RESEARCH

Spheres of Legislation: Polarization and Most Influential Nodes in Behavioral Context

Andrew C Phillips^{1†}
 , Mohammad T Irfan^{2*}
 and Luca Ostertag-Hill^{2†}

An earlier version of this article appeared as a full paper at the 8th International Conference on Complex Networks and Their Applications (COMPLEX NETWORKS 2019) in Lisbon, Portugal [PIOH20]. We have significantly extended that paper in this article. Following are the notable additions to the main body of the paper: Section 2.1: Preparing Congressional Roll-Call Data; LIG example and an example for each definitions in Section 3; Section 4.2: Cross-Validation and Model Selection for LIG; Section 5.2: Influence Weights and Thresholds; Section 6.1: PSNE Computation; and Figures 1, 2, 3, 7, 8, 11, 13, and 14. In addition, there are light revisions all throughout the conference version, but the following parts went through significant revision: Section 4.1: Learning Algorithm, Section 6.3: Game-Theoretic Vs. Structural Centrality Measures, and Section 7: Towards Richer Models. In this article, we are also able to include an Appendix with a detailed literature review and many figures and tables.

*Correspondence:

mirfan@bowdoin.edu

²Department of Computer Science, Bowdoin College, 255 Maine St, Brunswick, Maine 04011, USA

Full list of author information is available at the end of the article

[†]ACP and LOH worked on this research as undergraduate students at Bowdoin.

Abstract

Game-theoretic models of influence in networks often assume the network structure to be static. In this paper, we allow the network structure to vary according to the underlying behavioral context. This leads to several interesting questions on two fronts. First, how do we identify different contexts and learn the corresponding network structures using real-world data? We focus on the U.S. Senate and apply unsupervised machine learning techniques, such as fuzzy clustering algorithms and generative models, to identify spheres of legislation as context and learn an influence network for each sphere. Second, how do we analyze these networks in order to gain an insight into the role played by the spheres of legislation in various interesting constructs like polarization and most influential nodes? To this end, we apply both game-theoretic and social network analysis techniques. In particular, we show that game-theoretic notion of most influential nodes brings out the strategic aspects of interactions like bipartisan grouping, which structural centrality measures fail to capture.

Keywords: influence in networks; machine learning and networks; computational game theory; graphical games

1 Introduction

In recent times, the study of social influence has extended beyond mathematical sociology [Gra78, WF94] and has entered the realm of computation [KKET03, KKET05, CF07, LAH07, CF08, CCC+11, LCWZ13, IO14, BMA15, AAAY+17, IG18]. A computational study of “influence”—however we define it—is key to understanding the behavior of individuals embedded in networks. In this paper, we model and analyze social influence in a strategic setting where one’s behavior depends on others’ behavior. Since game theory reliably captures such interdependence of behavior in a population, we ground our computational approach in game

theory. The strategic setting of our interest here is the U.S. Senate. We model the influence structure among the senators by taking into account the relevant context, which we call the spheres of legislation. We learn these models of influence from the real-world behavioral data on Senate bills and voting records. Our particular focus is on analyzing machine learned influence networks to answer various questions on polarization and most influential nodes.

Interestingly, most computational models of influence assume a fixed network structure among individuals. We relax this simplifying assumption, allowing the network of influence to vary according to the spheres of legislation. For example, bills on finance may induce a very different influence network among senators than bills on defense, which may in turn have different impacts on inference problems like polarization and most influential nodes. One central question in this regard is: How do we identify different spheres of legislation that may have different implications on these inference problems? We address this in Section 2.

After identifying spheres of legislation, we can learn an influence network among the senators for each sphere by adopting game-theoretic models of strategic behavior. Broadly speaking, modeling and analyzing congressional voting behavior has been a trending topic in both political science and computer science [CJR04, Poo08, GB12, IO14, IG18], in part due to the availability of data.

In particular, we use the linear influence game (LIG) model of strategic behavior proposed by Irfan and Ortiz [IO11, IO14]. We learn these models using data from the spheres of legislation. In LIG, each senator exerts influence upon (and is subject to influences from) other senators in a network-structured way. The model focuses on interdependence among the senators and adopts the game-theoretic solution concept of *Nash equilibrium* to predict stable outcomes from a complex system of influences. This notion of Nash equilibrium leads to a definition of the most influential senators, where a group of senators is called most influential with respect to a desirable stable outcome (or desirable PSNE) if their support for that outcome influences enough other individuals to achieve that outcome. The LIG model will be elaborated in Section 3 and machine learning of this model using the spheres of legislation will be detailed in Section 4.

The main theme of this paper is how influence networks are affected by the underlying context, where the spheres of legislation represent the context. We should note here that contextual information has been considered before in the congressional setting. Recently, Irfan and Gordon [IG18] extended the LIG model to account for the bill context. They sought to combine both the social interactions and ideological leaning aspects of congressional voting. Although their model incorporates an aspect of bill context by assigning polarities to bills, it does not completely disentangle the influence network from the topics or subjects of the bills. *The influence network produced by their model does not change due to the bill context, whereas we allow the influence network to change based on the spheres of legislation.*

Additionally, Irfan and Gordon [IG18] focused on making predictions given a single bill, rather than analyzing the network as a whole. In Section 7, we briefly touch upon how their richer model can be applied to multiple spheres of legislation, thereby allowing the network to vary according to the context. The full exploration of their model within the spheres of legislation remains open.

As alluded above, while game-theoretic *prediction* of congressional votes has been well studied using the LIG model and its extensions [IO11, IO14, IG18], an *analysis* of the machine learned networks of influence did not get much attention, which we address here. Similarly, algorithms for computing most influential nodes in a strategic setting have been studied before (e.g., [IO14]), but their structural analogs like centrality measures have not been explored in a comparative fashion. In other words, what do we gain by using a game-theoretic definition of most influential nodes as opposed to a structural definition? We address questions like this.

Furthermore, polarization in social networks has been well studied [New06, GMJCK13, BMA15], especially in the political arena [MPR16, ZFT+08, WPF+09, CRF+11, GAS+15]. Three salient points distinguish our approach from the rich body of literature: (1) Ours is a model-based approach, where networks are central to predicting collective outcomes, (2) we learn the networks using behavioral data, because the networks are not observable, and (3) we seek to show that polarization in Senate varies according to the spheres of legislation. We do not touch on the rising polarization in Senate over time, which by now is a well-settled matter [Far15].

The past two terms of Congress are especially interesting for analyzing network behavior and polarization. The 114th Congress ran from January 2015 to January 2017 and the 115th Congress ran from January 2017 to January 2019. In both terms, Republicans controlled the Senate, but the executive power was different. In the 114th Congress, Barack Obama (D) held the presidency; in the 115th, Donald Trump (R) held the presidency. Despite the two opposing parties holding presidency, both terms are perceived to be deeply polarized. Interestingly, when we study different influence networks *among the same group of senators* arising from different spheres of legislation we find that polarization is not really equally applicable. It very much depends on the sphere under consideration. Our aim is to put polarization and other inference questions like most influential nodes in context. Before continuing, we note that a detailed literature review is provided in Appendix A.

2 Spheres of Legislation

We use an unsupervised machine learning technique, namely fuzzy clustering, to assign bills to different spheres of legislation based on the bill subjects. We learn the linear influence game (LIG) models, analyze influence networks, compute equilibria, and find most influential senators for each sphere separately. By doing so, we are able to examine differences and make comparative judgments across the spheres. We first describe how we prepare the data for clustering.

2.1 Preparing Congressional Roll-Call Data

Our model relies on data obtained from the @unitedstates project’s Congress repository, (<https://github.com/unitedstates/congress>), a public domain program that allows easy access to official congressional data from the Congressional Research Service (CRS). In particular, we use bill data and roll-call data. Roll-call data contains each senators “yea,” “nay,” or abstaining votes, while bill data includes a list of subjects incident to the bill, among other attributes. These 820 subjects range from “Abortion” to “Zimbabwe,” and a multitude of subjects describes each bill. Additionally, each bill is assigned a single “top term,” the broad subject which best

describes the bill out of 23 possible top-level subjects. We use the roll-call data to represent senator voting behavior, and bill data to extract bill topics.

Working with the combined data from multiple terms presents a troubling problem for graph-based analysis: senators come and go. Seats in the United States Senate often change during midterm elections, when constituents have the chance to re-elect or replace incumbent senators. In the middle of a term, if a senator leaves their seat, a successor is appointed until the state can hold a special election to find a democratically-elected replacement. In the 2016 midterm election at the onset of the 115th congress, seven senate seats were changed; during the course of the 115th congress, due to cabinet appointments by President Trump, scandals, and a death, the senate saw seven more changes.

When a senator is not present for a vote, they neither influence nor can be influenced by other senators' votes during that roll call. Some senators in our dataset never once overlap with another; one left the senate before the other even joined. In order to reduce the number of these cases, we *combined* non-permanent senators under the following circumstances, given a departing senator *A* and an incoming senator *B*:

- 1 Senator *A* does not run during an election, and senator *B* of the same party is elected to replace them.
- 2 Senator *A* voluntarily or involuntarily steps down, and senator *B* of the same party is appointed as their replacement

In these circumstances, we assume that the incoming senators behave similarly to the departing senators. In other circumstances, such as when a senator loses their seat to member of the opposing party, we keep both senators in the dataset. Changes in senate membership, and the operations undertaken to reduce the number of total senators, are described in Table 1. Additionally, learning the LIG model requires data to be in the form of two discrete values: 1 (*yea*) or -1 (*nay*). When a senator is not present for a vote—either because they were absent on that day, or were not yet holding office—we fill in the missing data with the mean vote of their party.^[1]

Table 1 Changes in congressional membership within the 114th-115th congresses. When two senators are combined, the senator name in boldface indicates the label which was kept. This table should be referred to when examining figures and graphs which depict individual senators.

State	Incoming	Departing	Reason for Change	Operation
IN	Young (R)	Coats (R)	Midterm Election	Combined
LA	Kennedy (R)	Vitter (R)	Midterm Election	Combined
NV	Cortez Masto (D)	Reid (D)	Midterm Election	Combined
NH	Hassan (D)	Ayotte (R)	Midterm Election	
IL	Duckworth (D)	Kirk (R)	Midterm Election	
MD	Van Hollen (D)	Mikulski (D)	Midterm Election	Combined
CA	Harris (D)	Boxer (R)	Midterm Election	
AL	Strange (R)	Sessions (R)	Appointment	Combined
AL	Jones (D)	Strange (R)	Special Election	
MN	Smith (D)	Franken (D)	Appointment & Special Election	Combined
MS	Hyde-Smith (R)	Cochran (R)	Appointment & Special Election	Combined
AZ	Kyl (R)	McCain (R)	Appointment	Combined

^[1]Senators King I-ME and Sanders I-VT are considered Democrats in these circumstances; while unaffiliated with that party, their progressive ideologies match the average Democratic vote far more closely than the average Republican vote.

2.2 Clustering Algorithm

We seek to split the bills into a small number of broad categories, each of which encompasses many bills. Each bill has been tagged with a “top term” by @united-states. The top term corresponds to congress.gov’s tag of “policy area.” According to congress.gov, “one Policy Area term, which best describes an entire measure, is assigned to every public bill or resolution.” The policy area vocabulary consists of 32 terms.^[2] However, these top terms/policy areas are too specific to be used as clusters on their own. In fact, making each top term its own cluster would result in some clusters containing only one bill and others containing a hundred. This would be problematic, because the “outcome space” of LIGs is exponential in size and as a result, learning LIGs requires a relatively large amount of data.

Rather than manually re-categorizing bills, we took a statistical clustering approach to grouping, based on a bill’s assigned “top term” in addition to all subjects it contains. For each data point, we assigned each possible subject a weight: 0 if missing, 1 if present, or 10 if it is the “top term.” By including both measures of subjects (top and regular), we produce more meaningful categories than using top terms or bill subjects lists alone.

In data science, K-Means (KM) is often used as a simple yet effective clustering algorithm [Mac67]. In KM, n data points are partitioned into k clusters based on their Euclidean distance from cluster centers. In each iteration, every data point is assigned a cluster based on the closest centroid; then the centroid of each cluster is reset to the average position of each data point within that cluster. The process repeats until centroid positions converge. The problem of choosing k is left up to the researcher; generally, k is chosen by trial and error. Cluster membership in KM is crisp, meaning each data point belongs to one and only one cluster. While effective at producing distinct clusters, KM is not ideal for our purposes because bills often belong to multiple clusters. For example, a bill about increasing defense spending is about national security as well as economics.

The Fuzzy C-Means (FCM) clustering algorithm addresses this problem. FCM is an extension of KM which allows for overlaps in clusters [Bez81, Ros04]. The objective function in FCM is largely the same as in KM, with the addition of membership values w_{ij} and a fuzzifier m . Membership values describe how closely each data point i belongs to cluster j . The fuzzifier changes membership values: $m = 1$ results in crisp clusters ($w_{ij} \in \{0, 1\}$), and higher values of m result in fuzzier clusters. The FCM algorithm produces a list of cluster centers, describing the position of each centroid, as well as the fuzzy partition matrix, describing the membership degree of each bill to every cluster.

Iterating over a range of values, we found that number of clusters, $c = 4$ and $m = 1.3$ resulted in clusters which were relatively distinct, had intuitive descriptions and also contained an adequate number of bills for machine learning. Additionally, we experimented with the threshold values for cluster membership and settled on 0.15. That is, a bill is considered a member of a cluster if its membership value is above 0.15. Table 2 describes the results of our chosen FCM parameters. Each cluster is assigned a shorthand name describing its contents and is called a sphere of legislation in this paper. We next describe the model.

^[2]See <https://www.congress.gov/help/field-values/policy-area>.

Table 2 Summary of four spheres of legislation: Shorthand names and descriptions for each of the spheres of legislation identified by the FCM algorithm are shown here. The sum of the number of bills across all spheres (965) is greater than the total number of bills (722) because membership is fuzzy. Spheres 1 and 2 are relatively distinct from the rest, while Spheres 3 and 4 share a large number of bills.

Sphere#	Size	Name of Sphere	Sampling of Bill Subjects	Ovlp. 1	Ovlp. 2	Ovlp. 3	Ovlp. 4
1	105	Security & Armed Forces	Armed forces and national security (77), Emergency management (11), Transportation and public works (10)		7%	20%	20%
2	263	Economics & Finance	Economics and public finance (263)	3%		0%	0%
3	284	Energy & Infrastructure	Energy (69), Education (31), Taxation (28), Transportation and public works (27)	7%	0%		76%
4	313	Public Welfare	Health (52), Crime and law enforcement (43), Taxation (38), Education (31)	7%	0%	69%	

3 The LIG Model

We represent the senate influence network as a linear influence game (LIG) [IO11, IO14], one type of 2-action graphical game [KLS01]. Nodes represent senators, or *players*, and are connected by directed edges. Edge weights represent the influence exerted by the source node upon the target. Influence weights can be negative, positive, or zero. The directed edges are allowed to be asymmetric, meaning nodes A and B may exert different levels of influences on each other. Additionally, nodes have a threshold level, which represents “stubbornness.” Nodes with thresholds further from zero are more resistant to change. Absent influences, a node with negative threshold is predisposed to adopting action 1 (*yea* vote), and a node with positive threshold is predisposed to -1 (*nay* vote). The matrix of influence weights $\mathbf{W} \in \mathbf{R}^{n \times n}$ and the threshold vector $\mathbf{b} \in \mathbf{R}^n$ constitute the LIG model. The action $x_i \in \{1, -1\}$ chosen by each node i is the outcome of the model, as described below in game-theoretic terms.

Each node’s *best response* to other nodes’ actions depends on the net incoming influence and the node’s threshold. When the total incoming influence from nodes playing 1 minus the total incoming influence from nodes playing -1 exceeds the node’s threshold level, that node’s best response is 1. If below, it is -1 ; in the case of a tie, the node is indifferent and can play either. Note that the best responses of the nodes are interdependent. A vector of *mutual best responses* of all the nodes is a stable outcome of the model, formally known as a *pure strategy Nash equilibrium* (PSNE). It is stable because no node has any incentive to deviate from it. The LIG model adopts PSNE to represent stable collective outcomes from a complex network of influence. Before formally defining the technical terms, we illustrate the model using an example.

Example. Figure 3 illustrates the LIG model with a simple, 4-node example. Note that the LIG model allows edges of opposite polarities between two nodes. This is not shown in this example for simplicity. As explained in Figure 3, A and B playing 1 and C and D playing -1 is a PSNE, whereas all nodes playing 1 is not a PSNE.

As shown for node A in the above example, the process of adding up incoming influences from nodes playing 1, then subtracting influences from nodes playing -1 , and finally comparing the result with the threshold value is succinctly captured by the *influence function* defined in Definition 3.1 below. The best response calculation (e.g., node A’s best response is to play 1 if the total weighted influence on A exceeds its threshold) can be done using the *payoff function* defined in Definition 3.2. Finally,

Figure 1 LIG example. A four-node LIG is shown here. The directed edges are labeled with influence levels. Any absence of a directed edge implies an influence level of 0. Threshold values of 0 (for simplicity) are shown with a connector to each node. We assume binary actions $\{1, -1\}$. In this game, nodes A and B playing 1 and nodes C and D playing -1 is a pure strategy Nash equilibrium (PSNE). To see this, consider node A first. We add up the incoming influences from those nodes (in this case, B) that are playing 1 and then subtract from it the influences coming from nodes (in this case, C and D) playing -1 . We get $1 - (-2 - 1.5) = 4.5$, which is basically the total *weighted* influence on A. Since 4.5 is greater than A's threshold of 0, A's best response is 1. Similarly, it can be shown that B, C, D's best responses are 1, -1 , -1 , respectively, and therefore, this is a PSNE. Similarly, nodes A and B playing -1 and C and D playing 1 is another PSNE. As a negative example, all nodes playing 1 is not a PSNE. To see this, consider node A. The total weighted influence on A is $1 + (-2) + (-1.5) = -2.5$, which is less than A's threshold of 0. Therefore, A's best response is to play -1 , which violates the mutual best response condition for PSNE.

PSNE is formally defined in Definition 3.3. In the following formal definitions, we use the same notation as [IO14].

Definition 3.1 (Linear Influence Game (LIG) [IO14]) In *LIG*, the influence function of each individual i , given others' actions $\mathbf{x}_{-i}$, is defined as $f_i(\mathbf{x}_{-i}) \equiv \sum_{j \neq i} w_{ij} x_j - b_i$ where for any other individual j , $w_{ij} \in \mathbb{R}$ is a *weight* parameter quantifying the “influence factor” that j has on i , and $b_i \in \mathbb{R}$ is a *threshold* parameter for i 's level of “tolerance.”

Here, individuals receive influences from other players and have an influence threshold of their own, which accounts for their own resistance to external influence. The influence function f_i calculates the weighted sum of incoming influences on i , as described in the paragraph above Definition 3.1, and subtracts i 's threshold from it.

Example. In the LIG shown in Figure 3, when B plays 1 and C and D play -1 , the influence function of A is $1 \times 1 + (-1) \times (-2) + (-1) \times (-1.5) - 0 = 4.5$. In contrast, when B, C, and D play 1, the influence function of A is $1 \times 1 + 1 \times (-2) + 1 \times (-1.5) - 0 = -2.5$. Note that the influence function of A does not depend on A's action.

We next define the payoff of each player. The payoff function happens to be one of the main ingredients of any game-theoretic model.

Definition 3.2 (Payoff Function [IO14]) For an LIG, we define the *payoff* function $u_i : \{-1, 1\}^n \rightarrow \mathbb{R}$ as $u_i(x_i, \mathbf{x}_{-i}) \equiv x_i f_i(\mathbf{x}_{-i})$, where $\mathbf{x}_{-i}$ denotes the vector of a joint-action of all players except i and f_i is defined in Definition 3.1.

The payoff function quantifies the preferences of the players based on the actions of other players. Given the action of all other individuals $\mathbf{x}_{-i}$ and influence function $f_i(\mathbf{x}_{-i})$, an individual will prefer to choose either 1 or -1 as follows. When $f_i(\mathbf{x}_{-i})$ is negative, $x_i = -1$ will result in a positive payoff; when $f_i(\mathbf{x}_{-i})$ is positive, $x_i = 1$ will result in a positive payoff. Actions chosen in this fashion in order to result in a positive payoff (i.e., to maximize payoff) is defined as the *best response*.

Example. For the LIG shown in Figure 3, when A and B play 1 and C and D play -1 , A's payoff is $1 \times 4.5 = 4.5$. In this scenario, A is playing its best response, because if A were to play -1 , A's payoff would have been -4.5 . As another example, when everyone plays 1, A's payoff is $1 \times (-2.5) = -2.5$. Here, A is not playing its best response, because A could have gotten a payoff of 2.5 by switching to action -1 . Note that the payoff of a node does depend on the node's own action.

We next define pure-strategy Nash Equilibrium (PSNE) of an LIG. PSNE is one of the most central solution concepts in game theory. A PSNE signifies everyone playing their best responses simultaneously.

Definition 3.3 (Pure-Strategy Nash Equilibrium [IO14]) A *pure-strategy Nash equilibrium (PSNE)* of an LIG $\mathcal{G}$ is an action assignment $\mathbf{x}^* \in \{-1, 1\}^n$ that satisfies the following condition. Every player i 's action x_i^* is a simultaneous best-response to the actions $\mathbf{x}_{-i}^*$ of the rest.

Example. In our running example (Figure 3), nodes A and B playing 1 and nodes C and D playing -1 is a PSNE, because it can be verified that every player is playing their best response simultaneously. As another example, nodes A and B playing -1 and C and D playing 1 is also a PSNE. As shown in Figure 3, all nodes playing 1 cannot be a PSNE.

We adopt PSNE as the notion of stable outcomes arising from a network of influence. We are interested in questions like how the network changes based on the spheres of legislation and what impact the spheres have on polarization and most influential nodes. For these, we learn the networks using the spheres data.

4 Machine Learning

We use Honorio and Ortiz's machine learning algorithm in order to instantiate an LIG from raw roll-call data [HO15]. The goal of the algorithm is to capture as much of the ground-truth data as possible as PSNE (the *empirical proportion of equilibria*), without having so many total PSNE (the *true proportion of equilibria*) that

the model is meaningless. For example, if all influence weights and threshold levels are 0 (i.e., $\mathbf{W}=0$, $\mathbf{b}=0$), then all 2^n possible joint actions among n players would be PSNE, trivially covering all observed voting data. However, this is undesirable as it has no predictive power at all. Therefore, we would like to maximize the empirical proportion of equilibria while minimizing the true proportion. Following is a gist of Honorio and Ortiz's machine learning algorithm resulting from a very lengthy proof [HO15].

4.1 Learning Algorithm

To balance the true and empirical proportions of equilibria, the learning algorithm uses a generative mixture model that picks a joint action which is either a PSNE or non-PSNE of an LIG model $\mathcal{G}$ with probabilities q and $1 - q$, respectively. Of course, our goal is to learn the game $\mathcal{G}$. Let $\mathcal{NE}(\mathcal{G})$ denote the set of PSNE of $\mathcal{G}$ and $\mathcal{D} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(m)}\}$ be the dataset of m voting instances. The empirical proportion of equilibria, $\hat{\pi}(\mathcal{G})$ is the fraction of data captured as PSNE of $\mathcal{G}$. This is formally defined as follows, where $\mathbb{1}$ is the indicator function returning 1 if the condition is true, 0 otherwise.

$$\hat{\pi}(\mathcal{G}) \equiv \frac{1}{m} \sum_{\mathbf{x} \in \mathcal{D}} \mathbb{1}[\mathbf{x} \in \mathcal{NE}(\mathcal{G})].$$

The true proportion of equilibria, denoted by $\pi(\mathcal{G})$, is the fraction of all joint actions among n players that are PSNE, regardless of their existence in the voting instance data. This can be expressed as:

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

Given a set of voting instances $\mathcal{D}$, the average log-likelihood of the probabilistic generative model can be written as follows. Here, KL stands for the Kullback-Liebler divergence [CT06, Ch 2].

$$\hat{\mathcal{L}}(\mathcal{G}, q) = KL(\hat{\pi}(\mathcal{G}) \parallel \pi(\mathcal{G})) - KL(\hat{\pi}(\mathcal{G}) \parallel q) - n \log 2.$$

Leaving the rigorous mathematical proof [HO15] aside, we can intuitively see how maximizing the above log-likelihood achieves maximization of the empirical proportion of equilibria $\hat{\pi}(\mathcal{G})$ relative to the true proportion of equilibria $\pi(\mathcal{G})$. For this, note that the first term above, $KL(\hat{\pi}(\mathcal{G}) \parallel \pi(\mathcal{G}))$, is maximized by a game $\mathcal{G}$ that makes $\hat{\pi}(\mathcal{G})$ as big as possible while making $\pi(\mathcal{G})$ as small as possible. In other words, the game should capture as much of the data as possible as PSNE while keeping its total number of PSNE as small as possible.

Furthermore, the second term, $-KL(\hat{\pi}(\mathcal{G}) \parallel q)$ becomes 0 when $\hat{\pi}(\mathcal{G}) = q$. This indicates that the optimal mixture parameter q is $\hat{\pi}(\mathcal{G})$. This leaves learning $\mathcal{G}$ to maximize $KL(\hat{\pi}(\mathcal{G}) \parallel \pi(\mathcal{G}))$ as the main task, because we are maximizing the log-likelihood over all choices of $\mathcal{G}$ and q . The main challenge here is dealing with $\pi(\mathcal{G})$ due to the hardness of computing PSNE [IO14]. However, it can be shown that with high probability, maximizing a lower bound of the log-likelihood is equivalent

to maximizing $\hat{\pi}(\mathcal{G})$ over all choices of $\mathcal{G}$. This is equivalent to minimizing $1 - \hat{\pi}(\mathcal{G})$, which leads to the following loss minimization formulation.

$$\min_{\mathbf{w}, \mathbf{b}} \frac{1}{m} \sum_l \max_i \ell[x_i^{(l)}(\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} - b_i)].$$

Above, the *loss function* ℓ represents the errors in best responses. It is easy to explain the above by using the 0/1 loss function $l(z) \equiv \mathbb{1}[z < 0]$. Whenever any player in the l -th voting instance does not play its best response, $\max_i \ell[x_i^{(l)}(\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} - b_i)]$ is 1. When all players play their best responses, then $\max_i \ell[x_i^{(l)}(\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} - b_i)] = 0$, signifying a PSNE. For practical purposes of optimization, instead of the 0/1 loss function, a continuous loss function like the logistic loss function is used.

The final optimization problem is the following:

$$\min_{\mathbf{w}, \mathbf{b}} \frac{1}{m} \sum_l \max_i \ell[x_i^{(l)}(\mathbf{w}_{i,-i}^T \mathbf{x}_{-i}^{(l)} - b_i)] + \rho \|\mathbf{w}\|_1.$$

Here, m is the number of bills, ℓ is the typical logistic loss function, and ρ is an l_1 regularization parameter controlling the number of edges $\|\mathbf{w}\|_1$. That is, we prefer sparser networks if the solution quality is not degraded too much.

We solve the above optimization for each sphere of legislation and obtain an influence network. While doing this, we rigorously cross validate to avoid overfitting or underfitting as described in the next section.

4.2 Cross-Validation and Model Selection for LIG

In order to make use of the l_1 -regularized model, we must choose a regularization parameter ρ . High values of ρ assign a higher penalty to the number of edges in the graph and result in a sparser graph, while low values of ρ assign a lower penalty and result in a denser graph. While low values of ρ will be better fitted to the model, there is a risk of over-fitting – “memorizing” the data – which results in poor predictive performance on new data.

Additionally, the number of edges must be taken into consideration because the problem of computing equilibria is NP-hard [IO11, IO14]. In fact, it is likely that an extremely complex model would have so many edges that equilibria computation would never finish within a reasonable time-frame of several days. However, an exceedingly low number of edges would lead to an under-fit model, and could not be generalized to new data. Therefore, we must pick a ρ value which strikes a balance between computation time and the risks of over- and under-fitting.

We use cross-validation (CV) in order to determine the effectiveness of a given ρ value. In CV, a process essential to most machine learning applications, data is partitioned into two sets: training and validation. The model is trained using the training set and then employed to make predictions against the validation set. The performance of the model is measured by the error in the training and validation set. When a model is over-fit, validation error will be significantly higher than training error. When a model is under-fit, both validation and training error will be high. In

CV, researchers adjust the parameters of the machine learning algorithm in order to create the best model which neither under- nor over-fits the data.

With large datasets, training and validation sets are often created by splitting the data in half, or holding out some smaller proportion of the data.^[3] However, the four datasets generated by clustering method are too small to form informative predictions if they are further reduced by this straightforward partitioning. Instead, we used k -fold CV, which leverages re-sampling in order to form useful insights on small datasets. In k -fold CV, the dataset is randomized and split into k partitions. In one run of the k -fold CV, one of the k sets is chosen as the validation set, while the remaining $k - 1$ sets are combined to form the train set. On the next run, a different set is chosen as the validation set, and the others are used to train the model. Measures of accuracy and error from each run are averaged across the k runs. Choosing a value of k is arbitrary, but $k = 10$ is often used in research applications, colloquially known as 10-fold CV.

We ran 10-fold CV on each sphere with $0 < \rho < 0.01$, tracking three measures of model performance:

- 1 Number of edges in the training set graph
- 2 Best response (BR) error, or the percentage of senators not playing their best response, in training and validation sets
- 3 q , the proportion of votes recorded as PSNE, in training and validation sets.

We chose ρ values, shown in based on the following goals:

- 1 The graph is sparse enough to efficiently compute equilibria
- 2 The model neither over-fits nor under-fits the data (i.e., BR error is low, and the differences between training and validation sets for BR error and q are low)
- 3 The proportion of observed roll-call votes that are PSNE (q) according to the learned model is high

We next present the cross-validation results for Sphere 1.

Cross-Validation on Sphere 1 (Security & Armed Forces). As shown in Figure 2, the number of edges drastically decreases until $\rho = 0.000367$, then begins to decrease at a slower rate, reaching a reasonable number of edges between values of 0.002424 and 0.003455. BR error in both the training and validation set remains low until $\rho \Rightarrow 0.004$, then begins to increase, showing that the model performs well until that point. Until $\rho = 0.001512$, the drastic difference between training and validation q values shows that the model is over-fit, and the regression to $q = 0$ when $\rho > 0.007014$ shows that the model is under-fit. Between values of 0.002154 and 0.03455, all metrics are within an acceptable range.

While we leave the detailed cross-validation results for the other spheres to Appendix B, there are a lot of similarities among these results. Across all spheres, when $\rho = 0$, the learned model is basically memorizing the training data as training error is 0, validation error is relatively high, and the proportion q of data captured as PSNE is drastically higher for the training set than the validation set. This is

^[3]In previous LIG research on congressional networks, models have been trained on one session of a congressional term, and validated on the other. However, because our data combines two terms (i.e., four sessions), this was not practical.

the over-fitting regime. As ρ increases, validation and training errors begin to converge, as do the validation and training q values. At higher ρ values, validation and training errors are both prohibitively high and the learning enters the under-fitting regime.

Table 3 summarizes the ρ values that we have chosen according to the three criteria listed above. We use these values of ρ to produce the LIG models used throughout the rest of the paper.

Table 3 Chosen values of ρ for each sphere and the corresponding number of edges and the average best response (BR) error of validation sets.

Sphere	ρ	# Edges	BR Error (Validation)
1	0.002728	1191	7.07%
2	0.003888	1071	5.05%
3	0.003070	1280	5.56%
4	0.003888	1076	5.10%

5 Polarization in Context

Visualization of the machine learned networks clearly shows that the network structure varies according to the spheres of legislation. In all spheres, however, the force-directed drawing algorithm automatically distinguishes Republicans from Democrats. Figures 3 and 4 depict the LIG visualizations for Spheres 1 (Security & Armed Forces) and 2 (Economics & Finance) as representative examples. The visualizations for the remaining spheres can be found in Appendix F.3. In this section, we discover different degrees of polarization across the spheres by investigating cross-party (or cross-border) edges, influence weights and thresholds, and modularity measures. We begin with cross-party edges.

5.1 Cross-Party Edges

The boundary between the two parties is interesting for studying polarization. Even though negative edges more often occur at the boundary, the connectivity between the two parties varies a lot according to the spheres of legislation. These are depicted in Figures 5 and 6 for Spheres 1 and 2, respectively (others are in Appendix F.4).

Figure 6 shows the cross-party edges in Sphere 2 (Economics & Finance), which starkly contrasts those of Sphere 1 (Security & Armed Forces) shown in Figure 5. In Sphere 2, only 12 of the strongest 40% of edges are between members of different parties. Of these, 2/3 are negative, suggesting a very polarized network. Aside from two positive influences between Maine senators King (a left-leaning Independent) and Collins (a center-leaning Republican), the remaining two positive connections are the weakest of all connections shown for this sphere.

Similarly, examining inter-party edges reveals that Sphere 3 (Energy & Infrastructure) is also very polarized. While there are many edges between both parties in this network, about 70% of them are negative. Positive influences come from a few sources, again including the centrist Senator Collins. Incongruously, prominent right-wing senator Tom Cotton (R-AR) also exhibits positive influences with democratic senators. However, most other far-left or far-right leaning senators, including Sanders (I-VT) and Cruz (R-TX) only exhibit negative influences with the opposite party.

Sphere 4 (Public Welfare)’s inter-party edges strike a balance between the polarities exhibited by the previous three spheres. There are slightly more positive edges (9) than negative edges (7), but still a low number of edges overall. Again, there are positive influences between Maine senators King (I-ME) and Collins (R-ME), but also positive influences between Senator McConnell and Democratic senators King (D-ME) and Tester (D-MT).

Overall, each sphere exhibits some level of polarization, but some spheres are far more polarizing than others. Some senators are present in every sphere's inter-party boundary, whether for positive or negative influences. Maine Senators Collins (R) and King (D) often share positive influences with each other, as well as other senators. Senator Lee (R-UT), a conservative libertarian, always exhibits negative edges with members of the other party, although in Sphere 1, he also shares positive influences with senators Harris (D-CA) and Feinstein (D-CA). Meanwhile, left-wing icon Bernie Sanders (I-VT) exhibits the equivalent behavior, with only negative cross-party edges in all spheres *except* Sphere 1. These results suggest that Sphere 1 (Security & Armed Forces) is least polarized, whereas Spheres 2 (Economics & Finance) is highly polarized.

5.2 Influence Weights and Thresholds

We now take a closer look at the influence weights and thresholds of the machine learned models, beyond just the cross-party edges. Figure 7 shows a histogram of four different categories of edge weights: in the top row, Democrat to Democrat and Republican to Republican, and in the bottom row, Democrat to Republican and Republican to Democrat (note that the edges are directed). In each plot, the histograms for the four spheres are superimposed for the purpose comparison. For the intra-party edges (D-D and R-R), Spheres 2, 3, and 4 have very similar histograms and they are different from the histogram of Sphere 1 (Security & Armed Forces). At the peak, the number of intra-party edges in Sphere 1 is dominated by the other

spheres. However, for higher edge weights, Sphere 1 dominates the other spheres. This indicates that there are stronger D-D and R-R influences in Sphere 1 compared to the other spheres, which in turn may indicate more polarization in Sphere 1. Interestingly, if we look at the cross-party edges (D-R and R-D), we can see that Sphere 1 again dominates the other spheres in the positive influence weights regime. Note that in the bottom row of Figure 7, the peak of Sphere 3 dominates that of Sphere 1, but Sphere 3's peak is in the negative influence regime, whereas Sphere 1's peak is in the positive influence regime.^[4] All of these indicate that there are more positive influences *within and across* the two parties in Sphere 1 compared to the other spheres, which contributes to Sphere 1 being less polarized.

Of course, the influence weights cannot be read alone without considering thresholds, because the game-theoretic model accounts for both of these in predicting stable outcomes. Recall that the threshold *magnitude* signifies stubbornness or resistance to influence. More positive threshold values resist positively weighted influences by leaning to play -1 in the presence of (1) positive influence from those playing 1 and (2) negative influence from those playing -1 (in both cases, a neighbor's

^[4]Also note that, as shown in Table 3, Sphere 3 has the highest number of edges among all the spheres.

action times the influence from that neighbor is positive). More negative threshold values resist negatively weighted influences in a similar fashion.

Figure 8 shows the threshold histograms for the two parties. The most interesting aspect of these histograms is that for both Democratic and Republican senators, the threshold distribution is “flatter” in Sphere 1 (Security & Armed Forces) compared to the other spheres. This indicates that for both parties, the thresholds are more “uniformly distributed” in Sphere 1 than in the other spheres. In contrast, in Spheres 2 (Economics & Finance), 3 (Energy & Infrastructure), and 4 (Public Welfare), the threshold values of each party are concentrated one region, which indicates the similarity among the senators belonging to the same party. Together with negative cross-party edges and positive intra-party edges, this contributes to polarization in these spheres.

5.3 Modularity

Furthermore, a formal study of polarization rooted in network science produces similar results. Modularity [GN02, NG04, New06] has been widely used as a measure of polarization in networks. We apply the following definition of modularity derived for directed networks with signed weights [GJA09].

$$Q = \frac{1}{2w^+ + 2w^-} \sum_i \sum_j \left[w_{ij} - \left(\frac{w_i^{+,out} w_j^{+,in}}{2w^+} - \frac{w_i^{-,out} w_j^{-,in}}{2w^-} \right) \right] \times \delta(C_i, C_j).$$

Here, w_{ij} is the weight of edge i to j , $w_{ij}^+ = \max\{0, w_{ij}\}$, $w_{ij}^- = \max\{0, -w_{ij}\}$, and $2w^\pm$ is the total weight of all positive or negative edges, expressed by $\sum_i \sum_j w_{ij}^\pm$. Furthermore, $w_i^{\pm,out}$ is the weighted out-degree $\sum_k w_{ik}^\pm$ and $w_j^{\pm,in}$ is the weighted in-degree $\sum_k w_{kj}^\pm$. The Kronecker delta function $\delta(C_i, C_j)$ is 1 if i and j belong to the same party; it is 0 otherwise.

Applying this definition, We obtain the following modularity scores for the four spheres of legislation respectively: 0.7861, 0.8904, 0.8724, and 0.8857 (see Table 5.) This shows that Sphere 1 (Security & Armed Forces) is least polarized and Spheres

2 (Economics & Finance), 3 (Energy & Infrastructure), and 4 (Public Welfare) are much more polarized.

6 Most Influential Nodes in Context

There exists a number of centrality measures that are derived from a structural analysis of networks [Jac10]. However, our model is behavioral where nodes adopt their best responses to each other. In a strictly game-theoretic model of behavior, a set of nodes will be called most influential *with respect to achieving a desirable stable outcome* if their choice of actions leads the whole system of influence to that desirable stable outcome [IO11, IO14]. Here, a crucial aspect is a desirable stable outcome, represented by a PSNE. For example, let us say that our desirable stable outcome is to pass a bill by a 100–0 vote. A set of senators will be called most influential if their voting together influences every other senator to also vote for the bill, thereby having the desirable stable outcome as the *unique* PSNE outcome. This concept can be extended to other types of desirable stable outcomes like blocking a bill unanimously, passing a bill with at least 60 votes, forcing/avoiding a filibuster, etc. When there are multiple most influential sets, we naturally prefer smaller sets of most influential nodes.

The above concept of most influential nodes is centered around stable or PSNE outcomes. As we will see in Section 6.2, it requires computation of all PSNE. We next outline how we compute all PSNE for each sphere of legislation.

6.1 PSNE Computation

Once the LIG is instantiated by the machine learning algorithm (see Section 4), we can compute the set of all PSNE using the algorithm described in [IO14]. This is a backtracking search algorithm which takes advantage of the graph’s structure. We give a brief overview below.

The algorithm begins by selecting the node with the highest out-degree—the node which directly influences the most other nodes—and assigns it the action 1. It progressively selects new nodes and assigns them the action 1 until all nodes are assigned actions without any contradiction (indicating a PSNE) or it encounters a contradiction that guarantees that there is no PSNE with the actions assigned so far. It then revisits the most recent node and changes its action from 1 to -1 . After this, the algorithm again tries to make progress. In general, at any stage of the algorithm, we have a *partial joint action*, which is the action (1 or -1) of each node selected so far. If some node in the network is not playing its best response, the partial joint action cannot lead to a PSNE. When this occurs, the algorithm tries a different action for the most recently selected node v if it has not already done so. If trying a different action for v still leads to a contradiction, the algorithm backtracks by deselecting v and changing the action of the node that had been selected before v . When every node is playing their best response with respect to each other, we have reached a PSNE. Importantly, the algorithm always tries to reach a contradiction so that it can reduce the overall computation time by pruning large parts of the search tree. This process repeats until *all possible* PSNE have been found.

For each sphere, we ran the algorithm on Bowdoin College’s high-performance computing (HPC) grid. The number of PSNE created for each sphere’s LIG given

our chosen ρ values are summarized in Table 4. Note that the number of PSNE is a tiny fraction of the 2^{103} possible joint actions. These sets of all PSNE are necessary to compute the most influential senators, which we describe next.

Table 4 ρ values selected through cross validation and the corresponding PSNE counts across different spheres.

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
ρ	0.002728	0.003888	0.003070	0.003888
# PSNE	865,578	6,454,013	8,711,782	4,162,629

6.2 Computing Most Influential Nodes

Algorithmically, the most influential nodes problem asks for selecting a minimum set of nodes such that when they choose their actions according to the desirable stable outcome (e.g., voting *yea* when the desirable stable outcome is passing a bill unanimously), the desirable stable outcome becomes the only possible PSNE. An approximation algorithm for computing most influential senators was given by Irfan and Ortiz [IO14], which produces a directed acyclic graph (DAG). The algorithm requires precomputation of all PSNE, which is a provably hard problem [IO14]. We apply Irfan and Ortiz’s PSNE computation algorithm to the LIG for each sphere of legislation. Having computed all the PSNE, we then compute the DAG representing most influential sets of nodes. Figures 9 and 10 show the results of the most influential nodes algorithm for Spheres 1 and 3 respectively, where the desirable stable outcome is to achieve the most number of *yea* votes possible in any PSNE (that is, to gain the most support possible from the legislative body according to our model). The way to read Figures 9 and 10 is to inspect each DAG and find a top to bottom path. Each of these paths gives a most influential set.

The sets of most influential senators in each sphere support the inferences gained from analyzing the LIG networks. As illustrated in Figure 9, in Sphere 1 (Security & Armed Forces), 4 Republicans and 4 Democrats comprise a set of 8 most influential senators. In other words, 8 senators and more importantly, the balanced bipartisan groups of 8 senators shown in Figure 9 are sufficient to generate the maximum possible support for a bill in Sphere 1. As shown in Figure 10, in Sphere 3 (Energy & Infrastructure), 5 Republicans and 6 Democrats comprise a set of 11. This suggests that Sphere 3 is more polarized than Sphere 1, since it requires a larger body of influencing senators. The DAGs for the other spheres are shown in Appendix E.

6.3 Game-Theoretic Vs. Structural Centrality Measures

In the above game-theoretic formulation of most influential nodes, we find that each set of most influential senators across all spheres is comprised of an (almost) equal number of Democrats and Republicans. This signifies the need for bipartisan support to guarantee passing a bill with the maximum possible support under the PSNE constraints. As we show next, this also happens to be a distinguishing feature between game-theoretic and structural measures. Table 5 shows various centrality measures and other quantities computed for each sphere.

First, measures like diameter, average shortest path length and clustering coefficient reveal some, but not many, differences among the spheres. The network

Figure 9 Most influential nodes for Sphere 1 (Security & Armed Forces): Directed acyclic graphs (DAGs) representing sets of most influential nodes. Any top to bottom path gives a most influential set. Here, 4 Republicans and 4 Democrats are most influential.

diameter of Sphere 3 (Economics & Finance) is 4, and the network diameter of every other sphere is equal to 5. The average shortest path lengths between all four spheres are similar to one another, ranging between 2.2295 and 2.5476. Being close to half the size of the network diameter, these values suggests that most nodes in the network are well-connected, though not all. The average clustering coefficient is a measure of the density of triangles in a network. In more polarized networks, we might expect this value to be high, because senators who are closely aligned on partisan issues would be well-connected with each other. In each sphere, the average clustering coefficient is similar, but lower in Spheres 1 and 3 (0.0187 and 0.0174, respectively) than in Spheres 2 and 4 (0.0206 and 0.0218, respectively). These measures, however, do not give a direct indication of polarization, at least not as much as the modularity measure. We discussed the modularity values in Section 5.

We now focus on the widely applied structural measures of centrality. For each sphere, we show the top 10 most central senators with respect to four centrality measures: degree, closeness, betweenness, and eigenvector. The simplest measure is degree centrality, or the number of nodes each node is connected to (normalized by the maximum possible degree, $N - 1$ or 102 in our case). The next form is closeness centrality, or how close a node is, on average, from every other node in the network. The third form is betweenness centrality, which is the average number of times the node is present along the shortest path from any other two nodes. The final form is eigenvector centrality, which has a self-referential definition accounting for the centrality of a node's neighbors.

Most notably, these centrality measures do not capture the strategic aspects of behavior. Throughout most measures, Republican senators are over-represented, comprising the majority of the top ten most central nodes. *In contrast, the game-*

Figure 10 Most influential nodes for Sphere 3 (Energy & Infrastructure): The interpretation is similar to Figure 9. Here, 5 Republicans and 6 Democrats are most influential.

theoretic measure gives a balanced coalition between Democrats and Republicans. This is important because when networks are polarized, achieving a desirable stable outcome requires support from both sides.

7 Towards Richer Models: Ideal Point Models with Social Interactions

We also apply a richer model of influence recently proposed by Irfan and Gordon [IG18] that extends the LIG model by incorporating ideal points of the senators and polarities of bills [Poo08]. Their work showed the value of combining game-theoretic and statistical models for studying strategic interactions in context, but they assume the network to be fixed, regardless of the bill context. We use their model and allow the network to change based on the spheres of legislation. We also perform an analysis of the networks learned.

As a cautionary note, the way Irfan and Gordon’s model [IG18] combines networks with ideal points makes it difficult to disentangle the two. Analyzing the networks alone may be inconclusive, because ideal points also supply the model with predictive power. Moreover, the machine learning algorithm learns these two components simultaneously. With this caveat in mind, we give an analysis of the networks and the ideal points learned.

Analysis of Influence Networks. Figures 11 and 12 show the learned networks for Sphere 1 (Security & Armed Forces) and 2 (Economics & Finance) under this richer

Table 5 Network analysis of learned influence networks for different spheres of legislation. Various centrality measures and network-level properties are shown.

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
Number of Edges	1191	1071	1280	1076
Network Diameter	5	5	4	5
Avg. (Shortest) Path Length	2.2295	2.5132	2.1506	2.5476
Avg. Clustering Coefficient	0.1867	0.2057	0.174	0.2176
Modularity	0.7861	0.8904	0.8724	0.8857
Degree Centrality				
Degree (1)	0.5784: LEE R-UT	0.3529: TOOMEY R-PA	0.3725: LANKFORD R-OK	0.3725: COTTON R-AR
Degree (2)	0.4216: PAUL R-KY	0.3333: PERDUE R-GA	0.3627: SASSE R-NE	0.2941: LEAHY D-VT
Degree (3)	0.4118: SANDERS I-VT	0.3235: ENZI R-WY	0.3529: WARNER D-VA	0.2843: CAPITO R-WV
Degree (4)	0.3824: MORAN R-KS	0.3137: LANKFORD R-OK	0.3333: CAPITO R-WV	0.2843: MURKOWSKI R-AK
Degree (5)	0.3725: MANCHIN D-WV	0.3137: YOUNG R-IN	0.3235: TOOMEY R-PA	0.2745: AYOTTE R-NH
Degree (6)	0.3627: RUBIO R-FL	0.3039: COTTON R-AR	0.3235: MURKOWSKI R-AK	0.2745: SHELBY R-AL
Degree (7)	0.3529: CRUZ R-TX	0.2941: CASEY D-PA	0.3235: COTTON R-AR	0.2647: PERDUE R-TN
Degree (8)	0.3333: ALEXANDER R-TN	0.2843: CASSIDY R-LA	0.3137: BROWN D-OH	0.2549: ALEXANDER R-TN
Degree (9)	0.3333: ENZI R-WY	0.2843: WICKER R-MS	0.3137: FEINSTEIN D-CA	0.2549: PETERS D-MI
Degree (10)	0.3235: LEAHY D-VT	0.2745: CORKER R-TN	0.3137: PAUL R-KY	0.2549: PAUL R-KY
Closeness Centrality				
Closeness (1)	0.5862: LEE R-UT	0.5126: PERDUE R-GA	0.5514: WARNER D-VA	0.5204: COTTON R-AR
Closeness (2)	0.5635: RUBIO R-FL	0.4951: COTTON R-AR	0.5426: LANKFORD R-OK	0.5178: KIRK R-IL
Closeness (3)	0.5574: PAUL R-KY	0.4766: COLLINS R-ME	0.5368: SASSE R-NE	0.4951: MURKOWSKI R-AK
Closeness (4)	0.5455: SANDERS I-VT	0.47: ENZI R-WY	0.5368: BENNET D-CO	0.4928: AYOTTE R-NH
Closeness (5)	0.5455: BALDWIN D-WI	0.4636: SASSE R-NE	0.534: KING I-ME	0.4766: SULLIVAN R-AK
Closeness (6)	0.5397: ENZI R-WY	0.4636: MANCHIN D-WV	0.5231: MURKOWSKI R-AK	0.4744: MCCONNELL R-KY
Closeness (7)	0.5368: MORAN R-KS	0.4615: FLAKE R-AZ	0.5178: COTTON R-AR	0.4722: PAUL R-KY
Closeness (8)	0.5285: CORKER R-TN	0.4595: SHELBY R-AL	0.5152: BROWN D-OH	0.4636: COLLINS R-ME
Closeness (9)	0.5258: CASEY D-PA	0.4595: YOUNG R-IN	0.5126: CASEY D-PA	0.4554: CAPITO R-WV
Closeness (10)	0.5231: DURBIN D-IL	0.4595: HEITKAMP D-ND	0.51: CARPER D-DE	0.4554: SANDERS I-VT
Betweenness Centrality				
Betweenness (1)	0.0696: LEE R-UT	0.0538: PERDUE R-GA	0.0278: LANKFORD R-OK	0.0685: SANDERS I-VT
Betweenness (2)	0.0362: PERDUE R-GA	0.0468: HEITKAMP D-ND	0.0272: SASSE R-NE	0.0641: WYDEN D-OR
Betweenness (3)	0.033: MORAN R-KS	0.0452: GILLIBRAND D-NY	0.0265: WARNER D-VA	0.0559: COTTON R-AR
Betweenness (4)	0.0314: KING I-ME	0.045: ENZI R-WY	0.0226: BENNET D-CO	0.0533: MARKEY D-MA
Betweenness (5)	0.0301: MANCHIN D-WV	0.0434: COLLINS R-ME	0.0206: MURKOWSKI R-AK	0.0532: ALEXANDER R-TN
Betweenness (6)	0.0288: DURBIN D-IL	0.0383: MERKLEY D-OR	0.0196: BROWN D-OH	0.0421: MURKOWSKI R-AK
Betweenness (7)	0.0283: RUBIO R-FL	0.0382: COTTON R-AR	0.0194: CORNYN R-TX	0.0381: PAUL R-KY
Betweenness (8)	0.0283: PAUL R-KY	0.0381: SANDERS I-VT	0.0193: SCHATZ D-HI	0.0376: SASSE R-NE
Betweenness (9)	0.0253: SANDERS I-VT	0.0344: LEE R-UT	0.0187: SANDERS I-VT	0.0347: HARRIS D-CA
Betweenness (10)	0.0249: CRUZ R-TX	0.034: TESTER D-MT	0.0185: WICKER R-MS	0.032: AYOTTE R-NH
Eigenvector Centrality				
Eigenvector (1)	0.271: LEE R-UT	0.2114: COTTON R-AR	0.181: WARNER D-VA	0.2029: KIRK R-IL
Eigenvector (2)	0.2291: SANDERS I-VT	0.2061: PERDUE R-GA	0.1703: CORNYN R-TX	0.2017: HOEVEN R-ND
Eigenvector (3)	0.228: PAUL R-KY	0.1866: SULLIVAN R-AK	0.1675: LANKFORD R-OK	0.1727: GARDNER R-CO
Eigenvector (4)	0.1894: BALDWIN D-WI	0.1865: ENZI R-WY	0.1567: BENNET D-CO	0.1724: PORTMAN R-OH
Eigenvector (5)	0.1892: RUBIO R-FL	0.183: YOUNG R-IN	0.1551: JOHNSON R-WI	0.1715: CAPITO R-WV
Eigenvector (6)	0.1696: ENZI R-WY	0.1758: THUNE R-SD	0.1541: SASSE R-NE	0.1706: COTTON R-AR
Eigenvector (7)	0.1681: BARRASSO R-WY	0.1734: WICKER R-MS	0.1525: MURKOWSKI R-AK	0.1706: ROBERTS R-KS
Eigenvector (8)	0.161: CASEY D-PA	0.1674: MORAN R-KS	0.1454: KING I-ME	0.1683: FISCHER R-NE
Eigenvector (9)	0.1607: MORAN R-KS	0.1653: JOHNSON R-WI	0.1426: COTTON R-AR	0.1682: MURKOWSKI R-AK
Eigenvector (10)	0.1602: MANCHIN D-WV	0.163: GARDNER R-CO	0.1379: BOOZMAN R-AR	0.1646: ISAKSON R-GA

model (other spheres are in Appendix F.3). First, it is evident that the two parties are not as clustered as they were in the LIG model (compare with Figures 3 and 4). Second, a closer look at the cross-party edges shows that there are a lot more negative edges between the two parties under this richer model than there are under the LIG model. We show the cross-party edges for Spheres 1 and 2 in Figures 11 and 12, respectively (others are in Appendix F.4). These two differences can be attributed to using ideal points to discriminate the behaviors of opposing senators.

Polarization Metric Based on Modularity. The modularity framework discussed in Section 5 yields scores of 0.5392, 0.6801, 0.6887, and 0.6229, respectively. Both the ideal point metric and modularity scores indicate that Spheres 2 (Economics & Finance) and 3 (Energy & Infrastructure) are most polarizing, whereas Sphere 1 (Security & Armed Forces) is least polarizing. Sphere 4 (Public Welfare) sits in between. These results are somewhat similar to our earlier conclusions based on LIG without using ideal points. We include a broader analysis of the learned influence networks in Appendix F.5.

Polarization Metric Based on Ideal Points. We now apply the well-known ideal point-based polarization metric (i.e., distance between the means of the two parties) [MPR16] to calculate polarization levels across the four spheres. The ideal point distributions for two of the spheres are depicted in Figures 15 and 16 (others are in Appendix F.2). Applying the ideal point-based polarization metric, we obtain

values of 0.754, 1.235, 1.126, and 0.889 for Spheres 1 to 4 respectively. Evidently, Sphere 1 is least polarizing with respect to the ideal point distributions alone. Note that in our computation, we have not used the scaled versions of ideal point distributions shown in Figures 15 and 16. Instead, we have the non-scaled, machine learned ideal points. The scaled ideal points are amenable to comparison, but we have observed similar results for non-scaled versions.

We conclude this section by reiterating an earlier point. Investigating the influence networks and ideal points separately does not give us the complete picture since the model combines these two components together to make predictions. Therefore, we should also combine them in a meaningful way to infer polarization. We leave this as future work. We also leave open an exploration of the most influential nodes problem under this richer model.

8 Concluding Remarks and Research Outlook

In this paper, we have studied the linear influence game (LIG) model in the context of four spheres of legislation. We have done a thorough network analysis of the machine learned models for each sphere. Our analysis shows that contrary to the popular notion that the U.S. Congress is overly polarized these days, the measure of polarization varies according to the spheres of legislation. In fact, the two opposing parties tend to come together when dealing with bills in Sphere 1 (Security & Armed Forces). Therefore, the notion of polarization should be contextualized with respect to the spheres of legislation.

We have also shown that across all the spheres, the LIG model predicts that a set of most influential senators consists of a bipartisan coalition (which also differentiates game-theoretic and structural centrality measures). Despite this shared property

among the four spheres, the number of senators required to form a most influential set varies. Sphere 1 happens to require the least number of senators in its most influential set to achieve a desirable outcome of garnering the maximum support possible for a bill (under PSNE constraints). Again, this signifies that Sphere 1 is least polarized among the four spheres.

In sum, the consideration of different spheres of legislation reveals interesting aspects of polarization and most influential senators in Congress. Building upon this study, following are some interesting future directions.

- 1 The most pressing task is to fully explore the ideal point model with social interactions [IG18] for different spheres of legislation. We have briefly touched upon it in Section 7. However, as we have mentioned in that section, finding a behavioral definition of polarization that can meaningfully combine different constituent parts of the model, such as ideal points of senators, polarity of bills, influence weights, and threshold values, remains an open problem.
- 2 In a similar vein, the notion of *context* provides another interesting direction. In this paper, we use spheres of legislation as a contextual platform for learning, analyzing, and comparing influence networks. The main idea here is that depending on the sphere, the influence network would be different. In contrast, the ideal point model with social interactions also has a contextual element in terms of polarities of bills and ideal points of senators, but it keeps the network fixed. How we can synthesize these two diverging ways of capturing context and thereby give a deeper meaning to context remains open.
- 3 A detailed comparative study of the most influential nodes for different spheres under the richer model [IG18] is another interesting direction. In particular, what happens to the balanced, bipartisan composition of most influential sets

Figure 13 Cross-party edges in Sphere 1 (Security & Armed Forces) under the richer model. Only 40% of the strongest cross-party edges are shown. 123 cross-party edges are positive, and 47 are negative.

Figure 14 Cross-party edges in Sphere 2 (Economics & Finance) under the richer model. Only 40% of the strongest cross-party edges are shown. 86 cross-party edges are positive, and 57 are negative.

under the LIG model (see Section 6 when we incorporate additional contextual parameters like ideal points and polarities)?

- 4 On the computational front, Irfan and Gordon [IG18] showed promising results on improving the time required to compute all PSNE. Extending those results to the spheres of legislation setting is another promising direction. It would also be interesting to investigate *why* their model leads to drastic improvement in computational time.
- 5 Considering different modeling frameworks is yet another exciting direction. A particularly promising framework is probabilistic graphical models (PGMs). Whereas we are currently constructing the spheres of legislation first and then learning the LIG models for each sphere, PGMs may allow us to do both at the same time. This approach would not require us to split the data. Finally, exploring the recently proposed semi-supervised learning for studying polarization [GTMB19] in game-theoretic settings is also an interesting direction.

Figure 15 Ideal points of Sphere 1 (Security & Armed Forces): The ideal point distributions of Democratic (blue +) and Republican (red x) senators, scaled linearly between -1 and 1 are shown. The distance between the mean ideal points of the two parties is 0.754 . It shows that Democratic and Republican senators are ideologically close to each other when it comes to national security.

Figure 16 Ideal points of Sphere 2 (Economics & Finance): The distance between the mean ideal points of the two parties is 1.235 , which shows more polarization compared to Sphere 1 shown in Figure 15.

Declarations

Availability of data and materials

All data used in this paper is publicly available at www.congress.gov.

Competing interests

The authors declare that they have no competing interests.

Funding

This research is partially supported by NSF Award IIS-1910203.

Author's contributions

Phillips and Irfan contributed by writing the manuscript. Phillips performed most of the experiments. Irfan synthesized and guided the project, revised all iterations of the manuscript, and performed some of the experiments. Ostertag-Hill made valuable contributions by performing most of the experiments in Section 7 and some in Section 5.

Acknowledgements

We are grateful to Drs. Honorio and Ortiz for letting us use their codes [HO15]. We are also thankful to the anonymous reviewers of COMPLEX NETWORKS 2019 for their suggestions.

Author details

¹Pillpack, Inc., 120 Beacon St, Somerville, Massachusetts 02143, USA. ²Department of Computer Science, Bowdoin College, 255 Maine St, Brunswick, Maine 04011, USA.

References

- AAAY+17. Al Amin, M.T., Aggarwal, C., Yao, S., Abdelzaher, T., Kaplan, L.: Unveiling polarization in social networks: A matrix factorization approach. In: IEEE INFOCOM 2017-IEEE Conference on Computer Communications, pp. 1–9 (2017). IEEE
- Bez81. Bezdek, J.C.: Pattern Recognition with Fuzzy Objective Function Algorithms. Plenum Press, ??? (1981)
- BMA15. Bakshy, E., Messing, S., Adamic, L.A.: Exposure to ideologically diverse news and opinion on facebook. *Science* **348**(6239), 1130–1132 (2015)
- CCC+11. Chen, W., Collins, A., Cummings, R., Ke, T., Liu, Z., Rincon, D., Sun, X., Wang, Y., Wei, W., Yuan, Y.: Influence maximization in social networks when negative opinion may emerge and propagate. In: Proceedings of the Eleventh SIAM International Conference on Data Mining, pp. 379–390 (2011)
- CF07. Christakis, N.A., Fowler, J.H.: The Spread of Obesity in a Large Social Network over 32 Years. *N Engl J Med* **357**(4), 370–379 (2007). doi:10.1056/NEJMs066082. <http://content.nejm.org/cgi/reprint/357/4/370.pdf>
- CF08. Christakis, N.A., Fowler, J.H.: The Collective Dynamics of Smoking in a Large Social Network. *N Engl J Med* **358**(21), 2249–2258 (2008). doi:10.1056/NEJMs0706154. <http://content.nejm.org/cgi/reprint/358/21/2249.pdf>
- CJR04. Clinton, J., Jackman, S., Rivers, D.: The statistical analysis of roll call data. *American Political Science Review* **98**(2), 355–370 (2004). doi:10.1017/S0003055404001194
- CRF+11. Conover, M.D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., Flammini, A.: Political polarization on twitter. In: Fifth International AAAI Conference on Weblogs and Social Media (2011)
- CT06. Cover, T.M., Thomas, J.A.: Elements of Information Theory. Wiley-Interscience, ??? (2006)

- DMSK15. Desmarais, B.A., Moscardelli, V.G., Schaffner, B.F., Kowal, M.S.: Measuring legislative collaboration: The Senate press events network. *Social Networks* **40**(1), 43–54 (2015). doi:10.1016/j.socnet.2014.07.006
- Far15. Farina, C.R.: Congressional polarization: Terminal constitutional dysfunction. *Colum. L. Rev.* **115**, 1689 (2015)
- GAS+15. Garcia, D., Abisheva, A., Schweighofer, S., Serdült, U., Schweitzer, F.: Ideological and temporal components of network polarization in online political participatory media. *Policy & Internet* **7**(1), 46–79 (2015)
- GB11. Gerrish, S.M., Blei, D.M.: Predicting Legislative Roll Calls from Text. *ICML'11 Proceedings of the 28th International Conference on Machine Learning*, 489–496 (2011)
- GB12. Gerrish, S.M., Blei, D.M.: How they vote: Issue-adjusted models of legislative behavior. *Advances in Neural Information Processing Systems* **25**(1), 2762–2770 (2012). doi:10.1111/j.1751-1097.2010.00727.x
- GJA09. Gómez, S., Jensen, P., Arenas, A.: Analysis of community structure in networks of correlated data. *Institut des Systèmes Complexes* **1**(1), 2–3 (2009)
- GMJCK13. Guerra, P.C., Meira Jr, W., Cardie, C., Kleinberg, R.: A measure of polarization on social media networks based on community boundaries. In: *7th International AAAI Conference on Weblogs and Social Media* (2013)
- GN02. Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proceedings of the national academy of sciences* **99**(12), 7821–7826 (2002)
- Gra78. Granovetter, M.: Threshold Models of Collective Behavior Source. *American Journal of Sociology* **83**(6), 1420–1443 (1978)
- GTMB19. Glonek, M., Tuke, J., Mitchell, L., Bean, N.: Semi-supervised graph labelling reveals increasing partisanship in the United States Congress. *arXiv*, 1–14 (2019)
- HO15. Honorio, J., Ortiz, L.: Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *The Journal of Machine Learning Research* **16**, 1157–1210 (2015). doi:10.1016/j.jasoc.2015.04.032
- IG18. Irfan, M.T., Gordon, T.: The power of context in networks: Ideal point models with social interactions. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 910–918 (2018). International Foundation for Autonomous Agents and Multiagent Systems
- IO11. Irfan, M.T., Ortiz, L.E.: A game-theoretic approach to influence in networks. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pp. 688–694 (2011). <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3746>
- IO14. Irfan, M.T., Ortiz, L.E.: On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence* **215**, 79–119 (2014). doi:10.1016/j.artint.2014.06.004
- Jac10. Jackson, M.O.: *Social and Economic Networks*. Princeton university press, ??? (2010)
- KKET03. Kempe, D., Kleinberg, J., Éva Tardos: Maximizing the spread of influence through a social network. In: *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2003)
- KKET05. Kempe, D., Kleinberg, J., Éva Tardos: Influential nodes in a diffusion model for social networks. In: *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)* (2005)
- KLS01. Kearns, M., Littman, M., Singh, S.: Graphical models for game theory. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pp. 253–260 (2001)
- LAH07. Leskovec, J., Adamic, L.A., Huberman, B.A.: The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* **1**(1), 5 (2007)
- LCWZ13. Li, Y., Chen, W., Wang, Y., Zhang, Z.-L.: Influence diffusion dynamics and influence maximization in social networks with friend and foe relationships. In: *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pp. 657–666 (2013). ACM
- Mac67. MacQueen, J.: Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* **1**: *Statist.*, 281–297 (1967)
- MPR16. McCarty, N., Poole, K.T., Rosenthal, H.: *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press, ??? (2016)
- New06. Newman, M.E.: Modularity and community structure in networks. *Proceedings of the national academy of sciences* **103**(23), 8577–8582 (2006)
- NG04. Newman, M.E., Girvan, M.: Finding and evaluating community structure in networks. *Physical review E* **69**(2), 026113 (2004)
- PIOH20. Phillips, A.C., Irfan, M.T., Ostertag-Hill, L.: Spheres of legislation: Polarization and most influential nodes in behavioral context. *Complex Networks and Their Applications VIII* **882** (2020)
- Poo08. Poole, K.T.: A Spatial Model for Legislative Roll Call Analysis Author (s): Keith T . Poole and Howard Rosenthal Source : *American Journal of Political Science* , Vol . 29 , No . 2 , (May , 1985) , pp . 357-384 Published by : Midwest Political Science Association *St. Political Science* **29**(2), 357–384 (2008)
- Ros04. Ross, T.J.: *Fuzzy Logic with Engineering Applications*, 2nd edn., pp. 362–444. John Wiley & Sons, Ltd, ??? (2004). doi:10.1016/0042-6989(71)90197-0
- WF94. Wasserman, S., Faust, K.: *Social Network Analysis: Methods and Applications*. Cambridge University Press, ??? (1994)
- Woo08. Woon, J.: Bill sponsorship in Congress: The moderating effect of agenda positions on legislative proposals. *Journal of Politics* **70**(1), 201–216 (2008). doi:10.1017/S0022381607080140
- WPF+09. Waugh, A.S., Pei, L., Fowler, J.H., Mucha, P.J., Porter, M.A.: Party polarization in congress: A network science approach. *arXiv preprint arXiv:0907.3509* (2009)
- ZFT+08. Zhang, Y., Friend, A.J., Traud, A.L., Porter, M.A., Fowler, J.H., Mucha, P.J.: Community structure in congressional cosponsorship networks. *Physica A: Statistical Mechanics and its Applications* **387**(7), 1705–1712 (2008)

Appendix A: Literature Review

We first review the literature on models and algorithms. We then review the literature on polarization.

A.1 Models and Algorithms

Irfan and Ortiz propose *Linear Influence Games (LIGs)* [IO14], a type of 2-action graphical game [KLS01]. In an LIG, every node (or player) represents an individual with a binary action (1 or -1) and a threshold level representing their “stubbornness.” There is an underlying network structure among the nodes. The weight of each edge (u, v) represents the amount of influence that node u exerts on node v . These influence weights may be positive or negative,^[5] and are not required to be symmetric (meaning node u may exert more or less influence on node v than it receives). The best response of a node depends on its threshold level and the net influence on it. The net influence is found by calculating (1) the sum of all incoming influences from nodes playing action 1 and (2) the sum of all incoming influences from nodes playing action -1 , and then subtracting the second sum from the first. If this net influence exceeds the node’s threshold, the best response for that node is 1; if it does not, the best response is -1 . In the case of a tie, the node is indifferent between the two actions.

Instantiating an LIG, then, requires a matrix of influence weights $\mathbf{W} \in \mathbf{R}^{n \times n}$ and a threshold vector $\mathbf{b} \in \mathbf{R}^n$. Each outcome of an LIG is a joint action $\mathbf{x}$, which is basically a vector of actions of all players. For every individual player i , $\mathbf{x}_{-i}$ is the vector of all actions except for i ’s. Each player has an influence function $f_i(\mathbf{x}_{-i}) \equiv \sum_j x_j w_{ij} - b_i$ and a payoff function $u_i(x) \equiv x_i f_i(\mathbf{x}_{-i})$. A joint action $\mathbf{x}^*$ is a *pure-strategy Nash equilibrium* (PSNE) when every individual is playing their best response x_i^* – that is, when no player has an incentive to unilaterally deviate from their chosen action. With the United States Senate as an example, each node is an individual senator, and each edge is the influence a senator has upon another senator. A senator will vote *yea* (1) if their threshold has been met given all incoming influences from other senators, or *nay* (-1) if not. When all senators are playing their best responses in $\mathbf{x}$, the system is stable, and the network is in PSNE. The LIG model is further explained in Section 3.

While the matrix of influence weights and vector threshold values necessary for instantiating an LIG could be generated manually for very small instances, Honorio and Ortiz develop a method of instantiating an arbitrarily large LIG from raw, binary-action data via machine learning [HO15]. Only voting records are made available to the learning program; no other information is involved. Given this data, the program generates the influence weights $\mathbf{w}$ and influence thresholds $\mathbf{b}$ which define a game $\mathbf{G}$. The program seeks to instantiate an LIG where a high proportion of real-world data is accurately reflected as PSNE, without allowing so many PSNE that any joint action would be in equilibria. Finding the number of ground-truth joint actions represented as PSNE is computationally easy, but computing the total number of PSNE in a game is NP-hard, and therefore infeasible on large datasets. By proving a number of simplifying assumptions, they approximate the problem using *convex loss minimization*. In this function, parameters of the game are chosen

^[5]The absence of an edge between a pair of nodes is interpreted as a weight of zero.

so that the average error – the proportion of ground-truth joint actions which are not reflected as PSNE – is minimized. This algorithm is explained in Section 4.

The majority of research in analyzing and predicting legislative votes has not been in the game theory space. Rather, roll-call data is most often used in ideal point models, which estimate the ideal point of a legislature upon a scale of conservative to liberal extremes. Clinton *et al.* proposed Bayesian methods for ideal point estimation, which can be solved using Markov Chain Monte Carlo (MCMC) simulations [CJR04]. In contrast to prior methods, this MCMC-calculated Bayesian method is computationally efficient at large scale; other methods required small populations or made statistical compromises in order to be feasible. Regardless of methods used, ideal points range on an arbitrary scale of negative to positive. In practice, a negative ideal point represents “liberal” polarity, while a positive ideal point represents a “conservative” polarity. Their work is widely cited in later ideal point models which expand upon the original concept.

While the importance of roll-call data is widely recognized, it is also recognized that each vote is a member of a broader context with important characteristics. Grerish and Blei extend the traditional ideal point model, which relies solely on roll-call data, to account for the topics of bills [GB11]. Using a *Latent Dirichlet Allocation* (LDA) topic model, Gerrish and Blei integrate bill topics and political tone into their ideal point model. LDA topic models identify patterns in words, but labelling and interpreting these patterns is left up to the researchers. These bill topics may be, for example, national recognition (“people”, “month”, “recognize”, “history”, “week”, “woman”) or healthcare (“care”, “applicable”, “coverage”, “hospital”, “eligible”). They find that the model performs especially well when bills have bipartisan support or disapproval, or when bills face clearly partisan support and disapproval, but lose accuracy when bills receive mixed, nonpartisan support. Topic modelling is not the only method of inferring bill topics: The Congressional Research Service (CRS) assigns subject codes to every bill, out of close to a thousand possible codes. In an ensuing study on ideal point models, Gerrish and Blei note that using CRS subject codes rather than an LDA topic model also provided a good basis for their ideal point model [GB12].

In a recent paper, Irfan and Gordon add context to the LIG models [IG18]. By combining social interactions and context, they develop a model which performs better than the purely-behavioral model. They learn the ideal points of each senator while learning parameters for the LIG, and account for disparities in polarity across bill topics by utilizing the subject codes of each bill. They expand the influence function of every senator i to include the ideal point of that senator (p_i), and the polarity a_l of a bill l . The product of these two terms is added to the otherwise unchanged influence. When the signs of the polarity of the bill and the ideal point of the senator are the same (e.g., -1.5 and -0.5 , meaning both are liberal leaning), the signs cancel, increasing senator i ’s payoff for voting *yea*; when they differ, a negative value is added, decreasing senator i ’s payoff for voting *yea*.

Some researchers have taken other approaches to modeling congressional behavior. Woon utilizes both ideal points and game-theoretic concepts to analyze how bill sponsorship and co-sponsorship affects the content senators write in a bill [Woo08]. Woon argues that, when sponsoring a bill, legislators balance two opposing forces.

One pushes them towards writing median language, because they want a bill to pass without complications, and the other towards writing highly polarized language, because they wish to signal their beliefs to their constituents. As such, a legislator L will propose a bill with location y within a one-dimensional policy space. They also consider that another legislator, P , will be *pivotal* in allowing a bill's passage. That pivotal senator may choose either y or the status quo, q . P 's choice is known as the policy outcome and is denoted by x . The passage of a bill depends on senator L and P 's utility functions, which consider the distance between x and the ideal points of L and P , respectively. In addition, L 's utility function also considers the weight w that L places on being close to y , which is known as L 's position-taking. Woon extends the model to account for co-sponsorship of other legislators as well, each with their own utility functions. While our research focuses on legislative votes rather than policy proposals, Woon's research affirms the validity of combining contextual data and game-theoretic models, and puts forth bill sponsorship and co-sponsorship as another direction of future research.

Bill sponsorship and co-sponsorship is not the only method by which legislators may signal their preferences for a bill prior to voting. Desmarais et al. build upon prior bill co-sponsorship research to introduce co-participation in press events—called the joint press events network—as an indicator for voting behavior [DMSK15]. Using linear regression, they show a statistically significant positive relationship between press event co-participation and roll-call votes. While not focused on the computational aspects of congressional research, this study highlights the observation that “[l]egislation is often the end product of a lengthy collaborative effort.” Studies like this attempt to uncover ostensibly hidden mechanisms within that lengthy effort. This process starkly contrasts to the behavioral, game-theoretic approach, which makes no assumptions about the underlying mechanism or process, viewing them instead as a “black box”. This lack of assumptions is one of the key benefits of the game-theoretic approach.

Recently, a group of mathematicians took a very different approach to analyzing congressional voting networks from roll-call data. Glonek et al. introduce the Graph Labelling Semi-Supervised (GLaSS) method [GTMB19], a random-walk-based graph labelling method. They model both the House and Senate (from 1935-2017, in different trials) as a graph from roll-call data, where nodes are Democratic or Republican legislators (other parties are ignored), and their labels correspond to their parties. While every senator's party affiliation is known for validation purposes, the only labelled nodes in the graph are the Democratic and Republican party leaders; all other nodes are unlabelled. With the GLaSS method, those nodes are labelled based on the expected time to absorption in a discrete-time Markov chain (DTMC), where absorption states are labelled nodes (i.e. party leaders) and transient states are unlabeled nodes (i.e. other senators). By comparing the labels generated by the GLaSS method to the ground-truth labels of legislator, they measure polarization in Congress. When party affiliation can be accurately predicted by voting trends, Congress is more polarized; when there is some uncertainty, it is less so. Their results show that the U.S. Congress has become remarkably polarized in the past decade, with the model able to accurately predict every senator's affiliation in each term of Congress since 2007. In contrast to Glonek et al.'s stochastic

process-based approach, we model strategic interactions among senators in a game-theoretic fashion that allows us to infer joint behavioral outcomes. Additionally, Glonek *et al.*'s method relies on a model of binary party affiliation and consider nodes as labeled only by party affiliation rather than named as individual senators, which prevents further analysis of the model's network structure.

A.2 Literature Review: Polarization

While modularity [GN02, NG04, New06] has been widely used as a measure of polarization in networks, it is often not a definitive measure. [GMJCK13] presents a novel metric based on the edges incident on the boundary nodes. Like most other metrics of polarization, their metric is also structural in the sense that it does not take into account potentially different network structures among the same population induced by different behavioral contexts. One of the main goals of this paper is to analyze polarization within behavioral context.

Closely related to this paper is Waugh *et al.*'s work on polarization in Congress [WPF+09]. They first compute a weighted network among the members of Congress by counting how many times each pair of members voted the same way. They then compute the modularity of this network as a measure of polarization. Their work can be contrasted with McCarty *et al.*'s ideal point-based approach [MPR16], where the absolute value of the differences in mean ideal points of the two parties serves as a measure of polarization. In fact, our approach may be mistaken as a combination of these two approaches. First, we do compute influence networks among the senators, but these networks are learned from behavioral data. Moreover, there are positive as well as negative edge weights in our networks, whereas Waugh *et al.*'s networks have only non-negative edge weights by definition [WPF+09]. Second, the richer model [IG18] we use combines influence networks with ideal points in such a way that we cannot talk about either networks or ideal points in isolation of the other.

Zhang *et al.* [ZFT+08] study polarization in the U.S. Congress, the same setting as ours. However, theirs is based on cosponsorship networks, which is observed from data. In contrast, ours is based on networks of influence, which have been learned using roll-call and bill-text data. Furthermore, one of the central aspects of our work is to show that polarization in Senate varies according to the spheres of legislation. We do not touch on the rise in polarization in Senate over time, which by now is a well-settled matter [Far15].

Behavioral aspects of polarization among political parties have been studied before, but at an empirical level. Garcia *et al.* analyze multiplex networks consisting of comments, likes, and supports levels among multiple political parties in Switzerland [GAS+15]. In contrast, ours is a model-based approach where polarization can be thought as an inference question.

On a broader level, there have been numerous studies on political polarization. Notably, Conover *et al.* [CRF+11] give evidence of polarization in twitter network based on retweet networks. Interestingly, the opposite happens in mention networks (where ideologically opposing individuals mention each other to start conversations).

There has been some interesting work on the behavioral side of a polarized environment. Bakshy *et al.* [BMA15] use large-scale Facebook data to show that the consumption of politically "hard content" is largely controlled by individuals' own choices and not by algorithmically fed news rankings.

On the computational side, algorithmic approaches to polarization extend beyond modularity. Al Amin et al. [AAAY+17] give a matrix factorization-based algorithm to uncover polarization in twitter networks.

Appendix B: Detailed Cross-Validation Results

In this section, describe the cross-validation results on learning LIG models for Spheres 2, 3, and 4.

Cross-Validation on Sphere 2 (Economics & Finance). The number of edges decreases smoothly, reaching a reasonable number of edges when $0.002424 \leq \rho \leq 0.006236$. Best response error converges for both training and validation sets around $\rho = 0.001512$, and remains at an acceptably low value until $\rho \geq 0.007017$. Until around $\rho = 0.001512$, the high values of training q relative to validation q show that the model is over-fit, and when $\rho > 0.006236$, q 's regression to 0 shows that the model is under-fit. Within this range, both training and validation q are relatively high, at around 0.22. The acceptable range, then, is between 0.002424 and 0.006236 (Figure 17).

Cross-Validation on Sphere 3 (Energy & Infrastructure). The number of edges again decreases smoothly, and is reasonable when $0.002728 < \rho < 0.005541$. Best response error for training and validation sets converges and remains low when $0.001914 < \rho < 0.006236$. As the large difference between training q value and validation q illustrates, the model is over-fit until $q \geq 0.001061$. When $\rho > 0.005541$ and the q value starts to decrease, the model is under-fit. The acceptable range is between 0.002728 and 0.005541 (Figure 18).

Cross-Validation on Sphere 4 (Public Welfare). The number of edges again decreases smoothly, and is reasonable when $0.002728 < \rho < 0.005541$. Best response error for training and validation sets converges and remains low when $0.001914 < \rho < 0.006236$. The high training q relative to validation q shows that the model is over-fit until $\rho \geq 0.001701$, and remains steady until the model begins to become under-fit when $\rho \geq 0.006236$. The acceptable range is between 0.002728 and 0.005541 (Figure 19).

Appendix C: Visualization of LIG Networks

Within the main body of the paper, the LIG networks for Spheres 1 and 2 were shown in Figures 3 and 4, respectively. Here, we show the LIG networks for the other two spheres. These are shown in Figures 20 and 21.

Figure 20 LIG network for Sphere 3 (Energy & Infrastructure).

Figure 21 LG network for Sphere 4 (Public Welfare).

Appendix D: Visualization of Cross-Party Edges

The cross-party edges for Spheres 1 and 2 were shown in Figures 5 and 5, respectively. Here, we show the cross-party edges of the remaining two spheres.

Appendix E: Visualization of Most Influential Nodes

The most influential nodes for Spheres 1 and 3 were shown in Figures 9 and 10, respectively. Here, we show the most influential nodes for the remaining two spheres in Figures 24 and 25.

Figure 24 Sphere 2 (Economics & Finance), $\rho=0.003888$. 10 total levels; 5 Republicans and 4 Democrats.

Figure 25 Sphere 4 (Public Welfare), $\rho=0.003888$. 9 total levels; 5 Republicans and 4 Democrats.

Appendix F: Application of Ideal Point Models with Social Interactions

F.1 Cross-Validation and Model Selection

Here, we show 10-fold cross-validation results of applying the ideal point model with social interactions [IG18] to Sphere 2 (Economics & Finance) only. The results for the other spheres are similar. Using these results, we choose the regularization parameters for the number of edges (ρ) and ideal points and polarities (ρ'). The cross-validation experiments were done over the following range of ρ and ρ' values: $0.001 < \rho < 0.0035$ and $0.0004 < \rho' < 0.0005$. The chosen values of ρ and ρ' are listed in Table 6.

Table 6 Summary of cross-validation results for ideal point model with social interactions.

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
ρ	0.002447	0.002447	0.002184	0.002053
ρ'	0.000395	0.000395	0.000374	0.000395
# Edges	1132	1234	1220	1499

Figures 26, 27, 28, and 29 show the plots for the best response (BR) training error, BR validation error, number of PSNE, and proportion of data captured as PSNE for Sphere 2 (Economics & Finance).

Figure 26 The BR training error is acceptable within the 2-5% range.

Figure 27 The BR validation error is acceptable within the 6-21% range.

Figure 28 The number of edges is acceptable within the 1000-1500 range.

Figure 29 The proportion of data captured as PSNE is acceptable within the 20-45% range.

F.2 Ideal Point Distributions

The plots of ideal point distributions are shown in this section. We first show the scaled versions of the ideal point distributions for Spheres 3 and 4 in Figures 30 and 31, respectively. Note that the ideal point distributions for Spheres 1 and 2 have been shown earlier in Figures 15 and 16, respectively. We also show the non-scaled versions of the ideal point distributions of the four spheres here (Figures 32 to 35).

Figure 30 The ideal points of senators in Sphere 3 (Energy & Infrastructure), scaled linearly. Democratic senators are blue and Republican senators red.

Figure 31 The ideal points of senators in Sphere 4 (Public Welfare), scaled linearly. Democratic senators are blue and Republican senators red.

Figure 32 Non-scaled version of the ideal points of senators in Sphere 1 (Security & Armed Forces). Democratic senators are blue and Republican senators red. The scaled version is shown in Figure 15.

Figure 33 Non-scaled version of the ideal points of senators in Sphere 2 (Economics & Finance). Democratic senators are blue and Republican senators red. The scaled version is shown in Figure 16.

Figure 34 Non-scaled version of the ideal points of senators in Sphere 3 (Energy & Infrastructure). Democratic senators are blue and Republican senators red.

Figure 35 Non-scaled version of the ideal points of senators in Sphere 4 (Public Welfare). Democratic senators are blue and Republican senators red.

F.3 Visualization of Networks

In this section, we show the networks for Spheres 3 and 4 learned under the ideal point model with social interactions (Figures 36 and 37). The networks for the first two spheres have been shown in Figures 11 and 12 within the main body of the paper.

Figure 36 Network for Sphere 3 (Energy & Infrastructure) under the ideal point model with social interactions. Only the strongest one third of the edges are shown.

Figure 37 Network for Sphere 4 (Public Welfare) under the ideal point model with social interactions. Only the strongest one third of the edges are shown.

F.4 Visualization of Cross-Party Edges

In this section, we show *Graphviz* visualizations of edges which connect members of opposite parties within the top 40% of all edges when we apply the richer model to Spheres 3 and 4 (Figures 38 and 39). Figures 13 and 14, shown in the main body of the paper, illustrates the cross-party edges for Spheres 1 and 2, respectively.

Figure 38 Sphere 3 (Energy & Infrastructure): 90 cross-party edges are positive, and 54 are negative.

Figure 39 Sphere 4 (Public Welfare): 101 cross-party edges are positive, and 56 are negative.

F.5 Analysis of Learned Networks Under the Richer Model

Table 7 shows the results of network analysis under the ideal point model with social interactions. This table can be compared with Table 5 shown in the main body of the paper. Table 8 shows the distance between the mean Democratic and Republican ideal points for each of the four spheres.

Table 7 Network analysis of influence networks for different spheres of legislation learned using Irfan and Gordon [IG18]'s model of ideal points with social interactions. Various centrality measures and network-level properties are shown.

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
Number of Edges	1102	1027	1106	1234
Avg. (Shortest) Path Length	2.1641	2.3898	2.2948	2.1806
Avg. Clustering Coefficient	0.165	0.1747	0.1754	0.1812
Modularity	0.5392	0.6801	0.6887	0.6229
Degree Centrality				
Degree (1)	0.451: MANCHIN D-WV	0.4118: MANCHIN D-WV	0.4902: TESTER D-MT	0.5392: CARPER D-DE
Degree (2)	0.4314: CRUZ R-TX	0.3922: KING I-ME	0.451: DONNELLY D-IN	0.4706: TESTER D-MT
Degree (3)	0.4118: CORKER R-TN	0.3824: HEITKAMP D-ND	0.3627: KING I-ME	0.451: MANCHIN D-WV
Degree (4)	0.402: LEE R-UT	0.3627: MENENDEZ D-NJ	0.3627: KLOBUCHAR D-MN	0.451: DURBIN D-IL
Degree (5)	0.3922: PAUL R-KY	0.3529: HELLER R-NV	0.3627: HEITKAMP D-ND	0.451: PAUL R-KY
Degree (6)	0.3922: ENZI R-WY	0.3431: FLAKE R-AZ	0.3627: LEE R-UT	0.4314: MCCASKILL D-MO
Degree (7)	0.3627: SCHATZ D-HI	0.3431: TESTER D-MT	0.3529: COLLINS R-ME	0.4118: HARRIS D-CA
Degree (8)	0.3529: HELLER R-NV	0.3333: DONNELLY D-IN	0.3529: MCCASKILL D-MO	0.3922: STABENOW D-MI
Degree (9)	0.3431: BALDWIN D-WI	0.3333: WYDEN D-OR	0.3529: WARNER D-VA	0.3725: HEITKAMP D-ND
Degree (10)	0.3431: TESTER D-MT	0.3333: FEINSTEIN D-CA	0.3529: CARPER D-DE	0.3627: LEE R-UT
Closeness Centrality				
Closeness (1)	0.5426: MANCHIN D-WV	0.516: MANCHIN D-WV	0.5543: DONNELLY D-IN	0.573: MCCASKILL D-MO
Closeness (2)	0.5231: HELLER R-NV	0.5133: HEITKAMP D-ND	0.5484: TESTER D-MT	0.5635: CARPER D-DE
Closeness (3)	0.5231: COLLINS R-ME	0.5106: FLAKE R-AZ	0.5455: KING I-ME	0.5574: TESTER D-MT
Closeness (4)	0.5204: BALDWIN D-WI	0.5106: GARDNER R-CO	0.5397: HELLER R-NV	0.5574: HARRIS D-CA
Closeness (5)	0.5152: LEE R-UT	0.5028: KLOBUCHAR D-MN	0.5312: HARRIS D-CA	0.5514: DURBIN D-IL
Closeness (6)	0.5152: BOOKER D-NJ	0.5002: KING I-ME	0.5231: KLOBUCHAR D-MN	0.5514: PAUL R-KY
Closeness (7)	0.51: CORKER R-TN	0.4927: MENENDEZ D-NJ	0.5204: DURBIN D-IL	0.5514: HELLER R-NV
Closeness (8)	0.51: SCHATZ D-HI	0.4927: HATCH R-UT	0.5178: HASSAN D-NH	0.5455: STABENOW D-MI
Closeness (9)	0.51: PORTMAN R-OH	0.4902: CARPER D-DE	0.51: BALDWIN D-WI	0.5368: MANCHIN D-WV
Closeness (10)	0.5075: CASEY D-PA	0.4878: HELLER R-NV	0.5075: KIRK R-IL	0.5368: KIRK R-IL
Betweenness Centrality				
Betweenness (1)	0.0552: MANCHIN D-WV	0.0483: MENENDEZ D-NJ	0.0521: LEE R-UT	0.0541: PAUL R-KY
Betweenness (2)	0.0513: CORKER R-TN	0.046: HELLER R-NV	0.0502: TESTER D-MT	0.0467: CARPER D-DE
Betweenness (3)	0.0365: CRUZ R-TX	0.0435: FLAKE R-AZ	0.0467: DONNELLY D-IN	0.0415: DURBIN D-IL
Betweenness (4)	0.0346: HELLER R-NV	0.0428: KING I-ME	0.039: KLOBUCHAR D-MN	0.0335: MORAN R-KS
Betweenness (5)	0.0339: ENZI R-WY	0.0427: LEE R-UT	0.0331: FLAKE R-AZ	0.0331: MCCASKILL D-MO
Betweenness (6)	0.0337: HEITKAMP D-ND	0.0415: HEITKAMP D-ND	0.0311: KIRK R-IL	0.0321: TESTER D-MT
Betweenness (7)	0.0334: KING I-ME	0.0401: TESTER D-MT	0.0293: HELLER R-NV	0.0305: MANCHIN D-WV
Betweenness (8)	0.0302: LEE R-UT	0.0369: MANCHIN D-WV	0.0292: PAUL R-KY	0.0281: LEE R-UT
Betweenness (9)	0.0264: BALDWIN D-WI	0.0314: CARPER D-DE	0.0286: SANDERS I-VT	0.028: MCCONNELL R-KY
Betweenness (10)	0.0262: PAUL R-KY	0.0284: DONNELLY D-IN	0.0286: HEITKAMP D-ND	0.0273: DONNELLY D-IN
Eigenvector Centrality				
Eigenvector (1)	0.2016: LEE R-UT	0.2159: MANCHIN D-WV	0.2174: TESTER D-MT	0.22: MCCASKILL D-MO
Eigenvector (2)	0.1908: MANCHIN D-WV	0.2128: HEITKAMP D-ND	0.1995: KING I-ME	0.2109: CARPER D-DE
Eigenvector (3)	0.1874: BALDWIN D-WI	0.1926: GARDNER R-CO	0.1801: DONNELLY D-IN	0.1908: HARRIS D-CA
Eigenvector (4)	0.1873: SCHATZ D-HI	0.1884: KING I-ME	0.1719: HEITKAMP D-ND	0.1852: PAUL R-KY
Eigenvector (5)	0.1819: COLLINS R-ME	0.1864: FLAKE R-AZ	0.169: CARPER D-DE	0.1766: TESTER D-MT
Eigenvector (6)	0.1675: SANDERS I-VT	0.1697: HATCH R-UT	0.1677: KIRK R-IL	0.1696: KIRK R-IL
Eigenvector (7)	0.1671: BOOKER D-NJ	0.1642: ALEXANDER R-TN	0.1643: HELLER R-NV	0.1674: STABENOW D-MI
Eigenvector (8)	0.1623: CORKER R-TN	0.1642: CARPER D-DE	0.1573: GRAHAM R-SC	0.166: MANCHIN D-WV
Eigenvector (9)	0.1577: HELLER R-NV	0.1633: LANKFORD R-OK	0.1549: KLOBUCHAR D-MN	0.157: DURBIN D-IL
Eigenvector (10)	0.1442: PAUL R-KY	0.1603: MENENDEZ D-NJ	0.1536: COLLINS R-ME	0.1566: WARREN D-MA

Table 8 Distance between the average Ideal Point (normalized) of Republican and Democratic Senators.

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
Distance	0.754	1.235	1.126	0.889