next, from convexity of $\ell(\cdot)$ and using the Cauchy-Schwartz inequality we have that

$$\begin{aligned}\ell(\widehat{\mathbf{v}}) - \ell(\mathbf{v}^*) &\geq \nabla \ell(\mathbf{v}^*)^T (\widehat{\mathbf{v}} - \mathbf{v}^*) \\ &\geq -\|\nabla \ell(\mathbf{v}^*)\|_\infty \|\Delta\|_1 \\ &\geq -\frac{\lambda}{2} \|\Delta\|_1,\end{aligned}\quad (16)$$

where in the last line we used the fact that $\lambda \geq \|\nabla \ell(\mathbf{v}^*)\|_\infty$. Thus, we have from (14), (15) and (16) that

$$\begin{aligned}\frac{1}{2} \|\Delta\|_1 &\geq \|\widehat{\mathbf{v}}\|_1 - \|\mathbf{v}^*\|_1 \\ \implies \frac{1}{2} \|\Delta\|_1 &\geq \|\Delta_{\overline{S^c}}\|_1 - \|\Delta_{\overline{S}}\|_1 \\ \implies \frac{1}{2} \|\Delta_{\overline{S^c}}\|_1 + \frac{1}{2} \|\Delta_{\overline{S}}\|_1 &\geq \|\Delta_{\overline{S^c}}\|_1 - \|\Delta_{\overline{S}}\|_1 \\ \implies 3 \|\Delta_{\overline{S}}\|_1 &\geq \|\Delta_{\overline{S^c}}\|_1.\end{aligned}\quad (17)$$

Finally, from (17) and Lemma 3 we have that

$$\begin{aligned}\|\Delta\|_1 &= \|\Delta_{\overline{S}}\|_1 + \|\Delta_{\overline{S^c}}\|_1 \\ &\leq 4 \|\Delta_{\overline{S}}\|_1 \leq 4\sqrt{|S|} \|\Delta_{\overline{S}}\|_2 \\ &\leq \frac{5C_{\min}}{D_{\max}}.\end{aligned}$$

□

Now we are ready to present our main result on recovering the true PSNE set.

Theorem 1. *If for all i , $|S_i| \leq k$, the minimum payoff $\rho_{\min} \geq 5C_{\min}/D_{\max}$, and the regularization parameter and the number of samples satisfy the following conditions:*

$$\begin{aligned}\nu\kappa + \sqrt{\frac{2}{m} \log \frac{6n^2}{\delta}} \leq \lambda \leq 2K + \nu\kappa - \sqrt{\frac{2}{m} \log \frac{6n^2}{\delta}} \quad (18) \\ m \geq \max \left\{ \frac{2}{K^2} \log \left(\frac{6n^2}{\delta} \right), \frac{2k}{C_{\min}} \log \left(\frac{3kn}{\delta} \right), \right. \\ \left. \frac{4k}{1-\nu} \log \left(\frac{3kn}{\delta} \right) \right\}, \quad (19)\end{aligned}$$

where $K \doteq 5C_{\min}^2/32kD_{\max} - \nu\kappa$, then with probability at least $1 - \delta$, for $\delta \in [0, 1]$, we recover the true PSNE set, i.e. $\mathcal{NE}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) = \mathcal{NE}(\mathbf{W}^*, \mathbf{b}^*)$.

Proof. From Cauchy-Schwartz inequality and Lemma 4 we have

$$|(\widehat{\mathbf{v}}_i - \mathbf{v}_i^*)^T \mathbf{z}_i| \leq \|\widehat{\mathbf{v}}_i - \mathbf{v}_i^*\|_1 \|\mathbf{z}_i\|_\infty \leq \frac{5C_{\min}}{D_{\max}}.$$

Therefore, we have that

$$(\mathbf{v}_i^*)^T \mathbf{z}_i - \frac{5C_{\min}}{D_{\max}} \leq \widehat{\mathbf{v}}_i^T \mathbf{z}_i \leq (\mathbf{v}_i^*)^T \mathbf{z}_i + \frac{5C_{\min}}{D_{\max}}.$$

Now, if $\forall \mathbf{x} \in \mathcal{NE}^*$, $(\mathbf{v}_i^*)^T \mathbf{z}_i \geq 5C_{\min}/D_{\max}$, then $\widehat{\mathbf{v}}_i^T \mathbf{z}_i \geq 0$. Using an union bound argument over all players i , we can show that the above holds with probability at least

$$1 - n(\delta + k \exp((-mC_{\min})/2k) + k \exp((-m(1-\nu))/4k)) \quad (20)$$

for all players. Therefore, we have that $\mathcal{NE}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}}) = \mathcal{NE}^*$ with high probability. Finally, setting $\delta = \delta'/3n$, for some $\delta' \in [0, 1]$, and ensuring that the last two terms in (20) are at most $\delta'/3n$ each, we prove our claim. □

To better understand the implications of the theorem above, we discuss some possible operating regimes for learning sparse linear influence games in the following paragraphs.

Remark 1 (Sample complexity for fixed q). *In the theorem above, if q is constant, which in turn makes ν constant, then $K = \Omega(1/k^2)$, and the sample complexity of learning sparse linear games grows as $\mathcal{O}(k^4 \log n)$. However, if q is small enough such that $\nu \leq 1/k$, then the constant D_{\max} is no longer a function of k and hence $K = \Omega(1/k)$. Therefore, the sample complexity scales as $\mathcal{O}(k^2 \log n)$ for exact PSNE recovery.*

The sample complexity of $\mathcal{O}(\text{poly}(k) \log n)$ for exact recovery of the PSNE set can be compared with the sample complexity of $\mathcal{O}(kn^3 \log^2 n)$ for the maximum likelihood estimate (MLE) of the PSNE set as obtained by Honorio [11]. Note that while the MLE procedure is consistent, i.e. MLE of the PSNE set is equal to the true PSNE set with probability converging to 1 as the number of samples tend to infinity, it is NP-hard³. In contrast, the logistic regression method is computationally efficient. Further, while the sample complexity of our method seems to be better than the empirical log-likelihood minimizer as given by Theorem 3 in [11], in this paper we restrict ourselves to LIGs with strictly positive payoff in the PSNE set — such games are invariably easier to learn than general LIGs considered by [11].

Honorio [11] also obtained lower bounds on the number of samples required by any conceivable method, for exact PSNE recovery, by scaling the parameter q with the number of players. Therefore, in order to compare our results with the information-theoretic limits of learning LIGs, we consider the regime where the parameter q scales with the number of players n .

Remark 2 (Sample complexity for q varying with number of players). *If we consider the regime where the signal level scales as $q = (|\mathcal{NE}^*|+1)/2^n$. Then, $D_{\max} = \mathcal{O}(k/(2^n - |\mathcal{NE}^*|))$, and as a result $K = \Omega(2^n/k^2)$. Therefore, the sample complexity, which is dominated by the second and third terms in (19), is given as $\mathcal{O}(k \log(kn))$. In general if $q = \Theta(\exp(-n))$, then the sample complexity for recovering the PSNE set exactly is $\mathcal{O}(k \log(kn))$.*

Once again we observe that even in the regime of q scaling exponentially with n , the sample complexity of $\mathcal{O}(k \log(kn))$ is better than the information theoretic limit of $\mathcal{O}(kn \log^2 n)$, by a factor of $\mathcal{O}(n \log n)$. This can be attributed to the fact that we consider restricted ensembles of games with strictly positive payoffs in the PSNE set, as opposed to general LIGs.

Further, from the aforementioned remarks we see that as the signal level q decreases, the sufficient number of samples

³Irfan and Ortiz [4] showed that counting the number of Nash equilibria is #P-complete. Therefore, computing the log-likelihood is NP-hard.




Fig. 1. The probability of exact recovery of the PSNE set computed across 40 randomly sampled LIGs as the number of samples is scaled as $\lfloor (C)(10^c)(k^2 \log(6n^2/\delta)) \rfloor$, where c is the control parameter and the constant C is 10000 for the $k = 1$ case and 1000 for the remaining two cases.

needed to recover the PSNE set reduces; and in the limiting case of q decreasing exponentially with the number of players, the sample complexity scales as $\mathcal{O}(k \log(kn))$. This seems counter-intuitive — with increased signal level, a learning problem should become easier and not harder. To understand this seemingly counter-intuitive behavior, first observe that the constant D_{\max}/C_{\min} can be thought of as the “condition number” of the loss function given by 5. Therefore, the sample complexity as given by Theorem 1 can be written as $\mathcal{O}(k^2(D_{\max}/C_{\min})^2 \log n)$. From (9), we see that as the signal level increases, the Hessian of the loss becomes more ill-conditioned, since the data set now comprises of many repetitions of the few joint-actions that are in the PSNE set; thereby increasing the dependency (D_{\max}) between actions of players in the sample data set.

V. EXPERIMENTS

In order to verify that our results and assumptions indeed hold in practice, we performed various simulation experiments. We generated random LIGs for n players and exactly k neighbors by first creating a matrix \mathbf{W} of all zeros and then setting k off-diagonal entries of each row, chosen uniformly at random, to -1 . We set the bias for all players to 0. We found that any odd value of k produces games with strictly positive payoff in the PSNE set. Therefore, for each value of k in $\{1, 3, 5\}$, and n in $\{10, 12, 15, 20\}$, we generated 40 random LIGs. The parameters q and δ were set to the constant value of 0.01 and the regularization parameter λ was set according to Theorem 1 as some constant multiple of $\sqrt{(2/m) \log(2n/\delta)}$. Figure 1 shows the probability of successful recovery of the PSNE, for various combinations of (n, k) , where the probability was computed as the fraction of the 40 randomly sampled LIGs for which the learned PSNE set matched the true PSNE set exactly. For each experiment, the number of samples was computed as: $\lfloor (C)(10^c)(k^2 \log(6n^2/\delta)) \rfloor$, where c is the control parameter and the constant C is 10000 for $k = 1$ and 1000 for $k = 3$ and 5. Thus, from Figure 1 we see that, the sample complexity of $\mathcal{O}(k^2 \log n)$ as given by Theorem 1 indeed holds in practice — i.e. there exists constants c and c' such that if the number of samples is less than $ck^2 \log n$, we fail to recover the PSNE set exactly with high probability, while if the number of samples is greater than $c'k^2 \log n$ then we are able to recover the PSNE set exactly, with high probability. Further, the scaling remains

consistent as the number of players n is changed from 10 to 20.

VI. CONCLUSION

In this paper, we presented a computationally efficient and statistically consistent method, based on ℓ_1 -regularized logistic regression, for learning linear influence games — a subclass of parametric graphical games with linear payoffs. Under some mild conditions on the true game, we showed that as long as the number of samples scales as $\mathcal{O}(\text{poly}(k) \log n)$, where n is the number of players and k is the maximum number of neighbors of any player; then we can recover the pure-strategy Nash equilibria set of the true game in polynomial time and with probability converging to 1 as the number of samples tend to infinity. An interesting direction for future work would be to consider structured actions — for instance permutations, directed spanning trees, directed acyclic graphs among others — thereby extending the formalism of linear influence games to the structured prediction setting. A more technical extension would be to consider a local noise model where the observations are drawn from the PSNE set but with each action independently corrupted by some noise. Other ideas that might be worth pursuing are: considering mixed strategies, correlated equilibria and epsilon Nash equilibria, and incorporating latent or unobserved actions and variables in the model.

REFERENCES

- [1] B. Blum, C. R. Shelton, and D. Koller, “A continuation method for nash equilibria in structured games,” *Journal of Artificial Intelligence Research*, vol. 25, pp. 457–502, 2006.
- [2] L. E. Ortiz and M. Kearns, “Nash propagation for loopy graphical games,” in *Advances in Neural Information Processing Systems*, 2002, pp. 793–800.
- [3] D. Vickrey and D. Koller, “Multi-agent algorithms for solving graphical games,” in *AAAI/IAAI*, 2002, pp. 345–351.
- [4] M. T. Irfan and L. E. Ortiz, “On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach,” *Artificial Intelligence*, vol. 215, pp. 79–119, 2014.
- [5] O. Ben-Zwi and A. Ronen, “Local and global price of anarchy of graphical games,” *Theoretical Computer Science*, vol. 412, no. 12, pp. 1196–1207, 2011.
- [6] J. Honorio and L. Ortiz, “Learning the structure and parameters of large-population graphical games from behavioral data,” *Journal of Machine Learning Research*, vol. 16, pp. 1157–1210, 2015.
- [7] P. Ravikumar, M. Wainwright, and J. Lafferty, “High-dimensional Ising model selection using L1-regularized logistic regression,” *The Annals of Statistics*, vol. 38, no. 3, pp. 1287–1319, 2010. [Online]. Available: <http://arxiv.org/abs/1010.0311>
- [8] M. Kearns, M. L. Littman, and S. Singh, “Graphical models for game theory,” in *Proceedings of the Seventeenth conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., 2001, pp. 253–260.
- [9] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Foundations of computational mathematics*, vol. 12, no. 4, pp. 389–434, 2012.
- [10] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, vol. 58, no. 301, pp. 13–30, 1963.
- [11] J. Honorio, “On the Sample Complexity of Learning Sparse Graphical Games,” 2016. [Online]. Available: <http://arxiv.org/abs/1601.07243>