Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data: Supplementary Material

Jean Honorio, Luis Ortiz
Department of Computer Science
Stony Brook University
Stony Brook, NY 11794
{jhonorio,leortiz}@cs.sunysb.edu

A On the Identifiability of Games

Several games with different coefficients can lead to the same PSNE set. As a simple example that illustrates the issue of identifiability, consider the three following LIGs with the same PSNE sets, i.e. $\mathcal{NE}(\mathbf{W}_k, \mathbf{0}) = \{(-1,-1,-1),(+1,+1,+1)\}$ for $k = 1,2,3$:

$$\mathbf{W}_1 = \begin{bmatrix} 0 & 0 & 0 \\ 1/2 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{W}_2 = \begin{bmatrix} 0 & 0 & 0 \\ 2 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \quad \mathbf{W}_3 = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

Clearly, using structural properties alone, one would generally prefer the former two models to the latter, all else being equal (e.g. generalization performance). A large number of the econometrics literature concerns the issue of identifiability of models from data. In typical machine-learning fashion, we side-step this issue by measuring the quality of our data-induced models via their generalization ability and invoke the principle of Ockham's razor to bias our search toward simpler models using well-known and -studied regularization techniques. In particular, we take the view that games are identifiable by their PSNE sets.

B Detailed Proofs

In this section, we show the detailed proofs of lemmas and theorems for which we provide only proof sketches.

B.1 Proof of Proposition 3

Proof. Note that $p_{(\mathcal{G},q)}(\mathbf{x}_1) = q/|\mathcal{NE}(\mathcal{G})| > p_{(\mathcal{G},q)}(\mathbf{x}_2) = (1-q)/(2^n - |\mathcal{NE}(\mathcal{G})|) \Leftrightarrow q > |\mathcal{NE}(\mathcal{G})|/2^n$ and given eq.(2), we prove our claim. $\square$

B.2 Proof of Proposition 4

Proof. Let $\mathcal{NE}_k \equiv \mathcal{NE}(\mathcal{G}_k)$. First, we prove the $\Rightarrow$ direction. By Definition 1, $\mathcal{G}_1 \equiv_{\mathcal{NE}} \mathcal{G}_2 \Rightarrow \mathcal{NE}_1 = \mathcal{NE}_2$. Note that $p_{(\mathcal{G}_k,q)}(\mathbf{x})$ in eq.(1) depends only on characteristic functions $1[\mathbf{x} \in \mathcal{NE}_k]$. Therefore, $(\forall \mathbf{x})\ p_{(\mathcal{G}_1,q)}(\mathbf{x}) = p_{(\mathcal{G}_2,q)}(\mathbf{x})$.

Second, we prove the $\Leftarrow$ direction by contradiction. Assume $(\exists \mathbf{x})\ \mathbf{x} \in \mathcal{NE}_1 \wedge \mathbf{x} \notin \mathcal{NE}_2$. $p_{(\mathcal{G}_1,q)}(\mathbf{x}) = p_{(\mathcal{G}_2,q)}(\mathbf{x})$ implies that $q/|\mathcal{NE}_1| = (1-q)/(2^n - |\mathcal{NE}_2|) \Rightarrow q = |\mathcal{NE}_1|/(2^n + |\mathcal{NE}_1| - |\mathcal{NE}_2|)$. Since $q > \max(\pi(\mathcal{G}_1), \pi(\mathcal{G}_2)) \Rightarrow q > \max(|\mathcal{NE}_1|, |\mathcal{NE}_2|)/2^n$ by eq.(2). Therefore $\max(|\mathcal{NE}_1|, |\mathcal{NE}_2|)/2^n < |\mathcal{NE}_1|/(2^n + |\mathcal{NE}_1| - |\mathcal{NE}_2|)$. If we assume that $|\mathcal{NE}_1| \geq |\mathcal{NE}_2|$ we reach the contradiction $|\mathcal{NE}_1| - |\mathcal{NE}_2| < 0$. If we assume that $|\mathcal{NE}_1| \leq |\mathcal{NE}_2|$ we reach the contradiction $(2^n - |\mathcal{NE}_2|)(|\mathcal{NE}_2| - |\mathcal{NE}_1|) < 0$. $\square$

B.3 Proof of Lemma 5

Proof. Let $\mathcal{NE} \equiv \mathcal{NE}(\mathcal{G})$, $\pi \equiv \pi(\mathcal{G})$ and $\widehat{\pi} \equiv \widehat{\pi}(\mathcal{G})$. First, for a non-trivial $\mathcal{G}$, $\log p_{(\mathcal{G},q)}(\mathbf{x}^{(l)}) = \log \frac{q}{|\mathcal{NE}|}$ for $\mathbf{x}^{(l)} \in \mathcal{NE}$, and $\log p_{(\mathcal{G},q)}(\mathbf{x}^{(l)}) = \log \frac{1-q}{2^n-|\mathcal{NE}|}$ for $\mathbf{x}^{(l)} \notin \mathcal{NE}$. The average log-likelihood $\widehat{\mathcal{L}}(\mathcal{G},q) = \frac{1}{m}\sum_l \log p_{\mathcal{G},q}(\mathbf{x}^{(l)}) = \widehat{\pi}\log\frac{q}{|\mathcal{NE}|} + (1-\widehat{\pi})\log\frac{1-q}{2^n-|\mathcal{NE}|} = \widehat{\pi}\log\frac{q}{\pi} + (1-\widehat{\pi})\log\frac{1-q}{1-\pi} - n\log 2$. By adding $0 = -\widehat{\pi}\log\widehat{\pi} + \widehat{\pi}\log\widehat{\pi} - (1-\widehat{\pi})\log(1-\widehat{\pi}) + (1-\widehat{\pi})\log(1-\widehat{\pi})$, this can be rewritten as $\widehat{\mathcal{L}}(\mathcal{G},q) = \widehat{\pi}\log\frac{\widehat{\pi}}{\pi} + (1-\widehat{\pi})\log\frac{1-\widehat{\pi}}{1-\pi} - \widehat{\pi}\log\frac{\widehat{\pi}}{q} - (1-\widehat{\pi})\log\frac{1-\widehat{\pi}}{1-q} - n\log 2$, and by using eq.(3) we prove our claim.

Note that by maximizing with respect to the mixture parameter q and by properties of the KL divergence, we get $KL(\widehat{\pi}\|\widehat{q}) = 0 \Leftrightarrow \widehat{q} = \widehat{\pi}$. We define our hypothesis space Υ given the conditions in Remark 2 and Propositions 3 and 4. For the case $\widehat{\pi} = 1$, we "shrink" the mixture parameter to $1 - \frac{1}{2m}$ in order to avoid generating an invalid PMF as discussed in Remark 2. □

B.4 Proof of Lemma 7

Proof. Let $\mathcal{NE} \equiv \mathcal{NE}(\mathcal{G})$, $\pi \equiv \pi(\mathcal{G})$, $\widehat{\pi} \equiv \widehat{\pi}(\mathcal{G})$, $\bar{\pi} \equiv \bar{\pi}(\mathcal{G})$, and $\mathbb{E}[\cdot]$ and $\mathbb{P}[\cdot]$ be the expectation and probability with respect to the ground truth distribution $\mathcal{Q}$ of the data.

First note that for any $\theta = (\mathcal{G}, q)$, we have $\bar{\mathcal{L}}(\theta) = \mathbb{E}[\log p_{(\mathcal{G},q)}(\mathbf{x})] = \mathbb{E}[1[\mathbf{x}\in\mathcal{NE}]\log\frac{q}{|\mathcal{NE}|} + 1[\mathbf{x}\notin\mathcal{NE}]\log\frac{1-q}{2^n-|\mathcal{NE}|}] = \mathbb{P}[\mathbf{x}\in\mathcal{NE}]\log\frac{q}{|\mathcal{NE}|} + \mathbb{P}[\mathbf{x}\notin\mathcal{NE}]\log\frac{1-q}{2^n-|\mathcal{NE}|} = \bar{\pi}\log\frac{q}{|\mathcal{NE}|} + (1-\bar{\pi})\log\frac{1-q}{2^n-|\mathcal{NE}|} = \bar{\pi}\log\left(\frac{q}{1-q}\cdot\frac{2^n-|\mathcal{NE}|}{|\mathcal{NE}|}\right) + \log\frac{1-q}{2^n-|\mathcal{NE}|} = \bar{\pi}\log\left(\frac{q}{1-q}\cdot\frac{1-\pi}{\pi}\right) + \log\frac{1-q}{1-\pi} - n\log 2$.

Similarly, for any $\theta = (\mathcal{G}, q)$, we have $\widehat{\mathcal{L}}(\theta) = \widehat{\pi}\log\left(\frac{q}{1-q}\cdot\frac{1-\pi}{\pi}\right) + \log\frac{1-q}{1-\pi} - n\log 2$. So that $\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta) = (\widehat{\pi} - \bar{\pi})\log\left(\frac{q}{1-q}\cdot\frac{1-\pi}{\pi}\right)$.

Furthermore, the function $\frac{q}{1-q}$ is strictly monotonically increasing for $0 \le q < 1$. If $\widehat{\pi} > \bar{\pi}$ then $-\epsilon/2 \le \widehat{\mathcal{L}}(\theta'') - \bar{\mathcal{L}}(\theta'') < \widehat{\mathcal{L}}(\theta') - \bar{\mathcal{L}}(\theta') < \widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta) \le \epsilon/2$. Else, if $\widehat{\pi} < \bar{\pi}$, we have $\epsilon/2 \ge \widehat{\mathcal{L}}(\theta'') - \bar{\mathcal{L}}(\theta'') > \widehat{\mathcal{L}}(\theta') - \bar{\mathcal{L}}(\theta') > \widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta) \ge -\epsilon/2$. Finally, if $\widehat{\pi} = \bar{\pi}$ then $\widehat{\mathcal{L}}(\theta'') - \bar{\mathcal{L}}(\theta'') = \widehat{\mathcal{L}}(\theta') - \bar{\mathcal{L}}(\theta') = \widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta) = 0$. □

B.5 Proof of Theorem 8

Proof. First our objective is to find a lower bound for $\mathbb{P}[\bar{\mathcal{L}}(\widehat{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) \ge -\epsilon] \ge \mathbb{P}[\bar{\mathcal{L}}(\widehat{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) \ge -\epsilon + (\widehat{\mathcal{L}}(\widehat{\theta}) - \widehat{\mathcal{L}}(\bar{\theta}))] \ge \mathbb{P}[-\widehat{\mathcal{L}}(\widehat{\theta}) + \bar{\mathcal{L}}(\widehat{\theta}) \ge -\frac{\epsilon}{2}, \widehat{\mathcal{L}}(\bar{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) \ge -\frac{\epsilon}{2}] = \mathbb{P}[\widehat{\mathcal{L}}(\widehat{\theta}) - \bar{\mathcal{L}}(\widehat{\theta}) \le \frac{\epsilon}{2}, \widehat{\mathcal{L}}(\bar{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) \ge -\frac{\epsilon}{2}] = 1 - \mathbb{P}[\widehat{\mathcal{L}}(\widehat{\theta}) - \bar{\mathcal{L}}(\widehat{\theta}) > \frac{\epsilon}{2} \vee \widehat{\mathcal{L}}(\bar{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) < -\frac{\epsilon}{2}]$.

Let $\widetilde{q} \equiv \max(1 - \frac{1}{2m}, \bar{q})$. Now, we have $\mathbb{P}[\widehat{\mathcal{L}}(\widehat{\theta}) - \bar{\mathcal{L}}(\widehat{\theta}) > \frac{\epsilon}{2} \vee \widehat{\mathcal{L}}(\bar{\theta}) - \bar{\mathcal{L}}(\bar{\theta}) < -\frac{\epsilon}{2}] \le \mathbb{P}[(\exists\theta \in \Upsilon, q \le \widetilde{q})\ |\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| > \frac{\epsilon}{2}] = \mathbb{P}[(\exists\theta, \mathcal{G} \in \mathcal{H}, q \in \{\pi(\mathcal{G}), \widetilde{q}\})\ |\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| > \frac{\epsilon}{2}]$. The last equality follows from invoking Lemma 7.

Note that $\mathbb{E}[\widehat{\mathcal{L}}(\theta)] = \bar{\mathcal{L}}(\theta)$ and that since $\pi(\mathcal{G}) \le q \le \widetilde{q}$, the log-likelihood is bounded as $(\forall\mathbf{x})\ -B \le \log p_{(\mathcal{G},q)}(\mathbf{x}) \le 0$, where $B = \log\frac{1}{1-\widetilde{q}} + n\log 2 = \log\max(2m, \frac{1}{1-\bar{q}}) + n\log 2$. Therefore, by Hoeffding's inequality, we have $\mathbb{P}[|\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| > \frac{\epsilon}{2}] \le 2e^{-\frac{m\epsilon^2}{2B^2}}$.

Furthermore, note that there are $2d(\mathcal{H})$ possible parameters θ, since we need to consider only two values of $q \in \{\pi(\mathcal{G}), \widetilde{q}\}$ and because the number of all possible games in $\mathcal{H}$ (identified by their Nash equilibria sets) is $d(\mathcal{H}) \equiv |\cup_{\mathcal{G}\in\mathcal{H}}\{\mathcal{NE}(\mathcal{G})\}|$. Therefore, by the union bound we get the following uniform convergence $\mathbb{P}[(\exists\theta, \mathcal{G} \in \mathcal{H}, q \in \{\pi(\mathcal{G}), \widetilde{q}\})\ |\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| > \frac{\epsilon}{2}] \le 4d(\mathcal{H})\mathbb{P}[|\widehat{\mathcal{L}}(\theta) - \bar{\mathcal{L}}(\theta)| > \frac{\epsilon}{2}] \le 4d(\mathcal{H})e^{-\frac{m\epsilon^2}{2B^2}} = \delta$. Finally, by solving for δ we prove our claim. □

B.6 Proof of Theorem 9

Proof. The logarithm of the number of possible pure-strategy Nash equilibria sets supported by $\mathcal{H}$ (i.e., that can be produced by some game in $\mathcal{H}$) is upper bounded by the VC-dimension of the class of neural networks with a single hidden layer of n units and $n + \binom{n}{2}$ input units, linear threshold activation functions, and constant output weights.

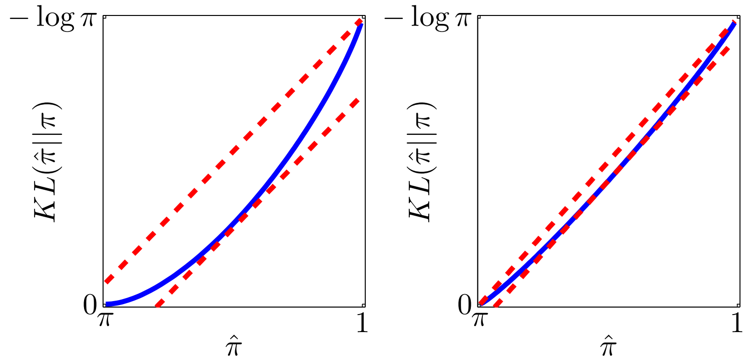

Figure 1: KL divergence (blue) and bounds in Lemma 10 (red) for $\pi = (3/4)^n$ where $n = 9$ (left) and $n = 36$ (right). Bounds are very informative when $n \to +\infty$.

For every linear influence game $\mathcal{G} = (\mathbf{W}, \mathbf{b})$ in $\mathcal{H}$, define the neural network with a single layer of n hidden units, n of the inputs corresponds to the linear terms $x_1, \ldots, x_n$ and $\binom{n}{2}$ corresponds to the quadratic polynomial terms $x_i x_j$ for all pairs of players (i, j), $1 \leq i < j \leq n$. For every hidden unit i, the weights corresponding to the linear terms $x_1, \ldots, x_n$ are $-b_1, \ldots, -b_n$, respectively, while the weights corresponding to the quadratic terms $x_i x_j$ are $-w_{ij}$, for all pairs of players (i, j), $1 \leq i < j \leq n$, respectively. The weights of the bias term of all the hidden units are set to 0. All n output weights are set to 1 while the weight of the output bias term is set to 0. The output of the neural network is $1[\mathbf{x} \notin \mathcal{NE}(\mathcal{G})]$. Note that we define the neural network to classify non-equilibrium as opposed to equilibrium to keep the convention in the neural network literature to define the threshold function to output 0 for input 0. The alternative is to redefine the threshold function to output 1 instead for input 0.

Finally, we use the VC-dimension of neural networks [Sontag, 1998]. □

B.7 Proof of Lemma 10

Proof. Let $\pi \equiv \pi(\mathcal{G})$ and $\widehat{\pi} \equiv \widehat{\pi}(\mathcal{G})$. Note that $\alpha(\pi) \equiv \lim_{\widehat{\pi} \to 0} KL(\widehat{\pi}\|\pi) = 0$ and $\beta(\pi) \equiv \lim_{\widehat{\pi} \to 1} KL(\widehat{\pi}\|\pi) = -\log \pi \leq n \log 2$. Since the function is convex we can upper-bound it by $\alpha(\pi) + (\beta(\pi) - \alpha(\pi))\widehat{\pi} = -\widehat{\pi} \log \pi$.

To find a lower bound, we find the point in which the derivative of the original function is equal to the slope of the upper bound, i.e. $\frac{\partial KL(\widehat{\pi}\|\pi)}{\partial \widehat{\pi}} = \beta(\pi) - \alpha(\pi) = -\log \pi$, which gives $\widehat{\pi}^* = \frac{1}{2-\pi}$. Then, the maximum difference between the upper bound and the original function is given by $\lim_{\pi \to 0} -\widehat{\pi}^* \log \pi - KL(\widehat{\pi}^*\|\pi) = \log 2$. □

B.7.1 Tightness of KL Divergence Bounds

The bounds on the KL divergence are very informative when $\pi(\mathcal{G}) \to 0$ (or in our setting when $n \to +\infty$), since $\log 2$ becomes small when compared to $-\log \pi(\mathcal{G})$, as shown in Figure 1.

B.8 Proof of Theorem 11

Proof. By applying the lower bound in Lemma 10 in eq.(5) to non-trivial games, we have $\widehat{\mathcal{L}}(\mathcal{G}, \widehat{q}) = KL(\widehat{\pi}(\mathcal{G})\|\pi(\mathcal{G})) - KL(\widehat{\pi}(\mathcal{G})\|\widehat{q}) - n \log 2 > -\widehat{\pi}(\mathcal{G}) \log \pi(\mathcal{G}) - KL(\widehat{\pi}(\mathcal{G})\|\widehat{q}) - (n+1) \log 2$. Since $\pi(\mathcal{G}) \leq \frac{\kappa^n}{\delta}$, we have $-\log \pi(\mathcal{G}) \geq -\log \frac{\kappa^n}{\delta}$. Therefore $\widehat{\mathcal{L}}(\mathcal{G}, \widehat{q}) > -\widehat{\pi}(\mathcal{G}) \log \frac{\kappa^n}{\delta} - KL(\widehat{\pi}(\mathcal{G})\|\widehat{q}) - (n+1) \log 2$. Regarding the term $KL(\widehat{\pi}(\mathcal{G})\|\widehat{q})$, if $\widehat{\pi}(\mathcal{G}) < 1 \Rightarrow KL(\widehat{\pi}(\mathcal{G})\|\widehat{q}) = KL(\widehat{\pi}(\mathcal{G})\|\widehat{\pi}(\mathcal{G})) = 0$, and if $\widehat{\pi}(\mathcal{G}) = 1 \Rightarrow KL(\widehat{\pi}(\mathcal{G})\|\widehat{q}) = KL(1\|1 - \frac{1}{2m}) = -\log(1 - \frac{1}{2m}) \leq \log 2$ and approaches 0 when $m \to +\infty$. Maximizing the lower bound of the log-likelihood becomes $\max_{\mathcal{G} \in \mathcal{H}} \widehat{\pi}(\mathcal{G})$ by removing the constant terms that do not depend on $\mathcal{G}$.

In order to prove $(\mathcal{G}, \widehat{q}) \in \Upsilon$ we need to prove $0 < \pi(\mathcal{G}) < \widehat{q} < 1$. For proving the first inequality $0 < \pi(\mathcal{G})$, note that $\widehat{\pi}(\mathcal{G}) \geq \gamma > 0$, and therefore $\mathcal{G}$ has at least one equilibria. For proving the third inequality $\widehat{q} < 1$, note that $\widehat{q} = \min(\widehat{\pi}(\mathcal{G}), 1 - \frac{1}{2m}) < 1$. For proving the second inequality $\pi(\mathcal{G}) < \widehat{q}$, we need to prove $\pi(\mathcal{G}) < \widehat{\pi}(\mathcal{G})$ and $\pi(\mathcal{G}) < 1 - \frac{1}{2m}$. Since $\pi(\mathcal{G}) \leq \frac{\kappa^n}{\delta}$ and $\gamma \leq \widehat{\pi}(\mathcal{G})$, it suffices to prove $\frac{(3/4)^n}{\delta} < \gamma \Rightarrow$

$\pi(\mathcal{G}) < \widehat{\pi}(\mathcal{G})$. Similarly we need to prove $\frac{(3/4)^n}{\delta} < 1 - \frac{1}{2m} \Rightarrow \pi(\mathcal{G}) < 1 - \frac{1}{2m}$. Putting both together, we have $\frac{(3/4)^n}{\delta} < \min(\gamma, 1 - \frac{1}{2m}) = \gamma$ since $\gamma < 1/2$ and $1 - \frac{1}{2m} \geq 1/2$. Finally, $\frac{(3/4)^n}{\delta} < \gamma \Leftrightarrow n > \log_\kappa(\delta\gamma)$. □

B.9 Proof of Lemma 13

Proof. Let $f_i(\mathbf{x}_{-i}) \equiv \mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i$. By noting that $\max(\alpha, \beta) = \max_{0 \leq u \leq 1}(\alpha + u(\beta - \alpha))$, we can rewrite $\widehat{\ell}(\mathbf{w}_{i,-i}, b_i) = \frac{1}{m}\sum_l \max_{0 \leq u^{(l)} \leq 1}(c^{(l)} + u^{(l)}(1 - x_i^{(l)} f_i(\mathbf{x}_{-i}^{(l)}) - c^{(l)}))$.

Note that $\widehat{\ell}$ has the minimizer $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$ if and only if $\mathbf{0}$ belongs to the subdifferential set of the non-smooth function $\widehat{\ell}$ at $(\mathbf{w}_{i,-i}, b_i) = \mathbf{0}$. In order to maximize $\widehat{\ell}$, we have $c^{(l)} > 1 - x_i^{(l)} f_i(\mathbf{x}_{-i}^{(l)}) \Leftrightarrow u^{(l)} = 0$, $c^{(l)} < 1 - x_i^{(l)} f_i(\mathbf{x}_{-i}^{(l)}) \Leftrightarrow u^{(l)} = 1$ and $c^{(l)} = 1 - x_i^{(l)} f_i(\mathbf{x}_{-i}^{(l)}) \Leftrightarrow u^{(l)} \in [0; 1]$. The previous rules simplify at the solution under analysis, since $(\mathbf{w}_{i,-i}, b_i) = \mathbf{0} \Rightarrow f_i(\mathbf{x}_{-i}^{(l)}) = 0$.

Let $g_j(\mathbf{w}_{i,-i}, b_i) \equiv \frac{\partial \widehat{\ell}}{\partial w_{ij}}(\mathbf{w}_{i,-i}, b_i)$ and $h(\mathbf{w}_{i,-i}, b_i) \equiv \frac{\partial \widehat{\ell}}{\partial b_i}(\mathbf{w}_{i,-i}, b_i)$. By making $(\forall j \neq i)\ 0 \in g_j(\mathbf{0}, 0)$ and $0 \in h(\mathbf{0}, 0)$, we get $(\forall j \neq i)\ \sum_l x_i^{(l)} x_j^{(l)} u^{(l)} = 0$ and $\sum_l x_i^{(l)} u^{(l)} = 0$. Finally, by noting that $x_i^{(l)} \in \{-1, 1\}$, we prove our claim. □

B.10 Proof of Lemma 15

Proof. Note that $\widehat{\ell}$ has the minimizer $(\mathbf{w}^*_{i,-i}, b^*_i) = \mathbf{0}$ if and only if the gradient of the smooth function $\widehat{\ell}$ is $\mathbf{0}$ at $(\mathbf{w}_{i,-i}, b_i) = \mathbf{0}$. Let $g_j(\mathbf{w}_{i,-i}, b_i) \equiv \frac{\partial \widehat{\ell}}{\partial w_{ij}}(\mathbf{w}_{i,-i}, b_i)$ and $h(\mathbf{w}_{i,-i}, b_i) \equiv \frac{\partial \widehat{\ell}}{\partial b_i}(\mathbf{w}_{i,-i}, b_i)$. By making $(\forall j \neq i)\ g_j(\mathbf{0}, 0) = 0$ and $h(\mathbf{0}, 0) = 0$, we get $(\forall j \neq i)\ \sum_l \frac{x_i^{(l)} x_j^{(l)}}{c^{(l)}+2} = 0$ and $\sum_l \frac{x_i^{(l)}}{c^{(l)}+2} = 0$. Finally, by noting that $x_i^{(l)} \in \{-1, 1\}$, we prove our claim. □

B.11 Proof of Lemma 17

Proof. Let $f_i(\mathbf{x}_{-i}) \equiv \mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i$. For proving Claim i, note that $1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})] = \min_j 1[x_j f_j(\mathbf{x}_{-j}) \geq 0] = 1[x_i f_i(\mathbf{x}_{-i}) \geq 0] \min_{j \neq i} 1[x_j f_j(\mathbf{x}_{-j}) \geq 0]$. Since all players except i are absolutely-indifferent, we have $(\forall j \neq i)\ (\mathbf{w}_{j,-j}, b_j) = \mathbf{0} \Rightarrow f_j(\mathbf{x}_{-j}) = 0 \Rightarrow \min_{j \neq i} 1[x_j f_j(\mathbf{x}_{-j}) \geq 0] = 1$. Therefore, $1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})] = 1[x_i f_i(\mathbf{x}_{-i}) \geq 0]$.

For proving Claim ii, by Claim i we have $|\mathcal{NE}(\mathcal{G})| = \sum_{\mathbf{x}} 1[x_i f_i(\mathbf{x}_{-i}) \geq 0]$. We can rewrite $|\mathcal{NE}(\mathcal{G})| = \sum_{\mathbf{x}} 1[x_i = +1] 1[f_i(\mathbf{x}_{-i}) \geq 0] + \sum_{\mathbf{x}} 1[x_i = -1] 1[f_i(\mathbf{x}_{-i}) \leq 0]$ or equivalently $|\mathcal{NE}(\mathcal{G})| = \sum_{\mathbf{x}_{-i}} 1[f_i(\mathbf{x}_{-i}) \geq 0] + \sum_{\mathbf{x}_{-i}} 1[f_i(\mathbf{x}_{-i}) \leq 0] = 2^{n-1} + \sum_{\mathbf{x}_{-i}} 1[f_i(\mathbf{x}_{-i}) = 0]$.

For proving Claim iii, by eq.(2) and Claim ii we have $\pi(\mathcal{G}) = \frac{|\mathcal{NE}(\mathcal{G})|}{2^n} = \frac{1}{2} + \frac{1}{2^n}\alpha(\mathbf{w}_{i,-i}, b_i)$, where $\alpha(\mathbf{w}_{i,-i}, b_i) \equiv \sum_{\mathbf{x}_{-i}} 1[\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i = 0]$. This proves the lower bound $\pi(\mathcal{G}) \geq \frac{1}{2}$. Geometrically speaking, $\alpha(\mathbf{w}_{i,-i}, b_i)$ is the number of vertices of the $(n-1)$-dimensional hypercube that are covered by the hyperplane with normal $\mathbf{w}_{i,-i}$ and bias b_i. Recall that $(\mathbf{w}_{i,-i}, b_i) \neq \mathbf{0}$. If $\mathbf{w}_{i,-i} = \mathbf{0}$ and $b_i \neq 0$ then $\alpha(\mathbf{w}_{i,-i}, b_i) = \sum_{\mathbf{x}_{-i}} 1[b_i = 0] = 0 \Rightarrow \pi(\mathcal{G}) = \frac{1}{2}$. If $\mathbf{w}_{i,-i} \neq \mathbf{0}$ then as noted in Aichholzer and Aurenhammer [1996] a hyperplane with $n-2$ zeros on $\mathbf{w}_{i,-i}$ (i.e. a $(n-2)$-*parallel hyperplane*) covers exactly half of the 2^{n-1} vertices, the maximum possible. Therefore, $\pi(\mathcal{G}) = \frac{1}{2} + \frac{1}{2^n}\alpha(\mathbf{w}_{i,-i}, b_i) \leq \frac{1}{2} + \frac{2^{n-2}}{2^n} = \frac{3}{4}$. □

B.12 Proof of Theorem 18

Proof. Let $y_i \equiv 1[x_i(\mathbf{w}_{i,-i}{}^{\mathrm{T}}\mathbf{x}_{-i} - b_i) \geq 0]$, $\mathcal{P} \equiv \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$ and $\mathcal{U}$ the uniform distribution for $\mathbf{x} \in \{-1, +1\}^n$. By eq.(2), $\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})] = \mathbb{E}_{\mathcal{P}}[\frac{1}{2^n}\sum_{\mathbf{x}}\prod_i y_i] = \mathbb{E}_{\mathcal{P}}[\mathbb{E}_{\mathcal{U}}[\prod_i y_i]] = \mathbb{E}_{\mathcal{U}}[\mathbb{E}_{\mathcal{P}}[\prod_i y_i]]$. Note that each y_i is independent since each $(\mathbf{w}_{i,-i}, b_i)$ is independently distributed. Therefore, $\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})] = \mathbb{E}_{\mathcal{U}}[\prod_i \mathbb{E}_{\mathcal{P}_i}[y_i]]$. Similarly each $z_i \equiv \mathbb{E}_{\mathcal{P}_i}[y_i]$ is independent since each $(\mathbf{w}_{i,-i}, b_i)$ is independently distributed. Therefore, $\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})] = \mathbb{E}_{\mathcal{U}}[\prod_i z_i] = \prod_i \mathbb{E}_{\mathcal{U}}[z_i] = \prod_i \mathbb{E}_{\mathcal{U}}[\mathbb{E}_{\mathcal{P}_i}[y_i]] = \prod_i \mathbb{E}_{\mathcal{P}_i}[\mathbb{E}_{\mathcal{U}}[y_i]]$. Note that $\mathbb{E}_{\mathcal{U}}[y_i]$ is the true proportion of equilibria of an influence game with non-absolutely-indifferent player i and absolutely-indifferent players $\forall j \neq i$, and therefore $1/2 \leq \mathbb{E}_{\mathcal{U}}[y_i] \leq 3/4$ by Claim iii of Lemma 17. Finally, we have $\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})] \geq \prod_i \mathbb{E}_{\mathcal{P}_i}[1/2] = (1/2)^n$ and similarly $\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})] \leq \prod_i \mathbb{E}_{\mathcal{P}_i}[3/4] = (3/4)^n$.

By Markov's inequality, given that $\pi(\mathcal{G}) \geq 0$, we have $\mathbb{P}_{\mathcal{P}}[\pi(\mathcal{G}) \geq c] \leq \frac{\mathbb{E}_{\mathcal{P}}[\pi(\mathcal{G})]}{c} \leq \frac{(3/4)^n}{c}$. For $c = \frac{(3/4)^n}{\delta} \Rightarrow \mathbb{P}_{\mathcal{P}}[\pi(\mathcal{G}) \geq \frac{(3/4)^n}{\delta}] \leq \delta \Rightarrow \mathbb{P}_{\mathcal{P}}[\pi(\mathcal{G}) \leq \frac{(3/4)^n}{\delta}] \geq 1 - \delta$. □

C Negative Results

C.1 Counting the Number of Equilibria is NP-hard

Here we provide a proof that establishes NP-hardness of counting the number of Nash equilibria, and thus also of evaluating the log-likelihood function for our generative model. A #P-hardness proof was originally provided by Irfan and Ortiz [2011], here we present a related proof for completeness. The reduction is from the *set partition problem* for a specific instance of a single *non-absolutely-indifferent* player.

Recall the *set partition problem*: given a set of positive numbers $\{a_1, \dots, a_n\}$, $\mathrm{SetPartition}(\mathbf{a})$ answers "yes" if and only if it is possible to partition the numbers into two disjoint subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$, $\mathcal{S}_1 \cup \mathcal{S}_2 = \{1, \dots, n\}$ and $\sum_{i \in \mathcal{S}_1} a_i - \sum_{i \in \mathcal{S}_2} a_i = 0$; otherwise it answers "no". The set partition problem is equivalent to the *subset sum problem*, in which given a set of positive numbers $\{a_1, \dots, a_n\}$ and a target sum $c > 0$, $\mathrm{SubSetSum}(\mathbf{a}, c)$ answers "yes" if and only if there is a subset $\mathcal{S} \subset \{1, \dots, n\}$ such that $\sum_{i \in \mathcal{S}} a_i = c$; otherwise it answers "no". The equivalence between set partition and subset sum follows from $\mathrm{SetPartition}(\mathbf{a}) = \mathrm{SubSetSum}(\mathbf{a}, \frac{1}{2}\sum_i a_i)$.

For clarity of exposition, we drop the subindices in the following lemma. Let $\mathbf{w} \equiv \mathbf{w}_{i,-i} \in \mathbb{R}^{n-1}$ and $b \equiv b_i \in \mathbb{R}$.

Lemma 1. *The problem of counting Nash equilibria considered in Claim ii of Lemma 17 reduces to the set partition problem. More specifically, given* $(\forall i)\ w_i > 0, b = 0$*, answering whether* $\sum_{\mathbf{x}} 1[\mathbf{w}^{\mathrm{T}}\mathbf{x} - b = 0] > 0$ *is equivalent to answering* $\mathrm{SetPartition}(\mathbf{w})$.

Proof. Let $\mathcal{S}_1(\mathbf{x}) = \{i | x_i = +1\}$ and $\mathcal{S}_2(\mathbf{x}) = \{i | x_i = -1\}$. We can rewrite $\sum_{\mathbf{x}} 1[\mathbf{w}^{\mathrm{T}}\mathbf{x} - b = 0]$ as a sum of *set partition* conditions, i.e. $\sum_{\mathbf{x}} 1[\sum_{i \in \mathcal{S}_1(\mathbf{x})} w_i - \sum_{i \in \mathcal{S}_2(\mathbf{x})} w_i = 0]$. Therefore, if no tuple $\mathbf{x}$ fulfills the condition, the sum is zero and $\mathrm{SetPartition}(\mathbf{w})$ answers "no". On the other hand, if at least one tuple $\mathbf{x}$ fulfills the condition, the sum is greater than zero and $\mathrm{SetPartition}(\mathbf{w})$ answers "yes". □

C.2 Computing the Pseudo-Likelihood is NP-hard

We show that evaluating the pseudo-likelihood function for our generative model is NP-hard. First, consider a non-trivial influence game $\mathcal{G}$ in which eq.(1) simplifies to $p_{(\mathcal{G},q)}(\mathbf{x}) = q\frac{1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})]}{|\mathcal{NE}(\mathcal{G})|} + (1-q)\frac{1[\mathbf{x} \notin \mathcal{NE}(\mathcal{G})]}{2^n - |\mathcal{NE}(\mathcal{G})|}$. Furthermore, assume the game $\mathcal{G} = (\mathbf{W}, \mathbf{b})$ has a single *non-absolutely-indifferent* player i and *absolutely-indifferent* players $\forall j \neq i$. Let $f_i(\mathbf{x}_{-i}) \equiv {\mathbf{w}_{i,-i}}^{\mathrm{T}}\mathbf{x}_{-i} - b_i$. By Claim i of Lemma 17, we have $1[\mathbf{x} \in \mathcal{NE}(\mathcal{G})] = 1[x_i f_i(\mathbf{x}_{-i}) \geq 0]$ and therefore $p_{(\mathcal{G},q)}(\mathbf{x}) = q\frac{1[x_i f_i(\mathbf{x}_{-i}) \geq 0]}{|\mathcal{NE}(\mathcal{G})|} + (1-q)\frac{1 - 1[x_i f_i(\mathbf{x}_{-i}) \geq 0]}{2^n - |\mathcal{NE}(\mathcal{G})|}$. Finally, by Lemma 1 computing $|\mathcal{NE}(\mathcal{G})|$ is NP-hard even for this specific instance of a single *non-absolutely-indifferent* player.

C.3 Counting the Number of Equilibria is not (Lipschitz) Continuous

We show that small changes in the parameters $\mathcal{G} = (\mathbf{W}, \mathbf{b})$ can produce big changes in $|\mathcal{NE}(\mathcal{G})|$. For instance, consider two games $\mathcal{G}_k = (\mathbf{W}_k, \mathbf{b}_k)$, where $\mathbf{W}_1 = \mathbf{0}, \mathbf{b}_1 = \mathbf{0}, |\mathcal{NE}(\mathcal{G}_1)| = 2^n$ and $\mathbf{W}_2 = \varepsilon(\mathbf{1}\mathbf{1}^{\mathrm{T}} - \mathbf{I}), \mathbf{b}_2 = \mathbf{0}, |\mathcal{NE}(\mathcal{G}_2)| = 2$ for $\varepsilon > 0$. For $\varepsilon \to 0$, any ℓ_p-norm $\|\mathbf{W}_1 - \mathbf{W}_2\|_p \to 0$ but $|\mathcal{NE}(\mathcal{G}_1)| - |\mathcal{NE}(\mathcal{G}_2)| = 2^n - 2$ remains constant.

C.4 The Log-Partition Function of an Ising Model is a Trivial Bound for Counting the Number of Equilibria

Let $f_i(\mathbf{x}_{-i}) \equiv {\mathbf{w}_{i,-i}}^{\mathrm{T}}\mathbf{x}_{-i} - b_i$, $|\mathcal{NE}(\mathcal{G})| = \sum_{\mathbf{x}} \prod_i 1[x_i f_i(\mathbf{x}_{-i}) \geq 0] \leq \sum_{\mathbf{x}} \prod_i e^{x_i f_i(\mathbf{x}_{-i})} = \sum_{\mathbf{x}} e^{\mathbf{x}^{\mathrm{T}}\mathbf{W}\mathbf{x} - \mathbf{b}^{\mathrm{T}}\mathbf{x}} = \mathcal{Z}(\frac{1}{2}(\mathbf{W} + \mathbf{W}^{\mathrm{T}}), \mathbf{b})$, where $\mathcal{Z}$ denotes the partition function of an Ising model. Given convexity of $\mathcal{Z}$ [Koller and Friedman, 2009] and that the gradient vanishes at $\mathbf{W} = \mathbf{0}, \mathbf{b} = \mathbf{0}$, we know that $\mathcal{Z}(\frac{1}{2}(\mathbf{W} + \mathbf{W}^{\mathrm{T}}), \mathbf{b}) \geq 2^n$, which is the maximum $|\mathcal{NE}(\mathcal{G})|$.

D Simultaneous SVM Dual

First, note that the primal problem in eq.(9) is equivalent to a linear program since we can set $\mathbf{W} = \mathbf{W}^+ - \mathbf{W}^-$, $\|\mathbf{W}\|_1 = \sum_{ij} w_{ij}^+ + w_{ij}^-$ and add the constraints $\mathbf{W}^+ \geq \mathbf{0}$ and $\mathbf{W}^- \geq \mathbf{0}$.

By Lagrangian duality, the dual of the problem in eq.(9) is the following linear program:

$$\begin{array}{ll} \max_{\boldsymbol{\alpha}} & \sum_{li} \alpha_{li} \\ \text{s.t.} & (\forall i)\ \| \sum_l \alpha_{li} x_i^{(l)} \mathbf{x}_{-i}^{(l)} \|_\infty \leq \rho \quad , \quad (\forall l, i)\ \alpha_{li} \geq 0 \\ & (\forall i)\ \sum_l \alpha_{li} x_i^{(l)} = 0 \quad , \qquad\qquad (\forall l)\ \sum_i \alpha_{li} \leq \frac{1}{m} \end{array} \tag{1}$$

Furthermore, strong duality holds in this case. Note that eq.(1) is equivalent to a linear program since we can transform the constraint $\|\mathbf{c}\|_\infty \leq \rho$ into $-\rho\mathbf{1} \leq \mathbf{c} \leq \rho\mathbf{1}$.

E Simultaneous Logistic Loss

Given that any loss $\ell(z)$ is a decreasing function, the following identity holds $\max_i \ell(z_i) = \ell(\min_i z_i)$. Hence, we can either upper-bound the max function by the logsumexp function or lower-bound the min function by a negative logsumexp. We chose the latter option for the logistic loss for the following reasons: Claim i of the following technical lemma shows that lower-bounding min generates a loss that is strictly less than upper-bounding max. Claim ii shows that lower-bounding min generates a loss that is strictly less than independently penalizing each player. Claim iii shows that there are some cases in which upper-bounding max generates a loss that is strictly greater than independently penalizing each player.

Lemma 2. *For the logistic loss $\ell(z) = \log(1 + e^{-z})$ and a set of $n > 1$ numbers $\{z_1, \dots, z_n\}$:*

$$\begin{array}{l} \text{i.}\ \ (\forall z_1, \dots, z_n)\ \max_i \ell(z_i) \leq \ell\left(-\log \sum_i e^{-z_i}\right) \\ \qquad < \log \sum_i e^{\ell(z_i)} \leq \max_i \ell(z_i) + \log n \\ \text{ii.}\ \ (\forall z_1, \dots, z_n)\ \ell\left(-\log \sum_i e^{-z_i}\right) < \sum_i \ell(z_i) \\ \text{iii.}\ (\exists z_1, \dots, z_n)\ \log \sum_i e^{\ell(z_i)} > \sum_i \ell(z_i) \end{array} \tag{2}$$

Proof. Given a set of numbers $\{a_1, \dots, a_n\}$, the max function is bounded by the logsumexp function by $\max_i a_i \leq \log \sum_i e^{a_i} \leq \max_i a_i + \log n$ [Boyd and Vandenberghe, 2006]. Equivalently, the min function is bounded by $\min_i a_i - \log n \leq -\log \sum_i e^{-a_i} \leq \min_i a_i$.

These identities allow us to prove two inequalities in Claim i, i.e. $\max_i \ell(z_i) = \ell(\min_i z_i) \leq \ell\left(-\log \sum_i e^{-z_i}\right)$ and $\log \sum_i e^{\ell(z_i)} \leq \max_i \ell(z_i) + \log n$. To prove the remaining inequality $\ell\left(-\log \sum_i e^{-z_i}\right) < \log \sum_i e^{\ell(z_i)}$, note that for the logistic loss $\ell\left(-\log \sum_i e^{-z_i}\right) = \log(1 + \sum_i e^{-z_i})$ and $\log \sum_i e^{\ell(z_i)} = \log(n + \sum_i e^{-z_i})$. Since $n > 1$, strict inequality holds.

To prove Claim ii, we need to show that $\ell\left(-\log \sum_i e^{-z_i}\right) = \log(1+\sum_i e^{-z_i}) < \sum_i \ell(z_i) = \sum_i \log(1 + e^{-z_i})$. This is equivalent to $1 + \sum_i e^{-z_i} < \prod_i (1 + e^{-z_i}) = \sum_{\mathbf{c} \in \{0,1\}^n} e^{-\mathbf{c}^{\mathrm{T}}\mathbf{z}} = 1 + \sum_i e^{-z_i} + \sum_{\mathbf{c} \in \{0,1\}^n, \mathbf{1}^{\mathrm{T}}\mathbf{c} > 1} e^{-\mathbf{c}^{\mathrm{T}}\mathbf{z}}$. Finally, we have $\sum_{\mathbf{c} \in \{0,1\}^n, \mathbf{1}^{\mathrm{T}}\mathbf{c} > 1} e^{-\mathbf{c}^{\mathrm{T}}\mathbf{z}} > 0$ because the exponential function is strictly positive.

To prove Claim iii, it suffices to find set of numbers $\{z_1, \dots, z_n\}$ for which $\log \sum_i e^{\ell(z_i)} = \log(n + \sum_i e^{-z_i}) > \sum_i \ell(z_i) = \sum_i \log(1 + e^{-z_i})$. This is equivalent to $n + \sum_i e^{-z_i} > \prod_i (1 + e^{-z_i})$. By setting $(\forall i)\ z_i = \log n$, we reduce the claim we want to prove to $n + 1 > (1 + \frac{1}{n})^n$. Strict inequality holds for $n > 1$. Furthermore, note that $\lim_{n \to +\infty} (1 + \frac{1}{n})^n = e$. □

F Extended Version of Experimental Results

For learning influence games we used our convex loss methods: independent and simultaneous SVM and logistic regression. Additionally, we used the (super-exponential) exhaustive search method only for $n \leq 4$. As a baseline, we used the sigmoidal maximum likelihood (NP-hard) only for $n \leq 15$ as well as the sigmoidal maximum empirical proportion of equilibria. Regarding the parameters α and β our sigmoidal function in eq.(7), we found experimentally that $\alpha = 0.1$ and $\beta = 0.001$ achieved the best results.

We compare learning influence games to learning Ising models. For $n \leq 15$ players, we perform exact ℓ_1-regularized maximum likelihood estimation by using the FOBOS algorithm [Duchi and Singer, 2009a,b] and exact gradients of the log-likelihood of the Ising model. Since the computation of the exact gradient at each step is NP-hard, we used this method only for $n \leq 15$. For $n > 15$ players, we use the Höfling-Tibshirani

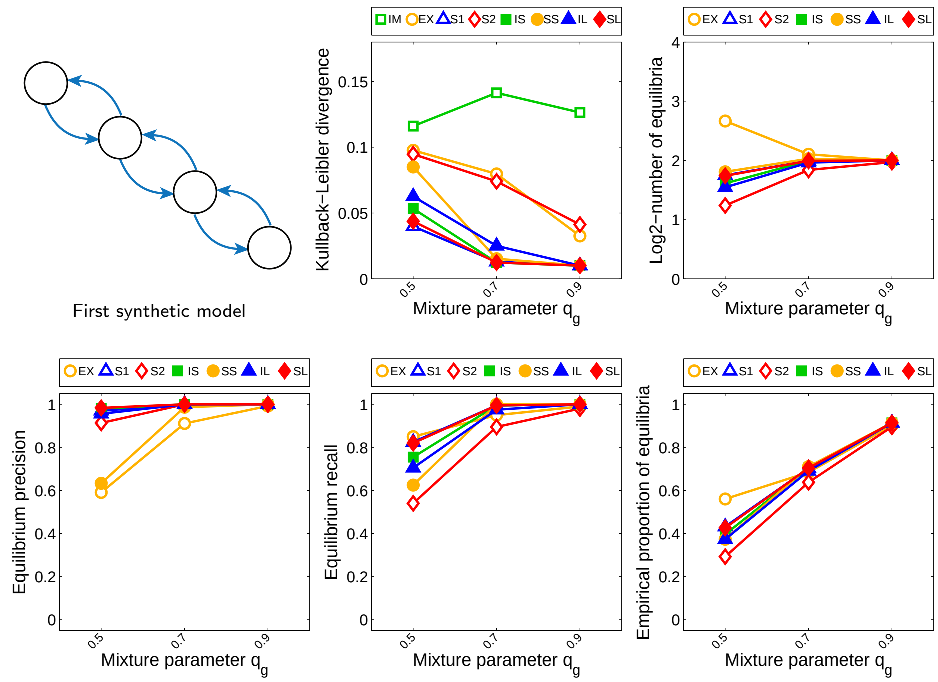

Figure 2: Closeness of the recovered models to the ground truth synthetic model for different mixture parameters q_g. Our convex loss methods (IS,SS: independent and simultaneous SVM, IL,SL: independent and simultaneous logistic regression) and sigmoidal maximum likelihood (S1) have lower KL than exhaustive search (EX), sigmoidal maximum empirical proportion of equilibria (S2) and Ising models (IM). For all methods, the recovery of equilibria is perfect for $q_g = 0.9$ (number of equilibria equal to the ground truth, equilibrium precision and recall equal to 1) and the empirical proportion of equilibria resembles the mixture parameter of the ground truth q_g.

method [Höfling and Tibshirani, 2009], which uses a sequence of first-order approximations of the exact log-likelihood. We also used a two-step algorithm, by first learning the structure by ℓ_1-regularized logistic regression [Wainwright et al., 2006] and then using the FOBOS algorithm [Duchi and Singer, 2009a,b] with belief propagation for gradient approximation. We did not find a statistically significant difference between the test log-likelihood of both algorithms and therefore we only report the latter.

Our experimental setup is as follows: after learning a model for different values of the regularization parameter ρ in a training set, we select the value of ρ that maximizes the log-likelihood in a validation set, and report statistics in a test set. For synthetic experiments, we report the Kullback-Leibler (KL) divergence, average precision (one minus the fraction of falsely included equilibria), average recall (one minus the fraction of falsely excluded equilibria) in order to measure the closeness of the recovered models to the ground truth. For real-world experiments, we report the log-likelihood. In both synthetic and real-world experiments, we report the number of equilibria and the empirical proportion of equilibria.

We first test the ability of the proposed methods to recover the ground truth structure from data. We use a small first synthetic model in order to compare with the (super-exponential) exhaustive search method. The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ has $n = 4$ players and 4 Nash equilibria (i.e. $\pi(\mathcal{G}_g)$=0.25), $\mathbf{W}_g$ was set according to Figure 2 (the weight of each edge was set to $+1$) and $\mathbf{b}_g = \mathbf{0}$. The mixture parameter of the ground truth q_g was set to 0.5,0.7,0.9. For each of 50 repetitions, we generated a training, a validation and a test set of 50 samples each. Figure 2 shows that our convex loss methods and sigmoidal maximum likelihood outperform (lower KL) exhaustive search, sigmoidal maximum empirical proportion of equilibria and Ising models. Note that the exhaustive search method which performs exact maximum likelihood suffers from over-fitting and consequently does not produce the lowest KL. From all convex loss methods, simultaneous logistic regression achieves the lowest KL. For all methods, the recovery of equilibria is perfect for $q_g = 0.9$ (number of equilibria equal to the ground truth, equilibrium precision and recall equal to 1). Additionally, the empirical proportion of equilibria resembles the mixture parameter of the ground truth q_g.

Next, we use a relatively larger second synthetic model with more complex interactions. We still keep the model small enough in order to compare with the (NP-hard) sigmoidal maximum likelihood method.

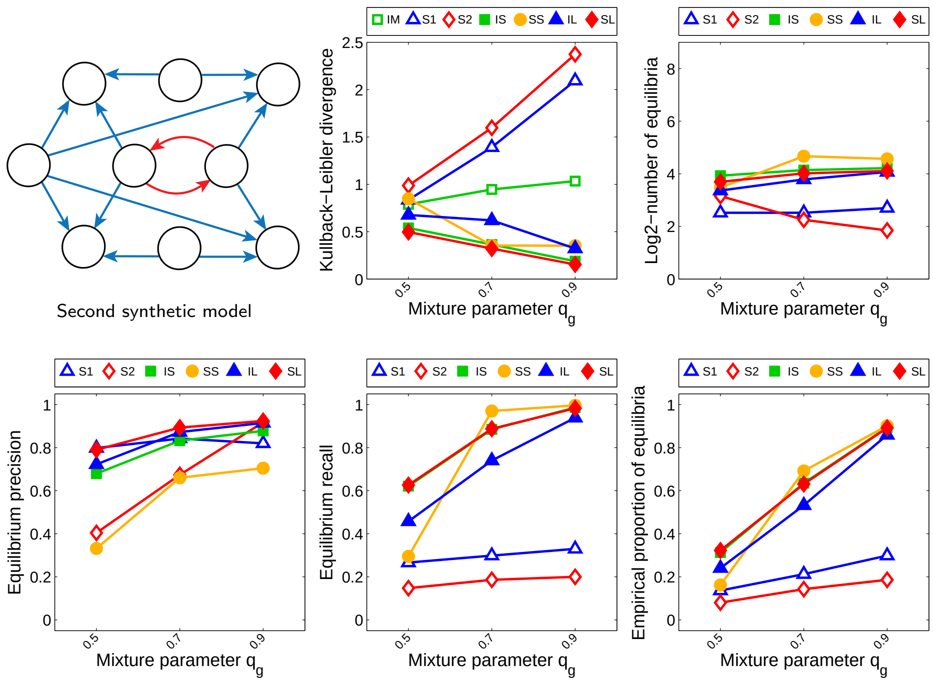

Figure 3: Closeness of the recovered models to the ground truth synthetic model for different mixture parameters q_g. Our convex loss methods (IS,SS: independent and simultaneous SVM, IL,SL: independent and simultaneous logistic regression) have lower KL than sigmoidal maximum likelihood (S1), sigmoidal maximum empirical proportion of equilibria (S2) and Ising models (IM). For convex loss methods, the equilibrium recovery is better than the remaining methods (number of equilibria equal to the ground truth, higher equilibrium precision and recall) and the empirical proportion of equilibria resembles the mixture parameter of the ground truth q_g.

The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ has $n = 9$ players and 16 Nash equilibria (i.e. $\pi(\mathcal{G}_g)$=0.03125), $\mathbf{W}_g$ was set according to Figure 3 (the weight of each blue and red edge was set to $+1$ and -1 respectively) and $\mathbf{b}_g = \mathbf{0}$. The mixture parameter of the ground truth q_g was set to 0.5,0.7,0.9. For each of 50 repetitions, we generated a training, a validation and a test set of 50 samples each. Figure 3 shows that our convex loss methods outperform (lower KL) sigmoidal methods and Ising models. From all convex loss methods, simultaneous logistic regression achieves the lowest KL. For convex loss methods, the equilibrium recovery is better than the remaining methods (number of equilibria equal to the ground truth, higher equilibrium precision and recall). Additionally, the empirical proportion of equilibria resembles the mixture parameter of the ground truth q_g.

In the next experiment, we show that the performance of convex loss minimization improves as the number of samples increases. We used random graphs with slightly more variables and varying number of samples (10,30,100,300). The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ contains $n = 20$ players. For each of 20 repetitions, we generate edges in the ground truth model $\mathbf{W}_g$ with a required density (either 0.2,0.5,0.8). For simplicity, the weight of each edge is set to $+1$ with probability $P(+1)$ and to -1 with probability $1-P(+1)$. Hence, the Nash equilibria of the generated games does not depend on the magnitude of the weights, just on their sign. We set the bias $\mathbf{b}_g = \mathbf{0}$ and the mixture parameter of the ground truth $q_g = 0.7$. We then generated a training and a validation set with the same number of samples. Figure 4 shows that our convex loss methods outperform (lower KL) sigmoidal maximum empirical proportion of equilibria and Ising models (except for the synthetic model with high true proportion of equilibria: density 0.8, $P(+1) = 0$, NE>1000). The results are remarkably better when the number of equilibria in the ground truth model is small (e.g. for NE<20). From all convex loss methods, simultaneous logistic regression achieves the lowest KL.

In the next experiment, we evaluate two effects in our approximation methods. First, we evaluate the impact of removing the true proportion of equilibria from our objective function, i.e. the use of maximum empirical proportion of equilibria instead of maximum likelihood. Second, we evaluate the impact of using convex losses instead of a sigmoidal approximation of the 0/1 loss. We used random graphs with varying number of players and 50 samples. The ground truth model $\mathcal{G}_g = (\mathbf{W}_g, \mathbf{b}_g)$ contains $n = 4, 6, 8, 10, 12$ players.

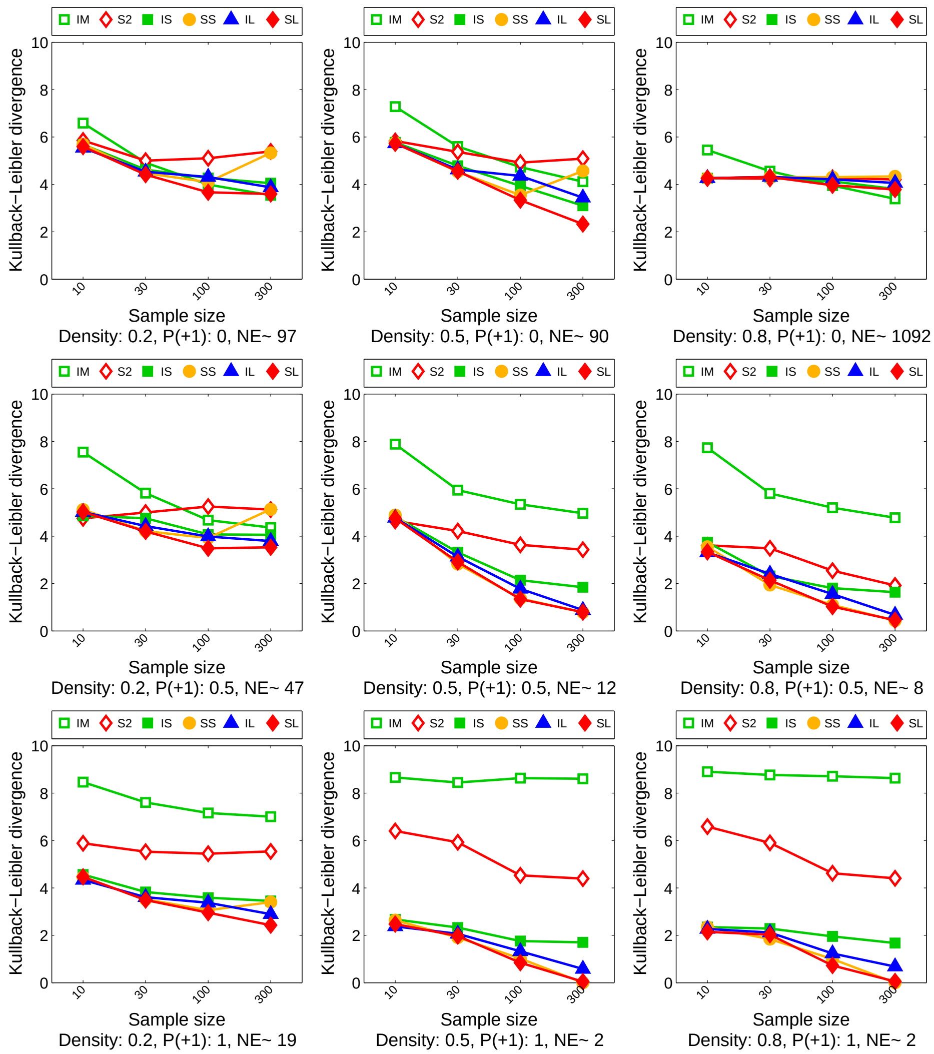

Figure 4: KL divergence between the recovered models and the ground truth for datasets of different number of samples. Each chart shows the density of the ground truth, probability $P(+1)$ that an edge has weight $+1$, and average number of equilibria (NE). Our convex loss methods (IS,SS: independent and simultaneous SVM, IL,SL: independent and simultaneous logistic regression) have lower KL than sigmoidal maximum empirical proportion of equilibria (S2) and Ising models (IM). The results are remarkably better when the number of equilibria in the ground truth model is small (e.g. for NE$<$ 20).

For each of 20 repetitions, we generate edges in the ground truth model $\mathbf{W}_g$ with a required density (either 0.2,0.5,0.8). As in the previous experiment, the weight of each edge is set to $+1$ with probability $P(+1)$ and to -1 with probability $1 - P(+1)$. We set the bias $\mathbf{b}_g = \mathbf{0}$ and the mixture parameter of the ground truth $q_g = 0.7$. We then generated a training and a validation set with the same number of samples. Figure 5 shows that in general, convex loss methods outperform (lower KL) sigmoidal maximum empirical proportion of equilibria, and the latter one outperforms sigmoidal maximum likelihood. A different effect is observed for mild (0.5) to high (0.8) density and $P(+1) = 1$ in which the sigmoidal maximum likelihood obtains the lowest KL. In a closer inspection, we found that the ground truth games usually have only 2 equilibria: $(+1, \ldots, +1)$ and $(-1, \ldots, -1)$, which seems to present a challenge for convex loss methods. It seems that for these specific cases, removing the true proportion of equilibria from the objective function negatively impacts the estimation process, but note that sigmoidal maximum likelihood is not computationally feasible for $n > 15$.

We used the U.S. congressional voting records in order to measure the generalization performance of convex loss minimization in a real-world dataset. The dataset is publicly available at `http://www.senate.gov/`. We used the first session of the 104th congress (Jan 1995 to Jan 1996, 613 votes), the first session of the 107th

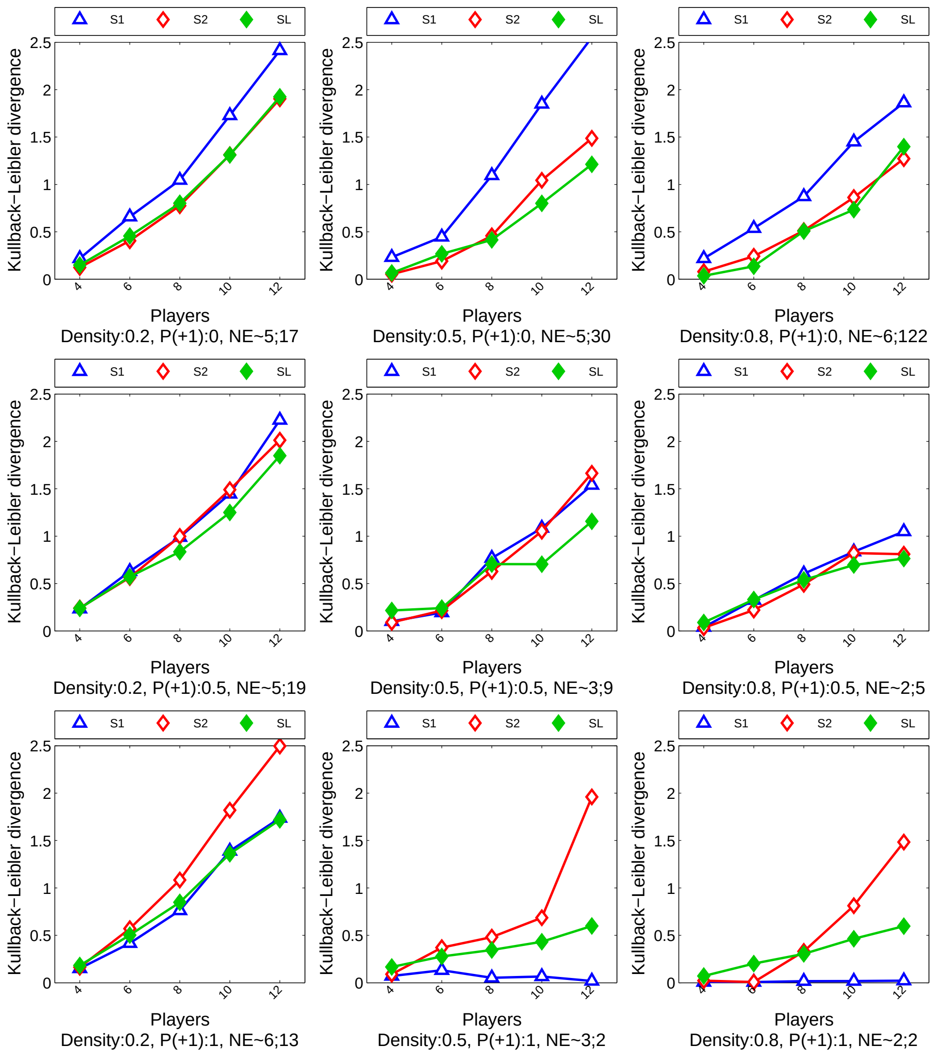

Figure 5: KL divergence between the recovered models and the ground truth for datasets of different number of players. Each chart shows the density of the ground truth, probability $P(+1)$ that an edge has weight $+1$, and average number of equilibria (NE) for $n = 2;n = 14$. In general, simultaneous logistic regression (SL) has lower KL than sigmoidal maximum empirical proportion of equilibria (S2), and the latter one has lower KL than sigmoidal maximum likelihood (S1). Other convex losses behave the same as simultaneous logistic regression (omitted for clarity of presentation).

congress (Jan 2001 to Dec 2001, 380 votes) and the second session of the 110th congress (Jan 2008 to Jan 2009, 215 votes). Following on other researchers who have experimented with this data set (e.g. Banerjee et al. [2008]), abstentions were replaced with negative votes. Since reporting the log-likelihood requires computing the number of equilibria (which is NP-hard), we selected only 20 senators by stratified random sampling. We randomly split the data into three parts. We performed six repetitions by making each third of the data take turns as training, validation and testing sets. Figure 6 shows that our convex loss methods outperform (higher log-likelihood) sigmoidal maximum empirical proportion of equilibria and Ising models. From all convex loss methods, simultaneous logistic regression achieves the lowest KL. For all methods, the number of equilibria (and so the true proportion of equilibria) is low.

We apply convex loss minimization to larger problems, by learning structures of games from all 100 senators. Figure 7 shows that simultaneous logistic regression produce structures that are sparser than its independent counterpart. The simultaneous method better elicits the bipartisan structure of the congress. We define the influence of player j to all other players as $\sum_i |w_{ij}|$ after normalizing all weights, i.e. for each player i we divide $(\mathbf{w}_{i,-i}, b_i)$ by $\|\mathbf{w}_{i,-i}\|_1 + |b_i|$. Note that Jeffords and Clinton are one of the 5 most directly-influential as well as 5 least directly-influenceable (high bias) senators, in the 107th and 110th

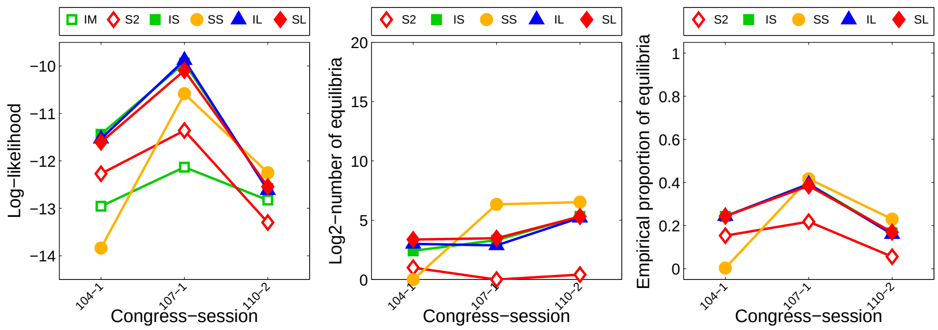

Figure 6: Statistics for games learnt from 20 senators from the first session of the 104th congress, first session of the 107th congress and second session of the 110th congress. The log-likelihood of our convex loss methods (IS,SS: independent and simultaneous SVM, IL,SL: independent and simultaneous logistic regression) is higher than sigmoidal maximum empirical proportion of equilibria (S2) and Ising models (IM). For all methods, the number of equilibria (and so the true proportion of equilibria) is low.

congress respectively. McCain and Feingold are both in the list of 5 most directly-influential senators in the 104th and 107th congress. McCain appears again in the list of 5 least influenciable senators in the 110th congress.

We test the hypothesis that influence between senators of the same party are stronger than senators of different party. We learn structures of games from all 100 senators from the 101th congress to the 111th congress (Jan 1989 to Dec 2010). The number of votes casted for each session were average: 337, minimum: 215, maximum: 613. Figure 8 validates our hypothesis and more interestingly, it shows that influence between different parties is decreasing over time. Note that the influence from Obama to Republicans increased in the last sessions, while McCain's influence to Republicans decreased.

G Discussion

There has been a significant amount of work for learning the structure of *probabilistic* graphical models from data. We mention only a few references that follow a maximum likelihood approach for Markov random fields [Lee et al., 2006], bounded tree-width distributions [Chow and Liu, 1968, Srebro, 2001], Ising models [Wainwright et al., 2006, Banerjee et al., 2008, Höfling and Tibshirani, 2009], Gaussian graphical models [Banerjee et al., 2006], Bayesian networks [Guo and Schuurmans, 2006, Schmidt et al., 2007] and directed cyclic graphs [Schmidt and Murphy, 2009].

Our approach learns the structure and parameters of games by maximum likelihood estimation on a related probabilistic model. Our probabilistic model does not fit into any of the types described above. Although a (directed) graphical game has a directed cyclic graph, there is a semantic difference with respect to graphical models. Structure in a graphical model implies a factorization of the probabilistic model. In a graphical game, the graph structure implies *strategic* dependence between players, and has no immediate probabilistic implication. Furthermore, our general model differs from Schmidt and Murphy [2009] since our generative model does not decompose as a multiplication of potential functions.

It is important to point out that our work is not in competition with the work in probabilistic graphical models, e.g. Ising models. Our goal is to learn the structure and parameters of games from data, and for this end, we propose a probabilistic model that is inspired by the concept of equilibrium in game theory. While we illustrate the benefit of our model in the U.S. congressional voting records, we believe that each model has its own benefits. If the practitioner "believes" that the data at hand is generated by a class of models, then the interpretation of the learnt model allows obtaining insight of the problem at hand. Note that none of the existing models (including ours) can be validated as the ground truth model that generated the real-world data, or as being more or less "realistic" with respect to other model. While generalization in unseen data is a very important measurement, a model with better generalization is not the "ground truth model" of the real-world data at hand. Finally, while our model is simple, it is well founded and we show that it is far from being computationally trivial. Therefore, we believe it has its own right to be analyzed.

The special class of graphical games considered here is related to the well-known *linear threshold model*

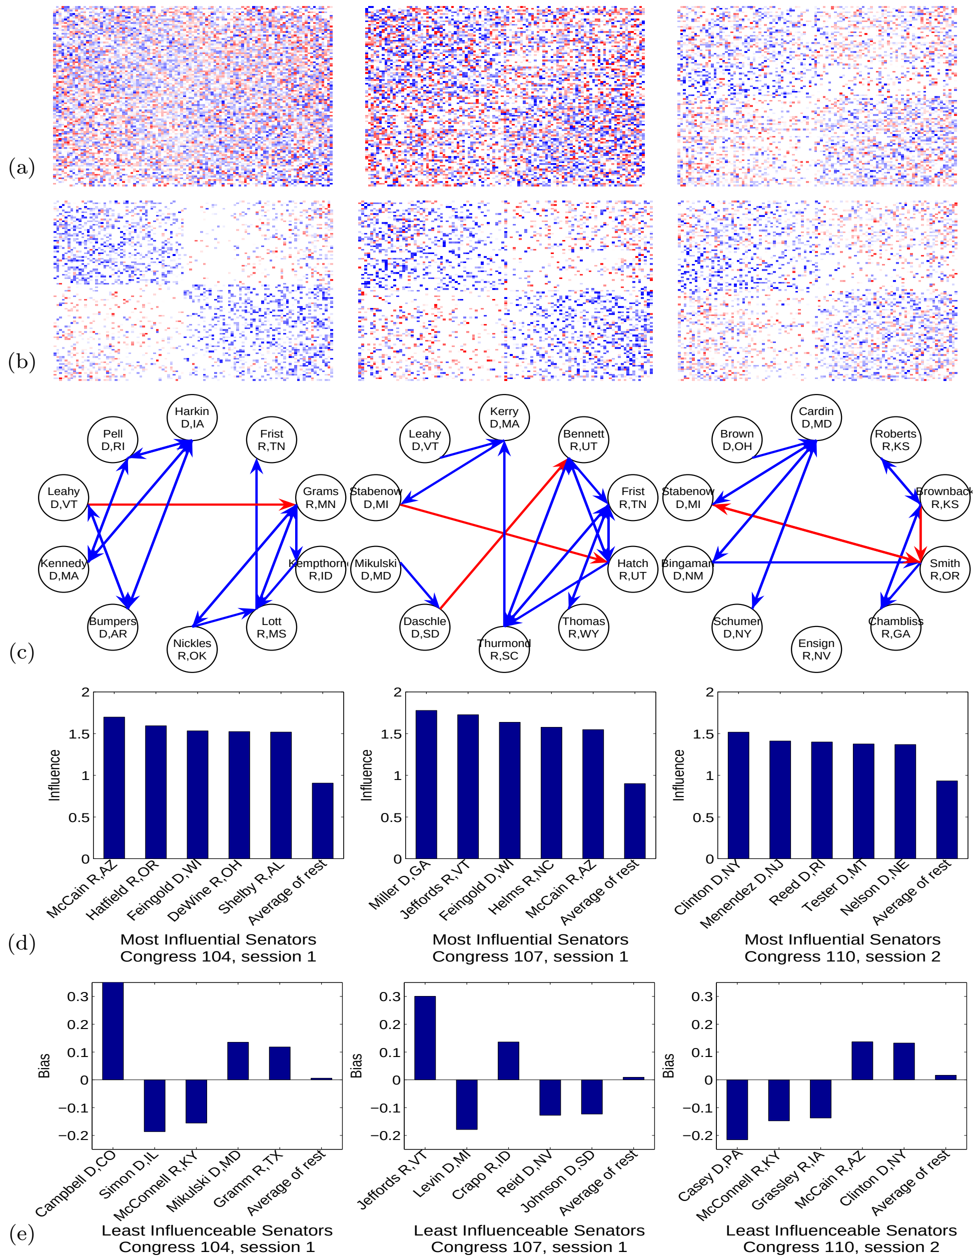

Figure 7: Matrix of influence weights for games learnt from all 100 senators, from the first session of the 104th congress (left), first session of the 107th congress (center) and second session of the 110th congress (right), by using our independent (a) and simultaneous (b) logistic regression methods. A row represents how every other senator influence the senator in such row. Positive influences are shown in blue, negative influences are shown in red. Democrats are shown in the top/left corner, while Republicans are shown in the bottom/right corner. Note that simultaneous method produce structures that are sparser than its independent counterpart. Partial view of the graph for simultaneous logistic regression (c). Most directly-influential (d) and least directly-influenceable (e) senators. Regularization parameter $\rho = 0.0006$.

(LTM) in sociology [Granovetter, 1978], recently very popular within the social network and theoretical computer science community [Kleinberg, 2007]. LTMs are usually studied as the basis for some kind of diffusion process. A typical problem is the identification of most influential individuals in a social network. An LTM is not in itself a game-theoretic model and, in fact, Granovetter himself argues against this view in the context of the setting and the type of questions in which he was most interested [Granovetter, 1978]. To the best of our knowledge, subsequent work on LTMs has not taken a strictly game-theoretic view either. Our model is also related to a particular model of *discrete choice with social interactions* in econometrics

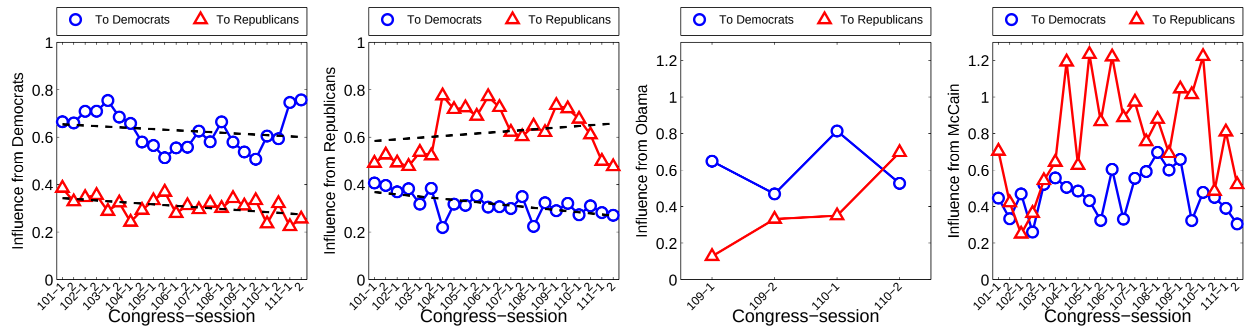

Figure 8: Direct influence between parties and influences from Obama and McCain. Games were learnt from all 100 senators from the 101th congress (Jan 1989) to the 111th congress (Dec 2010) by using our simultaneous logistic regression method. Direct influence between senators of the same party are stronger than senators of different party, which is also decreasing over time. In the last sessions, influence from Obama to Republicans increased, and influence from McCain to both parties decreased. Regularization parameter $\rho = 0.0006$.

(see, e.g. Brock and Durlauf [2001]). The main difference is that we take a strictly non-cooperative game-theoretic approach within the classical "static"/one-shot game framework and do not use a *random utility model*. In addition, we do not make the assumption of *rational expectations*, which is equivalent to assuming that all players use exactly the same mixed strategy. As an aside note, regarding learning of information diffusion models over social networks, [Saito et al., 2010] considers a dynamic (continuous time) LTM that has only positive influence weights and a randomly generated threshold value.

There is still quite a bit of debate as to the appropriateness of game-theoretic equilibrium concepts to model individual human behavior in a social context. Camerer's book on behavioral game theory [Camerer, 2003] addresses some of the issues. We point out that there is a broader view of behavioral data, beyond those generated by individual human behavior (e.g. institutions such as nations and industries, or engineered systems such as autonomous-response devices in residential or commercial properties that are programmed to control electricity usage based on user preferences). Our interpretation of Camerer's position is not that Nash equilibiria is universally a bad predictor but that it is not *consistently* the best, for reasons that are still not well understood. This point is best illustrated in Chapter 3, Figure 3.1 of Camerer [2003]. *Quantal response equilibria (QRE)* has been proposed as an alternative to Nash in the context of behavioral game theory. Models based on QRE have been shown superior during *initial play* in some experimental settings, but most experimental work assume that the game's payoff matrices are *known* and only the "precision parameter" is estimated, e.g. Wright and Leyton-Brown [2010]. Finally, most of the human-subject experiments in behavioral game theory involve only a handful of players, and the scalability of those results to games with *many* players is unclear.

In this work we considered pure-strategy Nash equilibria only. Note that the universality of mixed-strategy Nash equilibria does not diminish the importance of pure-strategy equilibria in game theory. Indeed, a debate still exist within the game theory community as to the justification for randomization, specially in human contexts. We decided to ignore mixed-strategies due to the significant added complexity. Note that we learn exclusively from observed joint-actions, and therefore we cannot assume knowledge of the internal mixed-strategies of players. We could generalize our model to allow for mixed-strategies by defining a process in which a joint mixed strategy $\mathcal{P}$ from the set of mixed-strategy Nash equilibrium (or its complement) is drawn according to some distribution, then a (pure-strategy) realization $\mathbf{x}$ is drawn from $\mathcal{P}$ that would correspond to the observed joint-actions.

In this paper we considered a "global" noise process, which is governed with a probability q of selecting an equilibrium. Potentially better and more natural "local" noise processes are possible, at the expense of producing a significantly more complex generative model than the one considered in this paper. For instance, we could use a noise process that is formed of many independent, individual noise processes, one for each player. As an example, consider a the generative model in which we first select an equilibrium $\mathbf{x}$ of the game and then each player i, independently, acts according to x_i with probability q_i and switches its action with probability $1 - q_i$. The problem with such a model is that it leads to a significantly more complex expression for the generative model and thus likelihood functions. This is in contrast to the simplicity afforded us by the generative model with a more global noise process defined above.

References

O. Aichholzer and F. Aurenhammer. Classifying Hyperplanes in Hypercubes. *SIAM Journal on Discrete Mathematics*, 9(2):225–232, 1996.

O. Banerjee, L. El Ghaoui, and A. d'Aspremont. Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data. *JMLR*, 2008.

O. Banerjee, L. El Ghaoui, A. d'Aspremont, and G. Natsoulis. Convex Optimization Techniques for Fitting Sparse Gaussian Graphical Models. *ICML*, 2006.

S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2006.

W. Brock and S. Durlauf. Discrete Choice with Social Interactions. *The Review of Economic Studies*, 68(2): 235–260, 2001.

C. Camerer. *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton University Press, 2003.

C. Chow and C. Liu. Approximating Discrete Probability Distributions with Dependence Trees. *IEEE Transactions on Information Theory*, 1968.

J. Duchi and Y. Singer. Efficient Learning using Forward-Backward Splitting. *NIPS*, 2009a.

J. Duchi and Y. Singer. Efficient Online and Batch Learning using Forward Backward Splitting. *JMLR*, 2009b.

M. Granovetter. Threshold Models of Collective Behavior. *The American Journal of Sociology*, 83(6): 1420–1443, 1978.

Y. Guo and D. Schuurmans. Convex Structure Learning for Bayesian Networks: Polynomial Feature Selection and Approximate Ordering. *UAI*, 2006.

H. Höfling and R. Tibshirani. Estimation of Sparse Binary Pairwise Markov Networks using Pseudo-likelihoods. *JMLR*, 2009.

M. Irfan and L. Ortiz. A Game-Theoretic Approach to Influence in Networks. *AAAI*, 2011.

J. Kleinberg. Cascading Behavior in Networks: Algorithmic and Economic Issues. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24, pages 613–632. Cambridge University Press, 2007.

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.

S. Lee, V. Ganapathi, and D. Koller. Efficient Structure Learning of Markov Networks Using ℓ_1-Regularization. *NIPS*, 2006.

K. Saito, M. Kimura, K. Ohara, and H. Motoda. Selecting Information Diffusion Models over Social Networks for Behavioral Analysis. *ECML*, 2010.

M. Schmidt and K. Murphy. Modeling Discrete Interventional Data using Directed Cyclic Graphical Models. *UAI*, 2009.

M. Schmidt, A. Niculescu-Mizil, and K. Murphy. Learning Graphical Model Structure Using ℓ_1-Regularization Paths. *AAAI*, 2007.

E. Sontag. VC Dimension of Neural Networks. In *Neural Networks and Machine Learning*, pages 69–95. Springer, 1998.

N. Srebro. Maximum Likelihood Bounded Tree-Width Markov Networks. *UAI*, 2001.

M. Wainwright, P. Ravikumar, and J. Lafferty. High dimensional Graphical Model Selection Using ℓ_1-Regularized Logistic Regression. *NIPS*, 2006.

J. Wright and K. Leyton-Brown. Beyond Equilibrium: Predicting Human Behavior in Normal Form Games. *AAAI*, 2010.