
Learning Graphical Games from Behavioral Data: Sufficient and Necessary Conditions

Asish Ghoshal
aghoshal@purdue.edu

Department of Computer Science, Purdue University, West Lafayette, IN, U.S.A. - 47907.

Jean Honorio
jhonorio@purdue.edu

Abstract

In this paper we obtain sufficient and necessary conditions on the number of samples required for exact recovery of the pure-strategy Nash equilibria (PSNE) set of a graphical game from noisy observations of joint actions. 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 show that one can efficiently recover the PSNE set of a linear influence game with $\mathcal{O}(k^2 \log n)$ samples, under very general observation models. On the other hand, we show that $\Omega(k \log n)$ samples are necessary for any procedure to recover the PSNE set from observations of joint actions.

1 Introduction and Related Work

Non-cooperative game theory is widely considered as an appropriate mathematical framework for studying *strategic* behavior in multi-agent scenarios. In Non-cooperative game theory, 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. On the other hand, Garg and Jaakkola [7] provide a discriminative approach to learn a class of graphical games called potential games. 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 obtain necessary and sufficient conditions for recovering the PSNE set of a graphical game in polynomial time. We also generalize the observation model from Ghoshal and Honorio [8], to arbitrary distributions that satisfy certain mild conditions. Our polynomial time method for recovering the PSNE set, which was proposed by Honorio and Ortiz [6], is based on 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 best achievable likelihood. In this paper we obtain the stronger guarantee of re-

covering the true PSNE set exactly.

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. [9] 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. [9] 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 conditions required by Ravikumar et. al. [9], such as mutual incoherence.

2 Preliminaries

In this section we provide some background information on graphical games introduced by Kearns et. al. [10].

2.1 Graphical Games

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

An important solution concept in the theory of non-cooperative games is that of *Nash equilibrium*. For a non-cooperative game, a joint action $\mathbf{x}^* \in \mathcal{X}$ is a pure-strategy Nash equilibrium (PSNE) if, for each player i , $x_i^* \in \operatorname{argmax}_{x_i \in \mathcal{X}_i} u_i(x_i, \mathbf{x}_{-i}^*)$, 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:

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

Graphical games, introduced by Kearns et al. [10], are game-theoretic analogues of graphical models. 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.

2.2 Linear Influence Games

Irfan and Ortiz [4] and Honorio and Ortiz [6], introduced a specific form of graphical games, called *Linear Influential Games*, characterized by binary actions, or pure strategies, and linear payoff functions. We assume, without loss of generality, that the joint action space $\mathcal{X} = \{-1, +1\}^n$. 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 is a linear function of the actions of the remaining players: $u_i(x_i, \mathbf{x}_{-i}) = x_i(\mathbf{w}_{-i}^T \mathbf{x}_{-i} - b_i)$, and the PSNE set is defined as follows:

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

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 $\operatorname{diag}(\mathbf{W}) = 0$. 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 a sparse game over n players where the in-degree of any vertex is at most k .

3 Problem Formulation

Now we turn our attention to the problem of learning graphical games from observations of joint actions only. Let $\mathcal{NE}^* \stackrel{\text{def}}{=} \mathcal{NE}(\mathcal{G}^*(n, k))$. 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 some distribution \mathcal{P} . We will use two specific distributions \mathcal{P}_g and \mathcal{P}_l , which we refer to as the *global* and *local* noise model, to provide further intuition behind our results. In the global noise model, we assume that a joint action is observed from the PSNE set with probability $q_g \in (|\mathcal{NE}^*|/2^n, 1)$, i.e.

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

In the above distribution, q_g can be thought of as the “signal” level in the data set, while $1 - q_g$ can be thought of as the “noise” level in the data set. In the local noise model we assume that the joint actions are drawn from the PSNE set with the action of

each player corrupted independently by some Bernoulli noise. Then in the local noise model the distribution over joint actions is given as follows:

$$\mathcal{P}_l(\mathbf{x}) = \frac{1}{|\mathcal{N}\mathcal{E}^*|} \sum_{\mathbf{y} \in \mathcal{N}\mathcal{E}^*} \prod_{i=1}^n q_i^{\mathbf{1}[x_i=y_i]} (1-q_i)^{\mathbf{1}[x_i \neq y_i]}, \quad (4)$$

where $q_i > 0.5$. While these noise models were introduced in [6], we obtain our results with respect to very general observation models, satisfying only some mild conditions. 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 game graph?’’ 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 global noise model (3) — an issue known as non-identifiability in the statistics literature. It is also easy to see that the same holds true for the local noise model. 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 :

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

Therefore, our goal in this paper is efficient and consistent recovery of the pure-strategy Nash equilibria set (PSNE) from observations of joint actions only; i.e., given a data set \mathcal{D} , drawn from some game $\mathcal{G}^*(n, k)$ according to the distribution \mathcal{P} , we infer a game $\widehat{\mathcal{G}}(n, k)$ from \mathcal{D} such that $\widehat{\mathcal{G}}(n, k) \equiv \mathcal{G}^*(n, k)$.

4 Method and Results

In order to efficiently learn games, we make a few assumptions on the probability distribution from which samples are drawn and also on the underlying game.

4.1 Assumptions

The following assumption ensures that the distribution \mathcal{P} assigns non-zero mass to all joint actions in \mathcal{X} and that the signal level in the data set is more than the noise level.

Assumption 1. *There exists constants \tilde{p}_{\min} , \tilde{p}_{\max} and p_{\max} such that the data distribution \mathcal{P} satisfies the following:*

$$0 < \frac{\tilde{p}_{\min}}{2^n - |\mathcal{N}\mathcal{E}^*|} \leq \mathcal{P}(\mathbf{x}) \leq \frac{\tilde{p}_{\max}}{2^n - |\mathcal{N}\mathcal{E}^*|}, \forall \mathbf{x} \in \mathcal{X} \setminus \mathcal{N}\mathcal{E}^*, \\ \frac{\tilde{p}_{\max}}{2^n - |\mathcal{N}\mathcal{E}^*|} < \mathcal{P}(\mathbf{x}) \leq p_{\max} \leq 1, \forall \mathbf{x} \in \mathcal{N}\mathcal{E}^*.$$

To get some intuition for the above assumption, consider the global noise model. In this case we have that $\tilde{p}_{\min} = \tilde{p}_{\max} = (1 - q_g)$, $p_{\max} = q_g/|\mathcal{N}\mathcal{E}^*|$, and $\forall \mathbf{x} \in \mathcal{N}\mathcal{E}^*$, $\mathcal{P}(\mathbf{x}) = p_{\max}$. For the local noise model, consider, for simplicity, the case when there are only two joint actions in the PSNE set: $\mathcal{N}\mathcal{E}^* = \{\mathbf{x}^1, \mathbf{x}^2\}$, such that $x_1^1 = +1, x_1^2 = -1$ and $x_i^1 = x_i^2 = +1$ for all $i \neq 1$. Then, $\tilde{p}_{\min} = 0.5 \times (1 - q_2) \times \dots \times (1 - q_n) \times (2^n - 2)$, $\tilde{p}_{\max} = 0.5 \times (1 - q_j) (\prod_{i \notin \{j, 1\}} q_i) \times (2^n - 2)$, where $q_j = \min\{q_2, \dots, q_n\}$, and $p_{\max} = 0.5 \times q_2 \times \dots \times q_n$.

Our next assumption concerns with the minimum payoff in the PSNE set.

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

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

where $C_{\min} > 0$ and D_{\max} are the minimum and maximum eigenvalue of the expected Hessian and scatter matrices respectively.

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. 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 $\mathcal{N}\mathcal{E}(\mathbf{W}^*, \mathbf{b}^*) \neq \mathcal{N}\mathcal{E}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$.

4.2 Method

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; 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:

$$\widehat{\mathbf{v}}_i = \underset{\mathbf{v}_i}{\operatorname{argmin}} \ell(\mathbf{v}_i, \mathcal{D}) + \lambda \|\mathbf{v}_i\|_1 \quad (5)$$

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

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 take a moment to introduce the expressions of the

gradient and the Hessian of the loss function (6), which will be useful later. The gradient and Hessian of the loss function for any vector \mathbf{v} and the data set \mathcal{D} is given as follows:

$$\nabla \ell(\mathbf{v}, \mathcal{D}) = \frac{1}{m} \sum_{l=1}^m \left\{ \frac{-\mathbf{z}^{(l)}}{1 + \exp(\mathbf{v}^T \mathbf{z}^{(l)})} \right\} \quad (7)$$

$$\nabla^2 \ell(\mathbf{v}, \mathcal{D}) = \frac{1}{m} \sum_{l=1}^m \eta(\mathbf{v}^T \mathbf{z}^{(l)}) \mathbf{z}^{(l)} \mathbf{z}^{(l)T}, \quad (8)$$

where $\eta(x) = 1/(e^{x/2} + e^{-x/2})^2$. Finally, \mathbf{H}_i^n 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. $\mathbf{H}_i^* \stackrel{\text{def}}{=} \mathbb{E}_{\mathcal{D}} [\mathbf{H}_i^m] = \mathbb{E}_{\mathcal{D}} [\nabla^2 \ell(\mathbf{v}_i^*, \mathcal{D})]$. In subsequent sections we drop the notational dependence of \mathbf{H}_i^* and \mathbf{z}_i on i to simplify notation.

We show that, under the aforementioned assumptions on the true game $\mathcal{G}^*(n, k) = (\mathbf{W}^*, \mathbf{b}^*)$, the parameters $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{b}}$ obtained using (6) induce the same PSNE set as the true game, i.e., $\mathcal{NE}(\mathbf{W}^*, \mathbf{b}^*) = \mathcal{NE}(\widehat{\mathbf{W}}, \widehat{\mathbf{b}})$.

4.3 Sufficient Conditions

In this section, we derive sufficient conditions on the number of samples for efficiently recovering the PSNE set of graphical games with linear payoffs. To start with, we make the following observation regarding the number of Nash equilibria of the game satisfying Assumption 2. The proof of the following proposition, as well as other missing proofs can be found in Appendix A.

Proposition 1. *The number of Nash equilibria of a non-trivial game ($|\mathcal{NE}^*| \in [1, 2^n - 1]$) satisfying Assumption 2 is at most 2^{n-1} .*

We will denote the fraction of joint actions that are in the PSNE set by $f_{\mathcal{NE}^*} \stackrel{\text{def}}{=} |\mathcal{NE}^*|/2^{n-1}$. By proposition 1, $f_{\mathcal{NE}^*} \in (0, 1]$. Then, our main strategy for obtaining sufficient conditions for exact PSNE recovery guarantees is to first show that under any data distribution \mathcal{P} that satisfies Assumption 1, the expected loss is smooth and strongly convex, i.e., the population Hessian matrix is positive definite and the population scatter matrix has eigenvalues bounded by a constant. Then using tools from random matrix theory, we show that the sample Hessian and scatter matrices are “well behaved”, i.e., are positive definite and have bounded eigenvalues respectively, 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 are “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.

The following lemma shows that the expected Hessian matrices for each player is positive definite and the maximum eigenvalues of the expected scatter matrices are bounded from above by a constant.

Lemma 1. *Let S be the support of the vector \mathbf{v} , i.e., $S \stackrel{\text{def}}{=} \{i \mid |v_i| > 0\}$. There exists constant $C_{\min} \geq \frac{\eta(\|\mathbf{v}^*\|_1) 2^n \tilde{p}_{\min}}{2^n - |\mathcal{NE}^*|} > 0$ and $D_{\max} \leq 2^n p_{\max}$, such that we have $\lambda_{\min}(\mathbf{H}_{SS}^*) = C_{\min}$ and $\lambda_{\max}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) = D_{\max}$.*

Proof.

$$\begin{aligned} \lambda_{\min}(\mathbf{H}_{SS}^*) &= \lambda_{\min} \left(\mathbb{E}_{\mathbf{x}} \left[\eta(\mathbf{v}^{*T} \mathbf{z}) \mathbf{z}_S \mathbf{z}_S^T \right] \right) \\ &= \eta(\|\mathbf{v}^*\|_1) \lambda_{\min}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]). \end{aligned}$$

Let $\mathbf{Z} \stackrel{\text{def}}{=} \{\mathbf{z}_S \mid \mathbf{x} \in \mathcal{X}\}$ and $\mathbf{P} \stackrel{\text{def}}{=} \text{Diag}((\mathcal{P}(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}})$, where \mathbf{z}_S denotes the feature vector for the i -th player constrained to the support set S for some i . Note that $\mathbf{Z} \in \{-1, 1\}^{2^n \times |S|}$; $\mathbf{P} \in \mathbb{R}^{2^n \times 2^n}$ and is positive definite by our assumption that the minimum probability $\frac{\tilde{p}_{\min}}{2^n - |\mathcal{NE}^*|} > 0$. Further note that the columns of \mathbf{Z} are orthogonal and $\mathbf{Z}^T \mathbf{Z} = 2^n \mathbf{I}_{|S|}$, where $\mathbf{I}_{|S|}$ is the $|S| \times |S|$ identity matrix. Then we have that

$$\begin{aligned} \lambda_{\min}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) &= \min_{\{\mathbf{y} \in \mathbb{R}^{|S|} \mid \|\mathbf{y}\|_2 = 1\}} \mathbf{y}^T \mathbf{Z}^T \mathbf{P} \mathbf{Z} \mathbf{y} \\ &= \min_{\{\mathbf{y}' \in \mathbb{R}^{2^n} \mid \mathbf{y}' = \mathbf{z}_S / \sqrt{2^n} \wedge \mathbf{y}' \in \mathbb{R}^{|S|} \wedge \|\mathbf{y}'\|_2 = 1\}} 2^n (\mathbf{y}')^T \mathbf{P} \mathbf{y}' \\ &\geq \min_{\{\mathbf{y}' \in \mathbb{R}^{2^n} \mid \|\mathbf{y}'\|_2 = 1\}} 2^n (\mathbf{y}')^T \mathbf{P} \mathbf{y}' \\ &= 2^n \lambda_{\min}(\mathbf{P}) = \frac{2^n \tilde{p}_{\min}}{2^n - |\mathcal{NE}^*|} \end{aligned}$$

Therefore, the minimum eigenvalue of \mathbf{H}_{SS}^* is lower bounded as follows:

$$\lambda_{\min}(\mathbf{H}_{SS}^*) = C_{\min} \geq \frac{\eta(\|\mathbf{v}^*\|_1) 2^n \tilde{p}_{\min}}{2^n - |\mathcal{NE}^*|} > 0.$$

Similarly, the maximum eigenvalue of $\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]$ can be bounded as $\lambda_{\max}(\mathbb{E}_{\mathbf{x}} [\mathbf{z}_S \mathbf{z}_S^T]) = \lambda_{\max}(\mathbf{Z}^T \mathbf{P} \mathbf{Z}) \leq 2^n p_{\max}$. \square

4.3.1 Minimum and Maximum Eigenvalues of Finite Sample Hessian and Scatter Matrices

The following technical lemma shows that the eigenvalue conditions of the expected Hessian and scatter matrices, hold with high probability in the finite sample case.

Lemma 2. If $\lambda_{\min}(\mathbf{H}_{SS}^*) \geq C_{\min}$, $\lambda_{\max}(\mathbb{E}_{\mathbf{x}}[\mathbf{z}_S \mathbf{z}_S^T]) \leq D_{\max}$, then we have that

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

with probability at least

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

respectively.

Proof. Let

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

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

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

Using the Matrix Chernoff bounds from Tropp [11], 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

$$\begin{aligned} \mu_{\max} &= \lambda_{\max}(\mathbb{E}_{\mathbf{x}}[\mathbf{z}_S \mathbf{z}_S^T]) \\ &\geq \lambda_{\min}(\mathbb{E}_{\mathbf{x}}[\mathbf{z}_S \mathbf{z}_S^T]) \geq \frac{2^n \tilde{p}_{\min}}{2^n - |\mathcal{N}\mathcal{E}^*|} \end{aligned}$$

Once again invoking Theorem 1.1 from [11] and setting $\delta = 1$ we have that

$$\begin{aligned} \Pr\{\lambda_{\max} \geq (1 + \delta)\mu_{\max}\} &\leq |S| \left[\frac{e^{\delta}}{(1 + \delta)^{1 + \delta}}\right]^{(m\mu_{\max})/R'} \\ \Pr\{\lambda_{\max} \geq 2\mu_{\max}\} &\leq |S| \left[\frac{e}{4}\right]^{(m\mu_{\max})/|S|} \\ &\leq |S| \exp\left(\frac{-m\mu_{\max}}{4|S|}\right) \\ &\leq |S| \exp\left(-\frac{m2^{n-2}\tilde{p}_{\min}}{|S|(2^n - |\mathcal{N}\mathcal{E}^*|)}\right) \\ &\leq |S| \exp\left(\frac{-m\tilde{p}_{\min}}{4|S|}\right) \end{aligned}$$

Therefore, we have that

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

□

4.3.2 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 (6) at the true vector \mathbf{v}^* , for all players.

Lemma 3. With probability at least $1 - \delta$ for $\delta \in (0, 1)$, we have that

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

where $\kappa = 1/(1 + \exp(\rho_{\min}))$, $\rho_{\min} \geq 0$ is the minimum payoff in the PSNE set, $f_{\mathcal{N}\mathcal{E}^*} = |\mathcal{N}\mathcal{E}^*|/2^{n-1}$, and

$$\nu \stackrel{\text{def}}{=} \kappa \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \mathcal{P}(x) + \frac{(\tilde{p}_{\max} - \tilde{p}_{\min})}{2 - f_{\mathcal{N}\mathcal{E}^*}} + \frac{f_{\mathcal{N}\mathcal{E}^*} \tilde{p}_{\min}}{2 - f_{\mathcal{N}\mathcal{E}^*}} \quad (9)$$

Proof. Consider the i -th player. Let $\mathbf{u}^m \stackrel{\text{def}}{=} \nabla \ell(\mathbf{v}_i^*, \mathcal{D})$ and u_j^m denote the j -th index of \mathbf{u}^m . For any subset $\mathcal{S}' \subset \mathcal{X}$ such that $|\mathcal{S}'| = 2^{n-1}$ define the function $g(\mathcal{S}')$ as follows:

$$g(\mathcal{S}') \stackrel{\text{def}}{=} \sum_{\mathbf{x} \in \mathcal{S}'} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) - \sum_{\mathbf{x} \in \mathcal{S}'^c} \mathcal{P}(\mathbf{x}) f(\mathbf{x}),$$

where \mathcal{S}'^c denotes the complement of the set \mathcal{S}' and $f(\mathbf{x}) = 1/(1 + \exp(\mathbf{v}_i^*{}^T \mathbf{z}_i(\mathbf{x})))$. For $\mathbf{x} \in \mathcal{N}\mathcal{E}^*$, $f(\mathbf{x}) \leq \kappa$, while for $\mathbf{x} \notin \mathcal{N}\mathcal{E}^*$ we have $1/2 \leq f(\mathbf{x}) \leq 1$. Lastly, let $\mathcal{S}_{ij} = \{\mathbf{x} \in \mathcal{X} | x_i x_j = +1\}$ and $\mathcal{S}_i = \{\mathbf{x} \in \mathcal{X} | x_i = +1\}$. From (7) we have that, for $j \neq n$, $|\mathbb{E}[u_j^m]| = |g(\mathcal{S}_{ij})|$, while for $j = n$ $|\mathbb{E}[u_j^m]| = |g(\mathcal{S}_i)|$. Thus we get

$$\|\mathbf{u}^m\|_{\infty} \leq \max_{\mathcal{S}' \subset \mathcal{X} | |\mathcal{S}'| = 2^{n-1}} g(\mathcal{S}') \quad (10)$$

Let \mathcal{S} be the set that maximizes (10), $A \stackrel{\text{def}}{=} \mathcal{S} \cap \mathcal{N}\mathcal{E}^*$ and $B \stackrel{\text{def}}{=} \mathcal{S}^c \cap \mathcal{N}\mathcal{E}^*$. Continuing from above,

$$\begin{aligned} |g(\mathcal{S})| &= \left| \sum_{\mathbf{x} \in \mathcal{S} \setminus A} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) + \sum_{\mathbf{x} \in A} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) \right. \\ &\quad \left. - \sum_{\mathbf{x} \in \mathcal{S}^c \setminus B} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) - \sum_{\mathbf{x} \in B} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) \right| \\ &\leq \kappa \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \mathcal{P}(x) + \left| \sum_{\mathbf{x} \in \mathcal{S} \setminus A} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) - \sum_{\mathbf{x} \in \mathcal{S}^c \setminus B} \mathcal{P}(\mathbf{x}) f(\mathbf{x}) \right| \end{aligned}$$

Assume that the first term inside the absolute value above dominates the second term, if not then we can proceed by reversing the two terms.

$$\begin{aligned} |g(\mathcal{S})| &\leq \kappa \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \mathcal{P}(x) + \frac{2^{n-1} \tilde{p}_{\max} - (2^{n-1} - |\mathcal{N}\mathcal{E}^*|) \tilde{p}_{\min}}{2^n - |\mathcal{N}\mathcal{E}^*|} \\ &= \kappa \sum_{\mathbf{x} \in \mathcal{N}\mathcal{E}^*} \mathcal{P}(x) + \frac{(\tilde{p}_{\max} - \tilde{p}_{\min}) + f_{\mathcal{N}\mathcal{E}^*} \tilde{p}_{\min}}{2 - f_{\mathcal{N}\mathcal{E}^*}} = \nu \end{aligned}$$

Also note that $|u_j^m| \leq \nu \leq 1$. Finally, from Hoeffding's inequality [12] and using a union bound argument over all players, we have that:

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

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

To get some intuition for the lemma above, consider the constant ν as given in (9). First, note that $\kappa \leq 1/2$. Also, as the minimum payoff ρ_{\min} increases, κ decays to 0 exponentially. Similarly, if the probability measure on the non-Nash equilibria set is close to uniform, meaning $\tilde{p}_{\max} - \tilde{p}_{\min} \approx 0$, then the second term in (9) vanishes. Finally, if the fraction of actions that are in the PSNE set ($f_{\mathcal{N}\mathcal{E}^*}$) is small, then the third term in (9) is small. Therefore, if the minimum payoff is high, the noise distribution, i.e., the distribution of the non-Nash equilibria joint actions, is close to uniform, and the fraction of joint actions that are in the PSNE set is small, then the expected gradient vanishes. 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 5, 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 4. *If the regularization parameter λ satisfies the following condition:*

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

then

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

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

Proof. The proof of this lemma follows the general proof structure of Lemma 3 in [9]. 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: $\tilde{f}(\Delta_S) = \ell(\mathbf{v}_S^* + \Delta_S) - \ell(\mathbf{v}_S^*) + \lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1)$, 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:

$$\begin{aligned} \tilde{f}(\Delta_S) &= \nabla \ell(\mathbf{v}_S^*)^T \Delta_S + \Delta_S^T \nabla^2 \ell(\mathbf{v}_S^* + \theta \Delta_S) \Delta_S \\ &\quad + \lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1), \end{aligned} \quad (11)$$

for some $\theta \in [0, 1]$. Next, we lower bound each of the terms in (11). Using the Cauchy-Schwartz inequality, the first term in (11) is bounded as follows:

$$\begin{aligned} \nabla \ell(\mathbf{v}_S^*)^T \Delta_S &\geq -\|\nabla \ell(\mathbf{v}_S^*)\|_\infty \|\Delta_S\|_1 \\ &\geq -\|\nabla \ell(\mathbf{v}_S^*)\|_\infty \sqrt{|S|} \|\Delta_S\|_2 \\ &\geq -b\sqrt{|S|} \left(\nu + \sqrt{\frac{2}{m} \log \frac{2n}{\delta}} \right), \end{aligned} \quad (12)$$

with probability at least $1 - \delta$ for $\delta \in [0, 1]$. It is also easy to upper bound the last term in equation 11, using the reverse triangle inequality as follows:

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

Which then implies the following lower bound:

$$\begin{aligned} \lambda(\|\mathbf{v}_S^* + \Delta_S\|_1 - \|\mathbf{v}_S^*\|_1) &\geq -\lambda \|\Delta_S\|_1 \\ &\geq -\lambda \sqrt{|S|} \|\Delta_S\|_2 \\ &= -\lambda \sqrt{|S|} b. \end{aligned} \quad (13)$$

Now we turn our attention to computing a lower bound of the second term of (11), 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 $\bar{\theta} \in [0, \theta]$. Finally, from Lemma 2 we have, with probability at least $1 - |S| \exp((-mC_{\min})/2|S|)$:

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

where we have defined

$$\begin{aligned} \mathbf{A}(\theta) \stackrel{\text{def}}{=} \frac{1}{m} \sum_{l=1}^m \eta'((\mathbf{v}_S^* + \theta\Delta_S)^T \mathbf{z}_S^{(l)}) \times \\ (\theta\Delta_S)^T \mathbf{z}_S^{(l)} \mathbf{z}_S^{(l)} (\mathbf{z}_S^{(l)})^T. \end{aligned}$$

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

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

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

$$\lambda_{\min}(\nabla^2 \ell(\mathbf{v}_S^* + \theta\Delta_S)) \geq \frac{C_{\min}}{4}. \quad (14)$$

Plugging back the bounds given by (12), (13) and (14) in (11) and equating to zero we get

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

Finally, coming back to our prior assumption we have

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

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

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

Lemma 5. *If the regularization parameter λ satisfies the following condition:*

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

then we have that

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

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

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:*

$$\begin{aligned} \nu + \sqrt{\frac{2}{m} \log \frac{6n^2}{\delta}} \leq \lambda \leq 2K + \nu - \sqrt{\frac{2}{m} \log \frac{6n^2}{\delta}} \quad (15) \\ m \geq \max \left\{ \frac{2}{K^2} \log \left(\frac{6n^2}{\delta} \right), \frac{2k}{C_{\min}} \log \left(\frac{3kn}{\delta} \right), \right. \\ \left. \frac{4k}{\bar{p}_{\min}} \log \left(\frac{3kn}{\delta} \right) \right\}, \quad (16) \end{aligned}$$

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

Proof. From Cauchy-Schwartz inequality and Lemma 5 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

$$1 - n(\delta + k \exp((-mC_{\min})/2k) + k \exp((-m\bar{p}_{\min})/4k)) \quad (17)$$

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 (17) are at most $\delta'/3n$ each, we prove our claim. \square

To better understand the implications of the theorem above, we instantiate it for the global and local noise model.

Remark 1 (Sample complexity under global noise model). *Recall that $D_{\max} \leq \min(k, 2^n p_{\max})$, and for the global noise model given by (3) $p_{\max} = q_g/|\mathcal{NE}^*|$. If q_g is constant, then $D_{\max} = k$. 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_g is small enough, i.e., $q_g = \mathcal{O}(|\mathcal{NE}^*|/2^n)$, then D_{\max} is no longer a function of k and $K = \Omega(1/k)$. Hence, the sample complexity scales as $\mathcal{O}(k^2 \log n)$ for exact PSNE recovery.*

Next, we consider the implications of Theorem 1 under the local noise model given by (4). we consider the regime where the parameter q scales with the number of players n .

Remark 2 (Sample complexity under local noise). *In the local noise model if the number of Nash-equilibria is constant, then $p_{\max} = \mathcal{O}(\exp(-n))$, and once again D_{\max} becomes independent of k , which results in a sample complexity of $\mathcal{O}(k^2 \log n)$.*

Also, observe the dependence of the sample complexity on the minimum noise level \bar{p}_{\min} . The number of samples required to recover the PSNE set increases as \bar{p}_{\min} decreases. From the aforementioned remarks we see that if the noise level is too low, i.e., $\bar{p}_{\min} \rightarrow 0$, then number of samples needed goes to infinity; This seems counter-intuitive — with reduced noise 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 (6). Then, the sample complexity as given by Theorem 1 can be written as $\mathcal{O}\left(\frac{k^2 D_{\max}^2}{C_{\min}^2} \log n\right)$.

Hence, we have that as the noise level gets too low, the Hessian of the loss becomes 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.

4.4 Necessary Conditions

In this section we derive necessary conditions on the number of samples required to learn graphical games. Our approach for doing so is information-theoretic: we treat the inference procedure as a communication channel and then use the Fano’s inequality to lower bound the estimation error. Such techniques have been widely used to obtain necessary conditions for model selection in graphical models, see e.g. [13, 14], sparsity recovery in linear regression [15], and many other problems.

Consider an ensemble \mathfrak{G}_n of n -player games with the in-degree of each player being at most k . Nature picks a true game $\mathcal{G}^* \in \mathfrak{G}$, and then generates a data set \mathcal{D} of m joint actions. A decoder is any function $\psi : \mathcal{X}^m \rightarrow \mathfrak{G}_n$ that maps a data set \mathcal{D} to a game, $\psi(\mathcal{D})$, in \mathfrak{G}_n . The minimum estimation error over all decoders ψ , for the ensemble \mathfrak{G}_n , is then given as follows:

$$p_{\text{err}} \stackrel{\text{def}}{=} \min_{\psi} \max_{\mathcal{G}^* \in \mathfrak{G}_n} \Pr \{ \mathcal{NE}(\psi(\mathcal{D})) \neq \mathcal{NE}(\mathcal{G}^*) \}, \quad (18)$$

where the probability is computed over the data distribution. Our objective here is to compute the number of samples below which PSNE recovery fails with probability greater than $1/2$.

Theorem 2. *The number of samples required to learn graphical games over n players and in-degree of at most k , is $\Omega(k \log n)$.*

Remark 3. *From the above theorem and from Theorem 1 we observe that the method of l_1 -regularized logistic regression for learning graphical games, operates close to the fundamental limit of $\Omega(k \log n)$.*

Results from simulation experiments for both global and local noise model can be found in Appendix B.

Concluding Remarks. 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. 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] V. Garg and T. Jaakkola, “Learning tree structured potential games,” *Neural Information Processing Systems*, vol. 29, 2016.
- [8] A. Ghoshal and J. Honorio, “From behavior to sparse graphical games: Efficient recovery of equilibria,” *arXiv preprint arXiv:1607.02959*, 2016.
- [9] 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>
- [10] 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.
- [11] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Foundations of computational mathematics*, vol. 12, no. 4, pp. 389–434, 2012.
- [12] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, vol. 58, no. 301, pp. 13–30, 1963.
- [13] N. P. Santhanam and M. J. Wainwright, “Information-theoretic limits of selecting binary graphical models in high dimensions,” *Information Theory, IEEE Transactions on*, vol. 58, no. 7, pp. 4117–4134, 2012.
- [14] W. Wang, M. J. Wainwright, and K. Ramchandran, “Information-theoretic bounds on model selection for Gaussian Markov random fields.” in *ISIT*. Citeseer, 2010, pp. 1373–1377.
- [15] M. J. Wainwright, “Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting,” *Information Theory, IEEE Transactions on*, vol. 55, no. 12, pp. 5728–5741, 2009.
- [16] B. Yu, *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*. New York, NY: Springer New York, 1997, ch. Assouad, Fano, and Le Cam, pp. 423–435. [Online]. Available: http://dx.doi.org/10.1007/978-1-4612-1880-7_29

Appendix A Detailed Proofs

Proof of Proposition 1. Let $|\mathcal{N}\mathcal{E}^*| > 2^{n-1}$. Then by the pigeon hole principle there are at least two joint actions \mathbf{x} and \mathbf{x}' in $\mathcal{N}\mathcal{E}^*$ such that $\mathbf{x} = -\mathbf{x}'$. Since the payoff is strictly positive, it follows that the bias b_i for each player must be 0. If the bias for all players is 0, then for each $\mathbf{x} \in \mathcal{N}\mathcal{E}^*$, $-\mathbf{x} \in \mathcal{N}\mathcal{E}^*$. Therefore, $|\mathcal{N}\mathcal{E}^*| = 2^n$. Since we have assumed that the game is non-trivial, we get a contradiction. \square

Proof of Lemma 5. Define $\Delta \stackrel{\text{def}}{=} \hat{\mathbf{v}} - \mathbf{v}^*$. 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,

$$\begin{aligned} \|\hat{\mathbf{v}}\|_1 &= \|\mathbf{v}^* + \Delta\|_1 = \|\mathbf{v}_{\bar{S}}^* + \Delta_{\bar{S}} + \Delta_{S^c}\|_1 \\ &= \|\mathbf{v}_{\bar{S}}^* - (-\Delta_{\bar{S}})\|_1 + \|\Delta_{S^c}\|_1 \\ &\geq \|\mathbf{v}^*\|_1 - \|\Delta_{\bar{S}}\|_1 + \|\Delta_{S^c}\|_1. \end{aligned} \quad (19)$$

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

$$\begin{aligned} \ell(\mathbf{v}^*) + \lambda \|\mathbf{v}^*\|_1 &\geq \ell(\hat{\mathbf{v}}) + \lambda \|\hat{\mathbf{v}}\|_1 \\ \implies \lambda(\|\mathbf{v}^*\|_1 - \|\hat{\mathbf{v}}\|_1) &\geq \ell(\hat{\mathbf{v}}) - \ell(\mathbf{v}^*). \end{aligned} \quad (20)$$

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

$$\begin{aligned} \ell(\hat{\mathbf{v}}) - \ell(\mathbf{v}^*) &\geq \nabla \ell(\mathbf{v}^*)^T (\hat{\mathbf{v}} - \mathbf{v}^*) \\ &\geq -\|\nabla \ell(\mathbf{v}^*)\|_\infty \|\Delta\|_1 \\ &\geq -\frac{\lambda}{2} \|\Delta\|_1, \end{aligned} \quad (21)$$

in the last line we used the fact that $\lambda \geq \|\nabla \ell(\mathbf{v}^*)\|_\infty$. Thus, we have from (19), (20) and (21) that

$$\begin{aligned} \frac{1}{2} \|\Delta\|_1 &\geq \|\hat{\mathbf{v}}\|_1 - \|\mathbf{v}^*\|_1 \\ \implies \frac{1}{2} \|\Delta\|_1 &\geq \|\Delta_{S^c}\|_1 - \|\Delta_{\bar{S}}\|_1 \\ \implies \frac{1}{2} \|\Delta_{S^c}\|_1 + \frac{1}{2} \|\Delta_{\bar{S}}\|_1 &\geq \|\Delta_{S^c}\|_1 - \|\Delta_{\bar{S}}\|_1 \\ \implies 3 \|\Delta_{\bar{S}}\|_1 &\geq \|\Delta_{S^c}\|_1. \end{aligned} \quad (22)$$

Finally, from (22) and Lemma 4 we have that

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

\square

Proof of Theorem 2. First, we construct a restricted ensemble of games $\tilde{\mathfrak{G}} \subset \mathfrak{G}$ as follows. Each game $\mathcal{G} \in \tilde{\mathfrak{G}}$ contains k , randomly chosen, *influential* players. The game graph for \mathcal{G} is then chosen to be a complete directed bipartite graph from the set of k influential players to the set of $n - k$ non-influential players. The edge weights are all set to -1 , the bias for the k influential players is set to $+1$, while the bias for the remaining $n - k$ players

is set to 0. Then it is clear that *each game in $\tilde{\mathfrak{G}}$ induces a distinct size-one PSNE set*. Specifically, for a game $\mathcal{G} \in \tilde{\mathfrak{G}}$, a joint action $\mathbf{x} \in \mathcal{NE}(\mathcal{G})$ is such that $x_i = -1$ if player i is influential in \mathcal{G} , otherwise $x_i = +1$. Also, note that the minimum payoff in the PSNE set of each game in $\tilde{\mathfrak{G}}$ is strictly positive, and is precisely 1. Finally, we assume that the data set is drawn according to the global noise model (3), with $q = 1/n$. Now let $\mathcal{G} \in \tilde{\mathfrak{G}}$ be a uniformly-distributed random variable corresponding to the game that was picked by nature. From the Fano's inequality, we have that:

$$p_{\text{err}} \geq 1 - \frac{I(\mathcal{D}; \mathcal{G}) + \log 2}{H(\mathcal{G})}, \quad (23)$$

where $I(\cdot)$ denotes mutual information and $H(\cdot)$ denotes entropy. Since, \mathcal{G} is uniformly distributed, we have that $H(\mathcal{G}) = \log |\tilde{\mathfrak{G}}| = \log \binom{n}{k} \geq k(\log n - \log k)$. Let $\mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_1}$ be the conditional distribution of the data set given a game $\mathcal{G}_1 \in \tilde{\mathfrak{G}}$. We bound the mutual information $I(\mathcal{D}; \mathcal{G})$ by a pairwise KL-based bound from [16] as follows:

$$I(\mathcal{D}; \mathcal{G}) \leq \frac{1}{|\tilde{\mathfrak{G}}|} \sum_{\mathcal{G}_1 \in \tilde{\mathfrak{G}}} \sum_{\mathcal{G}_2 \in \tilde{\mathfrak{G}}} \mathbb{KL}(\mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_1} \parallel \mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_2}). \quad (24)$$

Now from the fact that data are sampled i.i.d, we get:

$$\begin{aligned} & \mathbb{KL}(\mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_1} \parallel \mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_2}) \\ &= m \sum_{\mathbf{x} \in \mathcal{X}} \mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_1}(\mathbf{x}) \log \frac{\mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_1}(\mathbf{x})}{\mathcal{P}_{\mathcal{D}|\mathcal{G}=\mathcal{G}_2}(\mathbf{x})} \\ &= m \left\{ q \log \frac{q(2^n-1)}{1-q} + \frac{1-q}{2^n-1} \log \frac{1-q}{q(2^n-1)} \right\} \\ &= \frac{m(2^n q-1)}{2^n-1} \left(\log q - \log \left(\frac{1-q}{2^n-1} \right) \right) \\ &\leq m \log 2, \end{aligned} \quad (25)$$

where the last line comes from the fact that $q = 1/n$. Putting together (23), (24) and (25), and setting $p_{\text{err}} = 1/2$, we get

$$m \leq \frac{k \log n - k \log k - 2 \log 2}{2 \log 2}.$$

By observing that learning the ensemble \mathfrak{G} is at least as hard as learning a subset of the ensemble $\tilde{\mathfrak{G}}$, we prove our main claim. \square

Appendix B 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. For experiments involving the local noise model, we only used $n \in \{10, 12, 15\}$. The parameter δ was set to the constant value of 0.01. For the global noise model, the parameters q_g was set to 0.01, while for the local noise model we used $q_1 = \dots = q_n = 0.6$. 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.

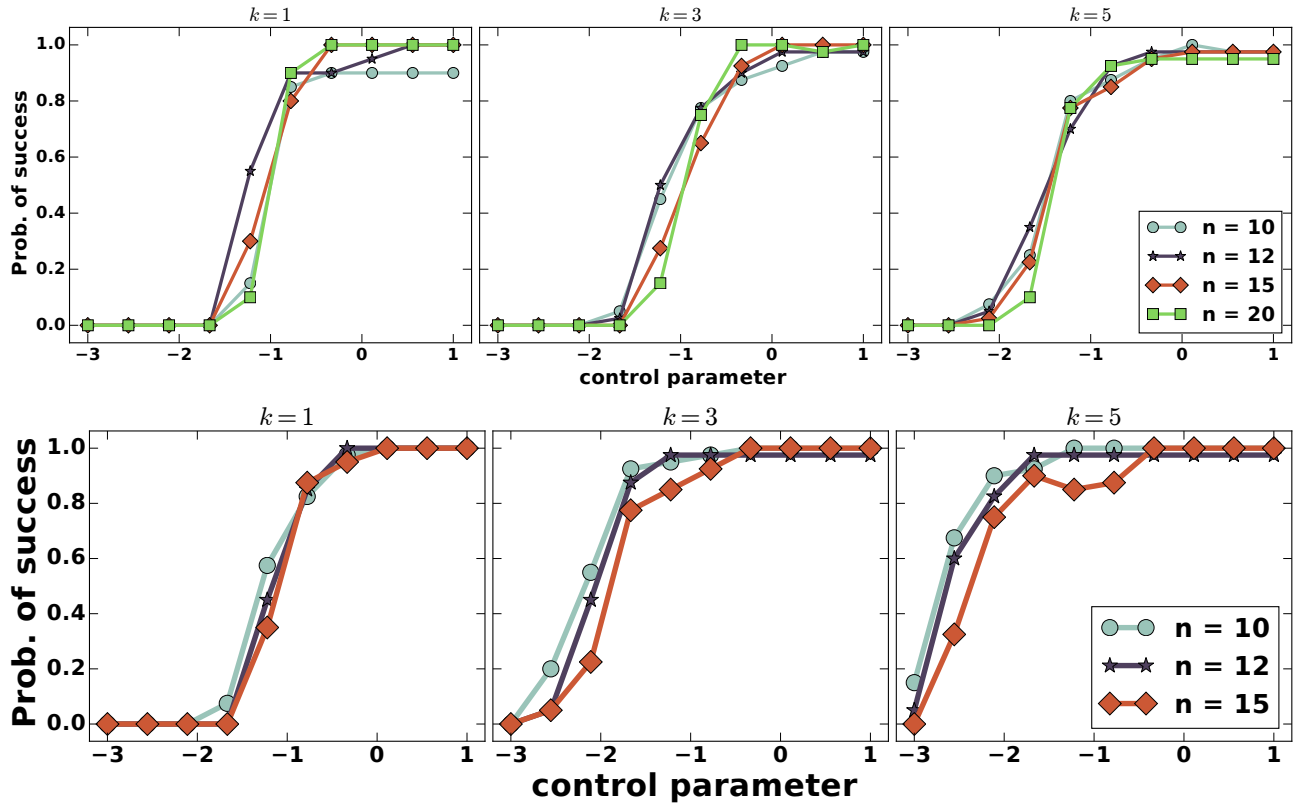


Figure 1: The probability of exact recovery of the PSNE set computed across 40 randomly sampled LIGs, using the global noise model (TOP) and local noise model (BOTTOM), as the number of samples is scaled as $\lfloor (C)(10^c)(k^2 \log(6n^2/\delta)) \rfloor$, where c is the control parameter and the constant C is 10000 for the $k = 1$ case and 1000 for the remaining two case. For the global noise model we set $q_g = 0.001$, while for the local noise model we used $q_1 = \dots = q_n = 0.6$.