The Power of Context in Networks: Ideal Point Models with Social Interactions

Mohammad T. Irfan
Bowdoin College
Brunswick, Maine
mirfan@bowdoin.edu

Tucker Gordon*
Booz Allen Hamilton Inc.
McLean, Virginia
gordon_tucker@bah.com

ABSTRACT

Game theory has been widely used for modeling strategic behaviors in networked multiagent systems. However, the context within which these strategic behaviors take place has received limited attention. We present a model of strategic behavior in networks that incorporates the behavioral context. We focus on the contextual aspects of congressional voting. A senator's decision to vote *yea* or *nay* on a bill comes as a result of their ideologies, agendas, and their interactions with other senators. One salient model in political science is the *ideal point model*, which assigns each senator and each bill a number on the real line of political spectrum. These points then allow for prediction of future voting behavior. We extend the classical ideal point model with network-structured interactions among senators. In contrast to the ideal point model's prediction of individual voting behavior, we predict *joint voting behaviors* in a game-theoretic fashion. Our model also includes the characteristics of a bill. This allows it to outperform previous models that solely focus on the networked interactions among senators with no bill-specific parameters. We focus on two fundamental questions: learning the model using real-world data and computing *stable outcomes* of the model in order to predict joint voting behaviors. We demonstrate the effectiveness of our model through experiments using data from the 114th U.S. Congress.

KEYWORDS

computational game theory; social networks; machine learning; congressional voting; influence in networks

ACM Reference Format:
Mohammad T. Irfan and Tucker Gordon. 2018. The Power of Context in Networks: Ideal Point Models with Social Interactions. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

We consider a networked multiagent system, where agents democratically vote on agendas. An agent's vote depends on her social interaction with others as well as the agenda on which she is voting. Here, the agenda gives a context to the voting behaviors of the agents. Context being a domain-specific construct, we focus on voting behavior in Congress. Our framework can also be applied to other settings where there are (1) network-structured interactions and (2) issues on which agents adopt binary behaviors. [1]

There have been multidisciplinary quests into modeling and predicting how legislators vote. The actual votes are easy to find by looking at the *roll call data* – a record of each individual's vote on any given bill, resolution, etc. Roll call data consists of affirmative *yea*s, negative *nay*s, and abstentions. The challenges in modeling and predicting how legislators vote come in identifying patterns in data, constructing models of voting behavior, and using these models to make predictions on future votes. Many factors play a role in senators' voting behaviors: perhaps the bill lines up closely with their political agenda, or is popular amongst their constituents. Perhaps they feel pressure from their party or other legislators to confirm a nomination or to pass a bill. Despite such pressure, senators sometimes vote against the party line. One recent example is Senators Collins (R-ME) and Murkowski (R-AK) voting against the nomination of Betsy DeVos for Secretary of Education. They could have genuine concern about the nominee due to their legislative positions on the left-to-right political spectrum (which indirectly reflects their constituents' views) or the nominee's position on relevant issues. Another recent example is the failure of the health-care bill in Senate, where Senator McCain (R-AZ) gave a dramatic, last-minute *nay* vote. There are other examples where senators voted against the party line due to the bill not lining up with their particular legislative positions (e.g., pro-choice Republicans or anti-gun-control Democrats). Therefore, a model of legislative voting behavior should take into account a senator's legislative position, the characteristics of the bills, *and* the network of influence among the senators. We present the first such model here.

Introduced by Davis et al. [3], one extremely popular approach to modeling legislative behavior is the *ideal point model*. It seeks to give a numeric representation of a senator's political leaning based on how liberal or conservative they are. In this model, each senator i has an ideal point $p_i \in \mathbb{R}$. Each bill l is characterized by its *polarity* $a_l \in \mathbb{R}$ and its *popularity* $b_l \in \mathbb{R}$, analogous to *discrimination* and *difficulty* in item-response theory [8]. Typically, p_i and a_l are negative for liberal positions and positive for conservative, and the b_l parameter corresponds to the fraction of senators voting *yea*. The probability of senator i voting *yea* on a bill l is given by the following logistic function σ.

$$p(x_{i,l} = yea \mid p_i, a_l, b_l) = \sigma(p_i a_l + b_l). \tag{1}$$

There are numerous applications of this model. Canes-Wrone et al. [1], Jenkins [13], and Schickler [21] each use ideal points to

*TG worked on this research as an undergraduate student at Bowdoin College.

[1] We borrow the term *the power of context* from Malcolm Gladwell's popular book *the Tipping Point*. We do not, however, deal with diffusion or contagion here.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

model the House of Representatives at different times in the U.S. history. Clinton et al. [2] use them to model the U.S. House, Senate, and even Supreme Court. Ideal point models provide informative results (e.g. a definition of what it means to be "moderate" or "extreme"). This model has been extended to factor in the actual content of bills into voting decisions, adjusting lawmaker's preferences in accordance with the bill [6].

For the most part, ideal point models consider each legislator's vote as an individual decision with respect to the issue being voted on. All elements of social interactions are lumped into the single popularity parameter b_l in Equation (1). In real world, however, the voting behaviors of legislators are often interdependent (e.g., deal-making), and as a result each voting outcome can be thought of as a joint behavioral outcome. Such systems of interdependent behaviors naturally fall in the territory of game theory.

Game-theoretic models have been proposed in the literature to model voting behaviors in Senate [10–12]. In the *linear influence game* (LIG) model, the actions of players (i.e., senators) influence those of others in a network-structured fashion, and every player acts in a way to maximize their payoffs with respect to others. In the context of legislative bodies, this means that lawmakers vote in response to the votes of others. For example, a Democratic senator may vote *yea* on a bill that many Republican senators oppose, and may vote *nay* if a close colleague opposes it. Within this context, votes occur in a group setting and one legislator's vote is impossible to predict without taking into account the votes of others. The LIG model tries to capture *strategic behavior* in the sense that lawmakers influence each other's votes and that the votes are rationalized through their network-structured interactions.

While both the LIG and ideal point models are predictive models, each is limited in some capacity. The ideal point model is statistical and provides insight into the ideologies of individual legislators, but lacks any notion of strategic interactions. The LIG model is game-theoretic and captures the effects of social influence on votes, but do not take into account the characteristics of the issue being voted on. Furthermore, the statistical ideal point model gives individual predictions, while the the game-theoretic LIG model gives joint predictions. Previous studies to extend the ideal point model focused on the textual content of the bills [6]. Here, we extend the ideal point model in the direction of network-structured strategic interactions among the senators. We also consider the content of the bills via the subject tags. We are not aware of any study on ideal point model with social interactions.

Our Contributions. We present a game-theoretic model of strategic behavior in Congress that extends the classical ideal point model to incorporate the subject areas of the bills *and* the social interactions among the senators. We learn the model using data from the 114th U.S. Senate (2015–17). We show the hardness of computing stable outcomes (or pure-strategy Nash equilibria) of the model and use an effective heuristic for fast computation of stable outcomes. We also compare and contrast our model with the ideal point model [2, 3] as well as the state-of-the-art LIG model [10–12]. We show that our model outperforms both across the whole political spectrum of bills. We show how our model can be used to identify the most influential Senators in Congress in the context of a bill.

Notation	**Description**
$x_i \in \mathbf{x}$	Vote of senator i
$w_{i,j} \in \mathbf{w}$	Influence on senator i from j
t_i	Threshold of senator i
p_i	Ideal point of senator i
a_l	Polarity of bill l

Table 1: Notation used in our model

2 AN IDEAL POINT MODEL WITH SOCIAL INTERACTIONS

We present a binary-action (+1 for *yea* and -1 for *nay*) graphical game [14] model with parametric payoff functions of quadratic forms. We use a weighted-directed graph to represent the network structure among the legislative body of N senators. We use $\mathcal{N}_i$ for the set of neighbors of node i in this graph. Below, we present the senator-specific and bill-specific model parameters. We then describe how senators vote according to our model.

Senator-Specific Parameters. Each node i is a senator and has two parameters: a threshold $t_i \in \mathbb{R}$ representing how "stubborn" node i is and an ideal point $p_i \in \mathbb{R}$ representing the senator's ideology in the political spectrum. Each directed edge from j to i is weighted by the amount of influence $w_{i,j} \in \mathbb{R}$ that node j has on node i (we allow negative influence weights). The influence weights together form a matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, with $\text{diag}(\mathbf{W}) = 0$.

Bill-Specific Parameters. Each bill l being voted on is parameterized by its polarity $a_l \in \mathbb{R}$ representing the bill's position in the liberal (negative numbers) to conservative (positive) spectrum.

Senators' Choices of Actions. Given the above parameters, the quantity of interest is the vote $x_i \in \{+1, -1\}$ of a senator i on a bill l. [2] In our model, x_i depends on: (1) the actions of other senators and their influences on i, (2) the polarity a_l of the bill l, and (3) threshold t_i. We combine these three factors in the following *influence function*.

$$\begin{aligned} f_i(\mathbf{x}_{-i}, l) &\equiv \sum_{j \in \mathcal{N}_i} \mathrm{w}_{ij} x_j + (p_i \cdot a_l) - t_i \\ &= \mathbf{w}_{i,-i}^T \mathbf{x}_{-i} + (p_i \cdot a_l) - t_i. \end{aligned} \tag{2}$$

In our model, if $f_i(\mathbf{x}_{-i}, l) > 0$), then Senator i will vote *yea* on bill l (i.e., $x_i = +1$). If $f_i(\mathbf{x}_{-i}, l) < 0$, then $x_i = -1$. Finally, if $f_i(\mathbf{x}_{-i}, l) = 0$, then $x_i \in \{+1, -1\}$ (signifying indifference). We next describe the intuition behind these.

The first term in (2) is collective, signifying how others' voting decisions influence senator i's voting decision. In contrast, the second term is specific to an individual senator and a particular bill, capturing how well a senator's ideal point matches the polarity of the bill in the political spectrum. The third term is purely an individualistic term modeling the level of "stubbornness" of a senator. We give more details below.

The first term in (2) gives the net influence on Senator i from others. This is essentially calculated by subtracting the sum of incoming influences from those who voted *nay* from the sum of incoming influences from those who voted *yea*. If the first term is

[2] We use x_i instead of $x_{i,l}$, since l will be clear from context.

positive, then $f_i(\mathbf{x}_{-i}, l)$ increases and Senator i will be more inclined to voting *yea* (the opposite holds when the first term is negative). The second term calculates how well Senator i's ideology aligns with the polarity of the bill l. Whenever the polarity of a bill lines up well with the senator's position, the corresponding p_i and a_l will have the same sign, which means that the $(p_i \cdot a_l)$ term will be some positive number. This will increase the value of $f_i(\mathbf{x}_{-i}, l)$, thereby making it more likely that the senator will vote *yea*. Alternatively, if p_i and a_l have different signs, it will decrease $f_i(\mathbf{x}_{-i}, l)$. Note that just like in the classical ideal point model, the second term $(p_i \cdot a_l)$ does not measure the Euclidean distance between the ideal point of a senator and the polarity of a bill. Neither does it give any comparative measure of ideal points between two different senators. It is indeed specific to a particular senator and a particular bill and measures how well the signs of their ideal point and polarity match. The third term t_i in (2) is the threshold value denoting "resistance" or "stubbornness" of Senator i in the face of influences.

Note that Senator i's choice of action depends on the other senators' choices and others' choices depend on i's choice. Game theory naturally models this mutual dependency among senators. We next define some game-theoretic terminology in our context, which we will use throughout the paper.

A vector of actions, one action per player/senator, is called a *joint action*. A joint action $\mathbf{x}^*$ is said to be a *pure-strategy Nash equilibrium (PSNE)* if no player has any incentive to unilaterally deviate from their action. Said differently, in a PSNE every player plays their *best response* x_i^* (i.e., voting in the optimal way) with respect to others' actions, simultaneously. We use PSNE to mathematically represent *stable outcomes* – no player has any incentive to unilaterally deviate from their action – that our model can generate. One inherent question here is: How can a player/senator calculate the best response to others? In game theory, the best response is an action that maximizes the player's *payoff function*, which we define now in our context.

Following is Senator i's payoff function. Note that this is a quadratic function of the actions.

$$u_i(x_i, \mathbf{x}_{-i}, l) \equiv x_i f_i(\mathbf{x}_{-i}, l) = x_i\left(\mathbf{w}_{i,-i}^T \mathbf{x}_{-i} + p_i \cdot a_l - t_i\right). \qquad (3)$$

The *best response* x_i^* of Senator i to the joint action $\mathbf{x}_{-i}$ of the other senators is the action that maximizes this payoff function. In other words, Senator i's best response x_i^* will be $+1$ if $f_i(\mathbf{x}_{-i}, l) \geq 0$. Conversely, it will be -1 if $f_i(\mathbf{x}_{-i}, l) \leq 0$. An equality indicates indifference between $+1$ and -1. This is summarized below.

$$\left.\begin{array}{l} f_i(\mathbf{x}_{-i}, l) > 0 \Rightarrow x_i^* = +1, \\ f_i(\mathbf{x}_{-i}, l) < 0 \Rightarrow x_i^* = -1, \text{and} \\ f_i(\mathbf{x}_{-i}, l) = 0 \Rightarrow x_i^* \in \{-1, +1\} \end{array}\right\} \Leftrightarrow x_i^* f_i(\mathbf{x}_{-i}, l) \geq 0.$$

This completes the description of our model. We are interested in learning the $w_{i,j}$, p_i, a_l, and t_i parameters of the model using roll call and bill data. We are then interested in computing *all* PSNE, which are the voting outcomes predicted by our model. We will see that learning the model does not readily provide predictions. In fact, the computation of PSNE is provably hard in our case.

Discussion

As a historical backdrop, our influence function (2) has not been arbitrarily designed. It is grounded in mathematical sociology as a threshold model of collective behavior [7] and has been very popular in studying influence. Particularly within computer science community, Kempe et al. [15]'s study of threshold models with the goal of identifying the most influential nodes in a network is the precursor to many studies on influence maximization [5, 16, 17]. In most threshold models, however, the influence weights are non-negative. Here, we do allow negative weights. Most importantly, in almost all of the previous threshold models, the context (e.g., the bill in our case) under which players choose actions has largely been ignored. Here, we incorporate context within the threshold model. We are not aware of any threshold model in our application area that does it.

The significance of modeling context is that we are able to capture those situations where a senator feels strongly in favor or against a bill. As a result, outside influence has less of an effect on their decision. For example, perhaps the bill concerns an issue that is key to a senator's platform, or maybe it is a bill that the senator proposed. Alternatively, the bill might run completely counter to a senator's position, or the bill's effect could be very harmful to the senator's constituents. In any of these cases, the model needs to be able to reflect a higher (or lower) likelihood of that senator voting *yea* (or *nay*). By adding ideal point and polarity parameters similar to those in the ideal point model, our model can account for those voting instances in which senators' decisions derive primarily from their alignment with the bill, and not from external influences.

In addition to bill-specific context, we also model social interactions. There are cases in which $p_i \cdot a_l \approx 0$, i.e., either Senator i's ideal point p_i or the polarity a_l of bill l (or both) are very close to 0. In such cases, the influence term $\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}$ in the payoff function becomes much more important. Moderate senators are particularly susceptible to these instances as their smaller p_i's mean that they are not pushed one way or another on very polarizing bills as much as more extreme senators are. This is where the real power of social interactions comes in, as negotiating with more moderate senators is a crucial part of passing legislation. This aspect is glaringly absent in the ideal point model and all of its immediate derivatives.

3 LEARNING THE MODEL

As a synopsis, we learn our game-theoretic model so that it captures as much of the observed roll call as possible in the form of PSNE. One trivial example that captures *all* observed roll call data as PSNE is an empty model with all parameters 0. In this case, any of the 2^N possible joint actions is a PSNE, including those observed in the roll call data. Obviously, this is not desirable. Therefore, we want to capture as many of the roll call instances as possible in the form of PSNE *while having as few PSNE as possible*. On the other spectrum, if we desire to have *exactly* the observed data as PSNE, it may very well be impossible to instantiate the game model to achieve that.

A natural way of approaching the middle ground between these two extremes is to use a generative mixture model [18] that generates the observed data by picking a PSNE with probability q and a non-PSNE with probability $1 - q$. Obviously, we would like q to be as high as possible, while avoiding overfitting. Honorio and Ortiz

[9] use this approach to learn threshold models from behavioral data (e.g., roll call data). We adapt their algorithms and use their notation here. In particular, in addition to the usual threshold model parameters, we also learn senator ideal points and bill polarities. We first formalize a desirable quality measure below.

Let $\mathcal{NE}(\mathcal{G})$ be the set of PSNE of our game model $\mathcal{G}$. As noted above, we want to maximize the "quality" of PSNE. We could have a huge set $\mathcal{NE}(\mathcal{G})$ which would increase the likelihood of finding an observed voting outcome as a PSNE, but the quality of $\mathcal{NE}(\mathcal{G})$ in that case would be very low, because most of the PSNE in that set would not match the data. To formalize the notion of the quality of PSNE, first, $\pi(\mathcal{G})$ below quantifies the size of $\mathcal{NE}(\mathcal{G})$ relative to total number of possibilities.

$$\pi(\mathcal{G}) \equiv |\mathcal{NE}(\mathcal{G})|/2^N. \tag{4}$$

If M is the number of observed bills voted on, $q \cdot M$ is the number of bills that the generative model picked from the set of PSNE. In order to maximize the quality of our model's $\mathcal{NE}(\mathcal{G})$, we want to maximize $q \cdot M$ relative to $\pi(\mathcal{G})$. This can be written as $\frac{q \cdot M}{|\mathcal{NE}(\mathcal{G})|/2^N}$. Since N and M are fixed, this expression is proportional to

$$\log_{10} \frac{q}{|\mathcal{NE}(\mathcal{G})|} \tag{5}$$

which allows us to measure the quality of $\mathcal{NE}(\mathcal{G})$ while taking into account its size.

3.1 Learning Algorithm

We first define a *loss function* to compute the errors in best responses. Let ℓ be the 0/1 loss function, $\ell(z) = 1[z < 0]$. Then

$$\ell(f_i(\mathbf{x}_{-i}), l) = \begin{cases} 1 & \text{if } i \text{ not playing best response} \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Equation (3) defines the payoff function of Senator i. For a given bill l and the actions of others, Senator i's best response maximizes the payoff function $u_i(x_i, \mathbf{x}_{-i}, l)$. Using the loss function given above, the number of bills that are not captured as PSNE is given by

$$\sum_l \max_i \ell \left[x_i^{(l)} (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} + (p_i \cdot a_l) - t_i) \right] \tag{7}$$

and the proportion of bills that are not captured as equilibria is given by

$$\frac{1}{M} \sum_l \max_i \ell \left[x_i^{(l)} (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} + (p_i \cdot a_l) - t_i) \right]. \tag{8}$$

Finally, the overall optimization function for the model consists of minimizing the proportion of bills which result in non-PSNE voting outcomes. It can be shown to be equivalent to maximizing (5).

$$\min_{\mathbf{W},\mathbf{t},\mathbf{p},\mathbf{a}} \frac{1}{M} \sum_l \max_i \ell \left[x_i^{(l)} (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} + (p_i \cdot a_l) - t_i) \right]. \tag{9}$$

Although we have several additional parameters, our implementation follows the general framework for threshold models [9]. We use a logistic loss function $\ell(z) = \log(1 + e^{-z})$ in order to produce a simpler minimization expression with a smooth upper bound. We also add an ℓ_1-norm regularizer $\rho||\mathbf{W}||_1$ (where $\rho > 0$) in order to bias the learning to sparse networks and to reduce the likelihood of over-fitting to the training data. We add another ℓ_1-norm regularizer $\rho'||\mathbf{p} \cdot \mathbf{a}||_1$ to penalize the ideal points and polarities as well. The minimization can be written in the following simultaneous logistic regression form:

$$\min_{\mathbf{W},\mathbf{t},\mathbf{p},\mathbf{a}} \frac{1}{M} \sum_l \log \left[1 + \sum_i e^{-x_i^{(l)} (\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} + (p_i \cdot a_l)) - t_i} \right] + \rho||\mathbf{W}||_1 + \rho'||\mathbf{p} \cdot \mathbf{a}||_1 \tag{10}$$

where $\rho, \rho' > 0$. We use the ℓ_1- projection method [22] for optimization. Please see [9] for details in the context of LIGs.

3.2 Issues with Learning

While implementing the learning algorithm, we had to address the following issues.

Anchoring Ideal Points. Traditionally, Democrats have negative ideal points and Republicans have positive ideal points. However, without biasing the optimization (10), it is possible to get the opposite signs. To avoid this, we set anchors, which is common in the literature [6]. We used ideal points generated by Poole et al. [20]'s method [3] to find the most extreme senators in each party for the 114th Senate and used those as anchors. This meant giving Warren (D-MA) and Sanders (I-VT) ideal points of -4, and Cruz (R-TX) and Shelby (R-AL) ideal points of +4.

Determining Validation Bill Polarities. We cross-validated our model parameters using held-out roll call data from the same Congress. However, when our model is given a new bill l', its polarity $a_{l'}$ is unknown and has not been learned as l' was not part of the training set. [4] We solved this problem using the subject codes provided by Congress.gov. These hand-entered codes are terms that describe the bills' contents (e.g. "Child safety and welfare", "Drug Enforcement Administration", or "Health care coverage and access"). Each bill usually contains multiple codes, with the median bill in the 114th Senate containing 60 subject codes out of a total of 737. For a new bill l', we find the Euclidean distance D between l' and each bill l in the training set. This procedure gives the bill l^* in the training set whose subject codes most closely match that of l'. Then, we assign l^*'s polarity (which was learned during training) to l'. This allows us to use the model with new bills, by approximating their polarities with bills that have already been seen by the model.

An alternative to using the subject codes would be to use the "topics" in a bill based on the bill text. This technique was used in the issue-adjusted ideal point model [6]. However, they also tested their model using the subject codes directly, instead than learning them using topic modeling. They find that the subject codes serve as a good approximation to topic modeling.

3.3 Model Selection

The interplay between the regularization parameters ρ and ρ' can lead the model to forming a mixture of an ideal point model and a social interaction model, or cause it to strongly shift towards one of the two. The ideal setting for ρ and ρ' produces a network that

[3] https://voteviewblog.com/2015/08/16/alpha-nominate-applied-to-the-114th-senate

[4] Note that the ideal point of the senators, once they are learned using the training data, can be used with held-out data.

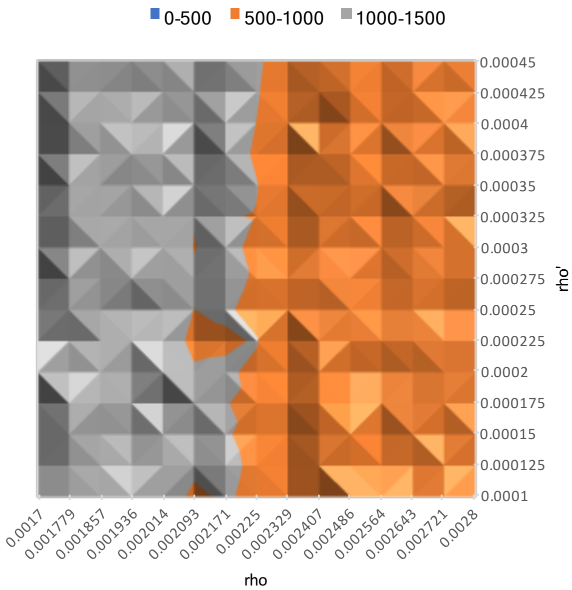

Figure 1: The number of edges in the influence network. As expected, the number of edges is mainly affected by ρ which penalizes dense graphs.

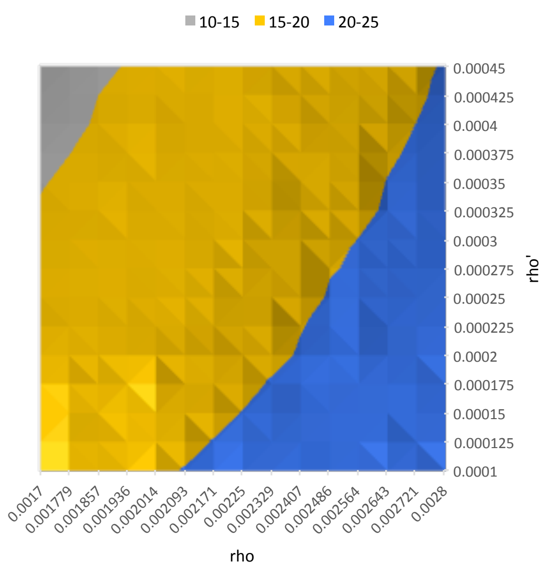

Figure 2: Validation error in best responses (in %). The blue region not only has high error but also degrades the quality measure given in Equation (5) by generating sparse graphs (due to high ρ) that have numerous PSNE. The gray region has low error, but produces very dense graphs that are computationally infeasible.

is sparse enough to avoid overfitting but large enough to represent the observed data, while also maintaining reasonable ideal points.

We trained the model using Session I voting data and rigorously performed cross-validation using Session II data. We tracked four different measures:

(1) The number of edges in the influence network: In order to avoid overfitting and burden on computing PSNE, we want this to be below 1,200. The plot is shown in Fig. 1.
(2) The proportion q_T of training votes that are PSNE: We desire a large q_T.
(3) Validation error: We consider multiple measures of errors in best responses with respect to held-out data. The error percentage is shown in Fig. 2.
(4) The Euclidean distance between the learned ideal points and the average ideal points: A reasonable assumption is that all Democrats should have $p_i < 0$ and all Republicans should have $p_i > 0$. The average ideal points p_i over 500 trials for a model with a ρ and ρ' values generating ideal points similar to Poole et al. [20]'s calculations had this property. We desire the learned ideal points to be close to those.

We have examined several other plots that are not shown here for space constraints. In sum, we settled on $\rho = 0.00225$ and $\rho' = 0.0004$. This produced a graph of 998 edges, and a validation error of 16.0%. It also produced ideal points close to the average.

4 EQUILIBRIA COMPUTATION

Our model gives joint predictions of senators' voting behaviors in the form of PSNE. As a result, the learned model is not immediately useful without computing the PSNE. We first show that this problem is provably hard.

Theorem 4.1. *It is NP-complete to decide whether there exists a PSNE; whether there exists a PSNE consistent with a designated set of senators voting yea; and whether there exists a PSNE with at least a given number of k senators voting yea.*

Theorem 4.2. *It is #P-complete to count the number of PSNE, even when the underlying network is bipartite or star.*

Proof Sketch. Given any LIG instance we can create an instance of our game model by assigning $a_l = 0$ for any bill l. The proofs of the above two theorems then follow from the hardness of LIGs [12]. □

To compute all PSNE of our model, we applied a backtracking search algorithm [12] that systematically explores the search space of 2^{100} due to 100 senators and binary actions. Whenever a node is assigned an action, we use the **NashProp** algorithm [12, 19] to propagate its action locally and see if it leads to any contradiction. Perhaps not surprisingly, non-zero values of the contextual parameters $\mathbf{p}$ and a_l significantly improve the runtime of the algorithm compared to when either of them is 0. To see why, we can view the effect of these two parameters as an adjustment to thresholds $\mathbf{t}$ by $-(\mathbf{p} \cdot a_l)$. When the threshold is adjusted this way, it increases or decreases for non-zero values of $\mathbf{p}$ and a_l. This in turn helps the backtracking search algorithm quickly find contradictions and prune large areas of the search space.

When run on a high performance computing grid with $a_l = 0$, the algorithm took 9 minutes to compute all PSNE in the 100-player game. When $a_l \approx -3.22$ (one of the largest negative polarities), it took only 0.49 seconds.

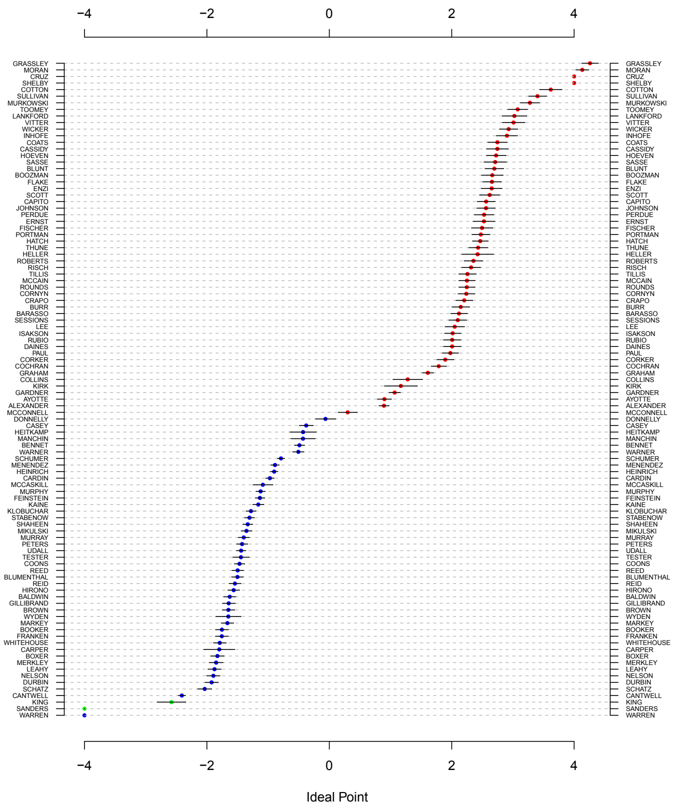

Figure 3: Average of learned ideal point values and 95% confidence intervals (over 500 samples) for senators in the 114th Congress. The results are consistent with the classical ideal point models without any social interactions component.

5 EXPERIMENTS: THE 114TH SENATE

We ran our model using roll call and bill-specific data from the 114th U.S. Senate (2015-2017). In this Senate, the same 100 senators kept their seats for both sessions. We considered all voting instances that had an associated bill (including amendments), and ignored all other types of votes (e.g. confirmations). We used most of the first session as a training set (308 voting instances) and the second session as a held-out set (145 voting instances). In total, there were 737 subject codes that appeared in either session.

5.1 Learned Ideal Points

With very few exceptions, our model assigns negative ideal points to Democrats and positive ideal points to Republicans. A plot of the learned ideal points is shown in Fig. 3. This is an encouraging result as it demonstrates the our model's ability to capture ideal point data in a manner similar to traditional ideal point models. The average ideal point for Democrats and Independents is -1.49 with a standard deviation of 0.76, and for Republicans 2.47 with a standard deviation of 0.79. One exception is Senator Collins (R-ME), who is assigned a negative polarity (and we know that she often sides with Democrats). Another exception is Senator Casey (D-PA), who is assigned a positive polarity. Note that even though the results are consistent with the ideal point literature, we must put them in context together with the influence network for our model.

5.2 Influence Network

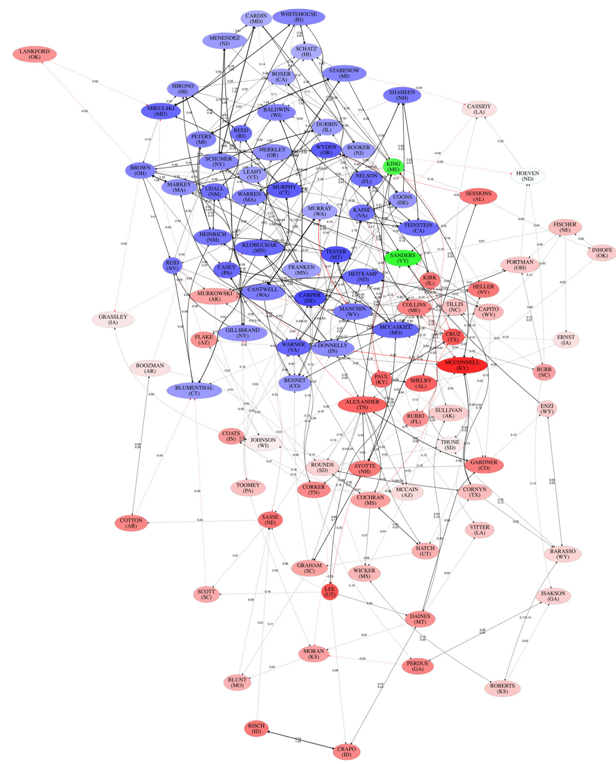

Figure 4: A bird's-eye-view of the learned influence network (force-directed drawing). only the strongest edges are shown. Each node represents a senator, with colors to indicate party and shade to indicate threshold t_i level (darker means higher threshold). Directed edge represent one senator's influence on another, with black edges representing positive influence and red edges negative.

Fig. 4 illustrates the network as a directed graph with nodes representing senators and edges representing influence. Only the strongest 30% incoming and outgoing edges for each node are shown. Nodes are colored according to their political party (Democrats – blue, Republicans – red, Independents – green) and shaded according to their effective threshold level t_i, adjusted by $-p_i \cdot a_l$. Darker colors represent higher thresholds (more stubborn) and lighter colors represent lower thresholds (more susceptible to influence). Note that the structure of this graph will not change for different bills, but senators' adjusted thresholds will, due to changes in bill polarity a_l. This means that the network would look the same except nodes would be shaded darker or lighter depending on the bill.

5.3 Comparison with Ideal Point Models

Clinton et al. [2] report an accuracy of 89.9% of the individual votes in the 106th U.S. House of Representatives. However, that is training accuracy, since their model is trained using *all* of the roll-call data without any held-out data. The challenge with measuring test accuracy with held-out data is that the polarity a_l of a bill l in the held-out data is unknown. We describe how we mitigate this challenge using the subject codes of bills in Section 3.2. Even if we apply the same procedure to the ideal point model, it will perform poorly with respect to the desirable quality measure $\log_{10} \frac{q}{|\mathcal{NE}(\mathcal{G})|}$ given in Equation (5). The reason is that the ideal point model is designed to give *individual* predictions, not *joint* predictions. [5] For the classical ideal point model, the network structure of social interactions is empty. As a result, it will lead to a staggering number of PSNE (computationally infeasible), which will exponentially bring down the above quality measure.

5.4 Comparison with LIG Model

We compare and contrast our model with the linear influence game (LIG) model [10–12], which is the state-of-the-art model on influence in Congress to the best of our knowledge. Since the LIG model does not have any notion of bill context while ours does, in order to compare ours with LIG we had to look into three distinct types of bills in the political spectrum: bills that are very much left-leaning (very negative polarity), bills that are in the middle (polarity close to 0), and bills that are very much right-leaning (very positive polarity). The results are presented below, followed by a more general comparison between our model and the LIG model. The bills used for comparison have been purposely held out from the training set.

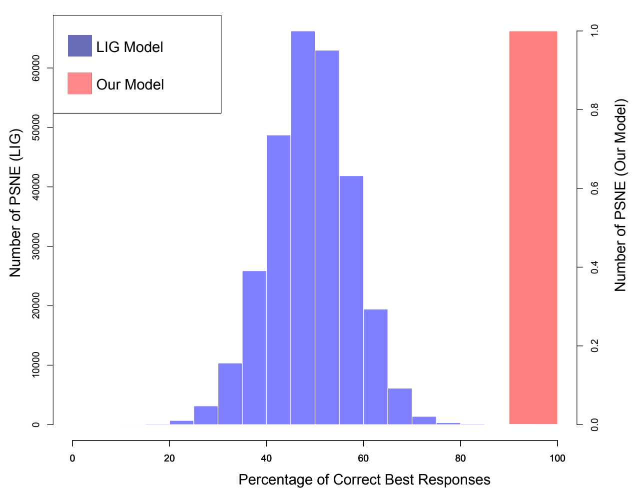

Figure 5: Vote # 49: Approval of Keystone XL Pipeline. Our model (red) produced *one single PSNE* that correctly predicted 91 votes. LIG's (blue) median correct prediction is 50 votes. Most importantly, the LIG model predicts several hundred thousand equilibrium outcomes, which makes it hard to pinpoint the actual outcome.

[5]Joining individual behaviors together does not lead to a joint voting behavior, because it does not model interdependence of actions. That is, if some senator were to change his/her vote, the ideal point model does not model which specific senators would be affected by this change and by how much.

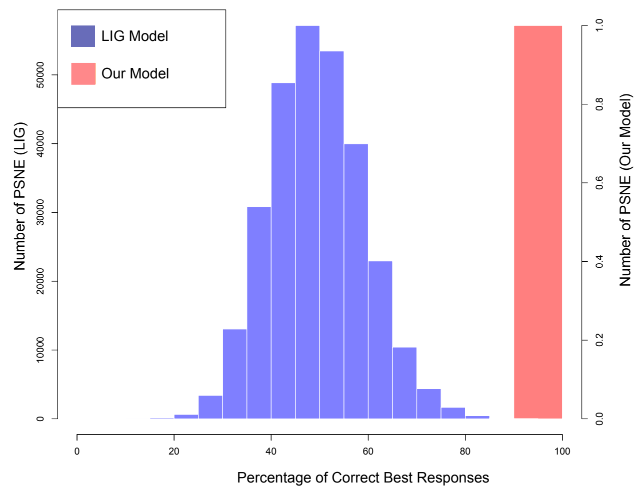

Figure 6: Vote #89: S. Amendment 777 to S.Con.Res.11. Our model (red) produced one PSNE that correctly predicted 92 votes. LIG's predictions are on the left (blue). Once again, our model gives drastically better quality PSNE compared to LIG.

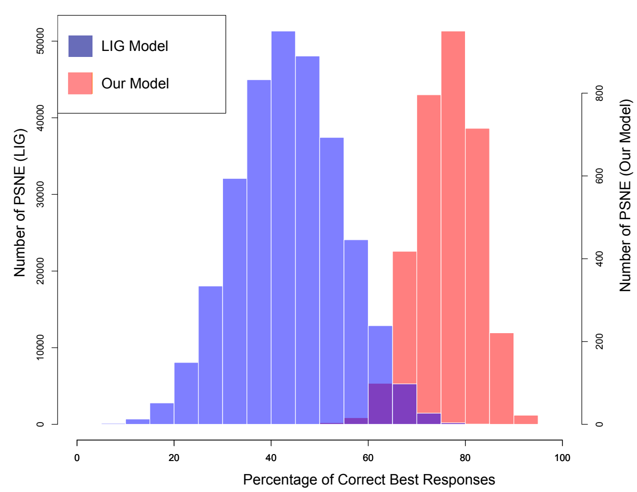

Figure 7: Vote #229: Motion to Invoke Cloture on the Motion Re: H.R. 1735. Our model predicts much fewer and a lot more accurate PSNE compared to LIG.

Right-Leaning Bill: Keystone XL Pipeline (a_l = 1.426). This was one of the most controversial bills of the 114th Senate [4], instigating protests from environmental groups and the public. This right-leaning bill passed by a 62–36 vote. Our model produced one single PSNE for this bill, which correctly predicted 91 votes. In contrast, the LIG model [12] produced 287,400 possible PSNE, where the median number of correctly predicted votes was 50. As shown in Fig. 5, the quality of our model's PSNE for this bill is significantly better than the state-of-the-art LIG model.

Left-Leaning Bill: Senate Amendment 777 (a_l = −3.7025). Proposed by Bernie Sanders, this amendment is a left-leaning one recognizing climate change and proposing actions to cut carbon pollution. It failed by a 49–50 vote. Similar to the Keystone bill, our model produces one single PSNE which correctly predicts 92 of the votes of Amdt. 777. In contrast, the LIG model gives a median number of correct votes of 50, among hundreds of thousands of PSNE. This is shown in Fig. 6.

The above two bills illustrate the LIG model's inability to capture the effect of bill content on voting behavior. First, 287,400 possible PSNE given by the LIG model are too many for any predictive model. Second, for more extreme legislations, lawmakers become more rigid in response to influence. Therefore, the number of PSNE should drop as it gets more difficult for legislators to influence each other. *Our model effectively captures this polarization by narrowing down the prediction.*

Middle-Ground Bill: Motion to Invoke the Cloture Motion ($a_l = -0.0297$). This is a near-unanimous bill that passed by 84–12, with the *nay* votes coming from both extremes: Cruz (R-TX), Paul (R-KY), Warren (D-MA), and Sanders (I-VT). With $a_l \approx 0$, our model produces 3,242 possible PSNE. This number, while a lot lower than LIG's prediction of 287,400 PSNE, stems from the minimization of the $(p_i \cdot a_l)$, and therefore a relaxation of the effective threshold parameter t_i (see Equation (2)). As shown in Fig. 7, the overall quality of the PSNE of our model is much better (i.e., histogram is on the right) than the quality of PSNE of LIG.

As the above experiments show, our model produces a much narrower band of PSNE as its stable-outcome prediction than the LIG model [10–12]. In particular, the prediction of our model is remarkably precise and high-quality if a bill falls in either end of the political spectrum. This is further evident from the following comprehensive comparison in terms of the quality of PSNE.

Quality of PSNE. Recall that one of our learning goals is to maximize the quality measure of PSNE defined by Equation (5). With this measure as a yardstick, our results can be comprehensively compared with the LIG model.

First, even on graphs of similar size, the number of PSNE is much higher for the LIG than for our model. We can compare the two by setting $a_l \approx 0$ in order to allow our model to generate a large number of PSNE (as described above). With a graph of ≈ 900 edges, the LIG model produces over 287,400 PSNE, taking many hours of computation time. For a similar sized graph and with $a_l \approx 0$, our model produces only 3,242 PSNE, taking only minutes. Second, based Equation (5), LIG's quality measure is −4.16, whereas ours is −2.69. Therefore, our model is over 20 times more precise on the held-out data than the LIG model. Note that non-zero a_l values result in a much lower number of PSNE in our case. As a result, the quality measure of our model is lot more pronounced in such cases.

5.5 Most Influential Senators

It is natural to ask who the most influential senators are for passing a bill. The implication is that when the most influential senators group up and vote *yea* on a bill, then everybody else would be influenced to vote *yea*. Given the set of all PSNE, we can compute a small set of most influential senators by applying the well-known greedy set-cover algorithm with a provable log-factor approximation guarantee [12]. According to our model of the 114th Senate, only four senators constitute the most influential set of senators in the sense that if they agree on a bill, others will also come along. One such most influential set is {Carper (D-DE), Schumer (D-NY), Cruz (R, TX), Peters (D, MI)} (note that the 114th Senate was Republican-majority). There are several other such sets, which are shown in Fig. 8(a). In contrast, due to its wide band of predictions, the LIG model gives a much bigger set of nine most influential senators. This is shown in Fig. 8(b). This not only highlights an improvement over the LIG model but also makes our framework applicable to the minimal targeted intervention problem where we are interested in intervening as few individuals as possible to achieve a desirable social outcome (e.g., reducing smoking, eating healthy, etc.).

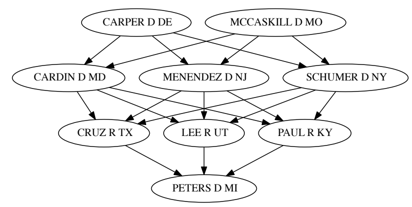

(a) Our Model

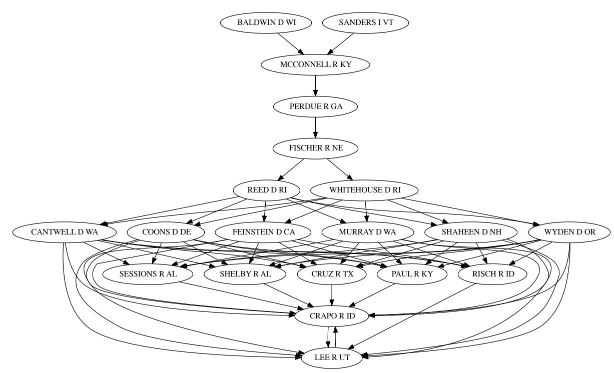

(b) LIG Model

Figure 8: A pictorial view of the sets of most-influential nodes. A directed acyclic graph illustrates all possible selections of most-influential senators, with each directed path from the top level to the bottom representing one possible set. Our model selects a much smaller set of most-influential senators than LIG.

6 CONCLUDING REMARKS

We have presented a model of strategic behavior grounded in threshold models. Unlike the existing literature on threshold models, ours takes into account the context under which the behaviors take place. This leads to significant improvements over previous models, both in terms of the number of PSNE as well as the quality of these PSNE. Among many exciting future directions is a Bayesian games perspective on Senate voting.

Acknowledgments

We would like to thank Drs. Jean Honorio and Luis E. Ortiz for providing us the machine learning codes for LIGs [9] and Ian Dieli for building a congressional database to help with this research.

REFERENCES

[1] Brandice Canes-Wrone, David W Brady, and John F Cogan. 2002. Out of step, out of office: Electoral accountability and House members' voting. *American Political Science Review* 96, 01 (2002), 127–140.

[2] Joshua Clinton, Simon Jackman, and Douglas Rivers. 2004. The statistical analysis of roll call data. *American Political Science Review* 98, 02 (2004), 355–370.

[3] Otto A. Davis, Melvin J. Hinich, and Peter C. Ordeshook. 1970. An Expository Development of a Mathematical Model of the Electoral Process. *American Political Science Review* 64, 2 (1970), 426âĂŞ448. https://doi.org/10.2307/1953842

[4] Peter Erickson and Michael Lazarus. 2014. Impact of the Keystone XL pipeline on global oil markets and greenhouse gas emissions. *Nature Climate Change* 4, 9 (2014), 778–781.

[5] Eyal Even-Dar and Asaf Shapira. 2007. A Note on Maximizing the Spread of Influence in Social Network s. In *Internet and Network Economics (WINE)*.

[6] Sean M Gerrish. 2013. *Applications of Latent Variable Models in Modeling Influence and Decision Making*. Technical Report. PRINCETON UNIV NJ.

[7] Mark Granovetter. 1978. Threshold Models of Collective Behavior. *The American Journal of Sociology* 83, 6 (1978), 1420–1443. http://www.jstor.org/stable/2778111

[8] Ronald K Hambleton and Hariharan Swaminathan. 2013. *Item response theory: Principles and applications*. Springer Science & Business Media.

[9] Jean Honorio and Luis E Ortiz. 2015. Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *Journal of Machine Learning Research* 16 (2015), 1157–1210.

[10] Mohammad Tanvir Irfan. 2013. *Causal Strategic Inference in Social and Economic Networks*. Ph.D. Dissertation. Stony Brook University, Department of Computer Science. http://www.bowdoin.edu/~mirfan/papers/Mohammad_Tanvir_Irfan_Dissertation.pdf.

[11] Mohammad T Irfan and Luis E Ortiz. 2011. A Game-Theoretic Approach to Influence in Networks. In *AAAI Conference on Artificial Intelligence*. http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3746

[12] Mohammad T Irfan and Luis E Ortiz. 2014. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence* 215 (2014), 79–119.

[13] Jeffery A Jenkins. 1999. Examining the bonding effects of party: A comparative analysis of roll-call voting in the US and Confederate Houses. *American Journal of Political Science* (1999), 1144–1165.

[14] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. 2001. Graphical Models for Game Theory. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 253–260.

[15] David Kempe, John Kleinberg, and Éva Tardos. 2003. Maximizing the Spread of Influence through a Social Network. In *Proceedings of the 9th ACM SIGKDD International Conference on Kno wledge Discovery and Data Mining*.

[16] Jon Kleinberg. 2007. Cascading Behavior in Networks: Algorithmic and Economic Issues. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, Chapter 24, 613–632.

[17] Elchanan Mossel and Sebastien Roch. 2010. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.* 39, 6 (2010), 2176–2188.

[18] Kevin P Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT press.

[19] Luis E. Ortiz and Michael Kearns. 2003. Nash Propagation for Loopy Graphical Games. In *Advances in Neural Information Processing Systems 15*, Suzanna Becker Becker, Sebastian Thrun Thrun, and Klaus Obermayer (Eds.). 817–824.

[20] Keith Poole, Howard Rosenthal, and Christopher Hare. 2015. Alpha-NOMINATE Applied to the 114th Senate. https://voteviewblog.com/2015/08/16/alpha-nominate-applied-to-the-114th-senate/. *Voteview* August 16 (2015).

[21] Eric Schickler. 2000. Institutional change in the House of Representatives, 1867–1998: a test of partisan and ideological power balance models. *American Political Science Review* 94, 02 (2000), 269–288.

[22] Mark Schmidt, Glenn Fung, and Romer Rosales. 2007. Fast optimization methods for l1 regularization: A comparative study and two new approaches. *Proceedings of the 18th European Conference on Machine Learning, ECML '07* (2007), 286–297. https://www.cs.ubc.ca/~schmidtm/Software/L1General.html.