

RESEARCH

Open Access



Spheres of legislation: polarization and most influential nodes in behavioral context

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

*Correspondence:

mirfan@bowdoin.edu

[†]Phillips and Ostertag-Hill worked on this research as undergraduate students at Bowdoin

¹ Department 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

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 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

Introduction

In recent times, the study of social influence has extended beyond mathematical sociology [16, 24, 50] and has entered the realm of computation [1, 4–7, 28, 30, 33, 34, 36, 37]. 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 "[Spheres of legislation](#)" section.

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, the topic of modeling and analyzing congressional voting behavior has been getting a lot of attention in both political science and computer science [9, 19, 28, 30, 45], 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 [29, 30]. 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 represented by a pure-strategy Nash equilibria (PSNE) if their support for that outcome influences enough other individuals to achieve that outcome. The LIG model will be elaborated in "[The LIG model](#)" section and machine learning of this model using the spheres of legislation will be detailed in "[Machine learning](#)" section.

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 [28] 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 [28] focused on making predictions given a single bill, rather than analyzing the network as a whole. In "[Towards richer models: ideal point models with social interactions](#)" section, 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 [28–30], 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., [30]), 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 [4, 16, 21, 22, 39, 41], especially in the political arena [10, 17, 27, 40, 52, 53] and often in a regional context [2, 43]. A detailed literature review is provided in Appendix A. 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 [15].

Two recent congressional terms—114th and 115th—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.

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.

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 contain senators' "yea," "nay," or abstaining votes, while bill data include 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

Table 1 Changes in congressional membership within the 114th–115th congresses

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

When two senators are combined, the senator name in italicface indicates the label which was kept. This table should be referred to when examining figures and graphs which depict individual senators

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. 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

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.

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 @unitedstates. 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 [38]. 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 that 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 [3, 47]. 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

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%	

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

The LIG model

We represent the senate influence network as a linear influence game (LIG) [29, 30], one type of 2-action graphical game [35]. 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 that 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. Fig. 1 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 Fig. 1, 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

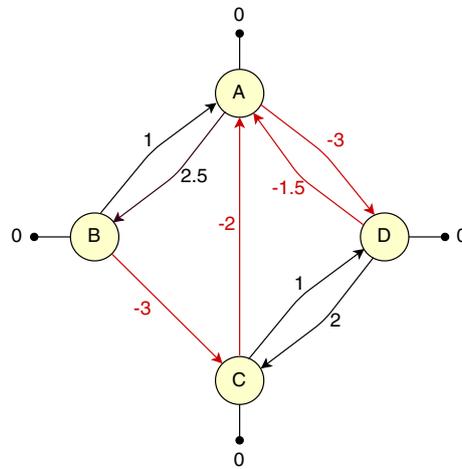


Fig. 1 LIG example. A four-node LIG is shown here. The directed edges are labeled with influence levels. Any absence of a directed edge implies an influence level of 0. Threshold values of 0 (for simplicity) are shown with a connector to each node. We assume binary actions $\{1, -1\}$. In this game, 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, and 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

influence function defined in Definition 3.1. 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, PSNE is formally defined in Definition 3.3. In the following formal definitions, we use the same notation as [30].

Definition 3.1 (Influence function [30]) 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 Fig. 1, 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 [30]) 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 Fig. 1, 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 [30]) 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 (Fig. 1), 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 Fig. 1, 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.

Machine learning

We use Honorio and Ortiz's machine learning algorithm to instantiate an LIG from raw roll-call data [26]. 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 [26].

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 [11, 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 [26] 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 [30]. 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 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.

Cross-validation and model selection for LIG

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 overfitting—“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 [29, 30]. 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) to determine the effectiveness of a given ρ value. In CV, a process essential to most machine learning applications, data are 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 overfit, 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 to create the best model which neither underfits nor overfits 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 to form useful insights on small datasets. In k -fold CV, the dataset is randomized and split into k partitions. In one run

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 combine two terms (i.e., four sessions), this was not practical.

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 overfits 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 Fig. 2, the number of edges drastically decreases until $\rho = 0.000367$ and 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 $\rho \geq 0.004$ and 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 overfit, 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 the overfitting 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.

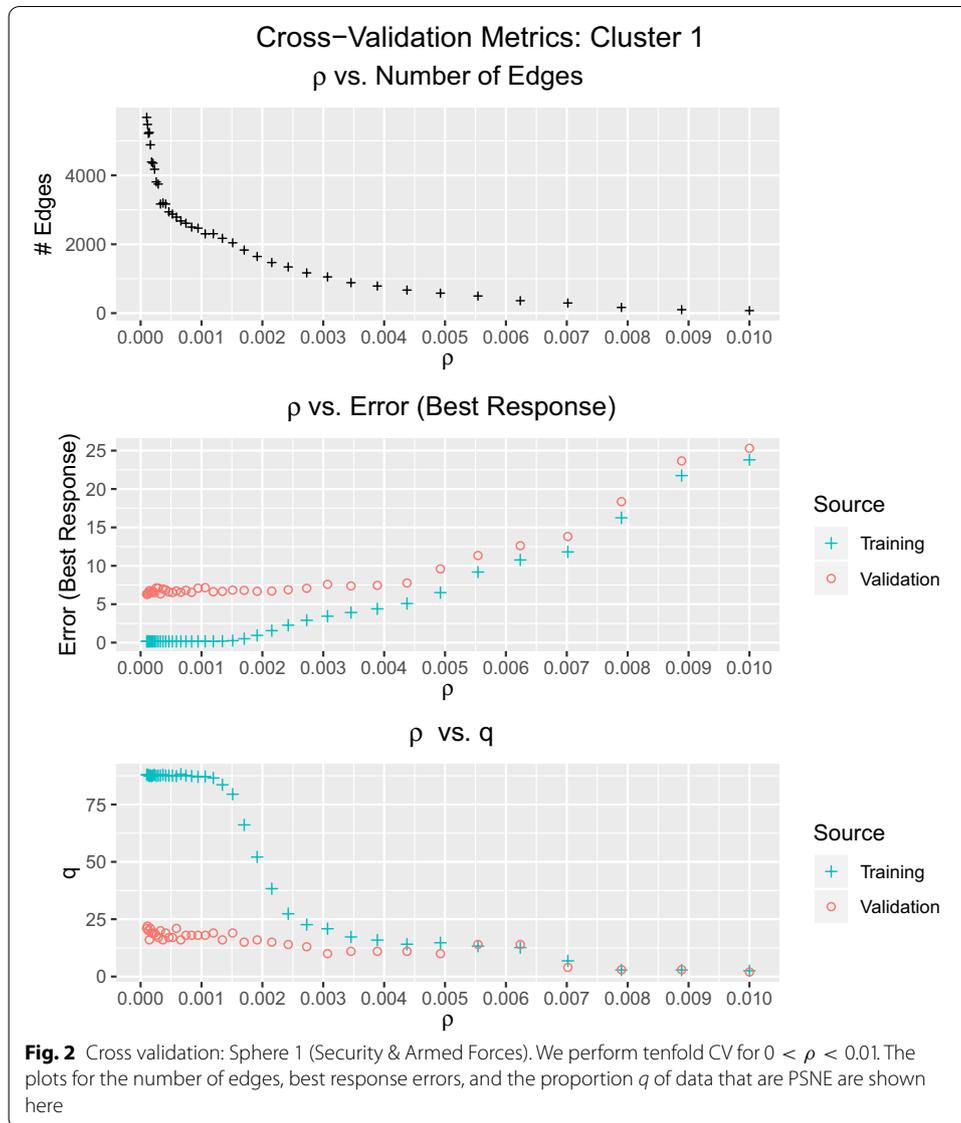
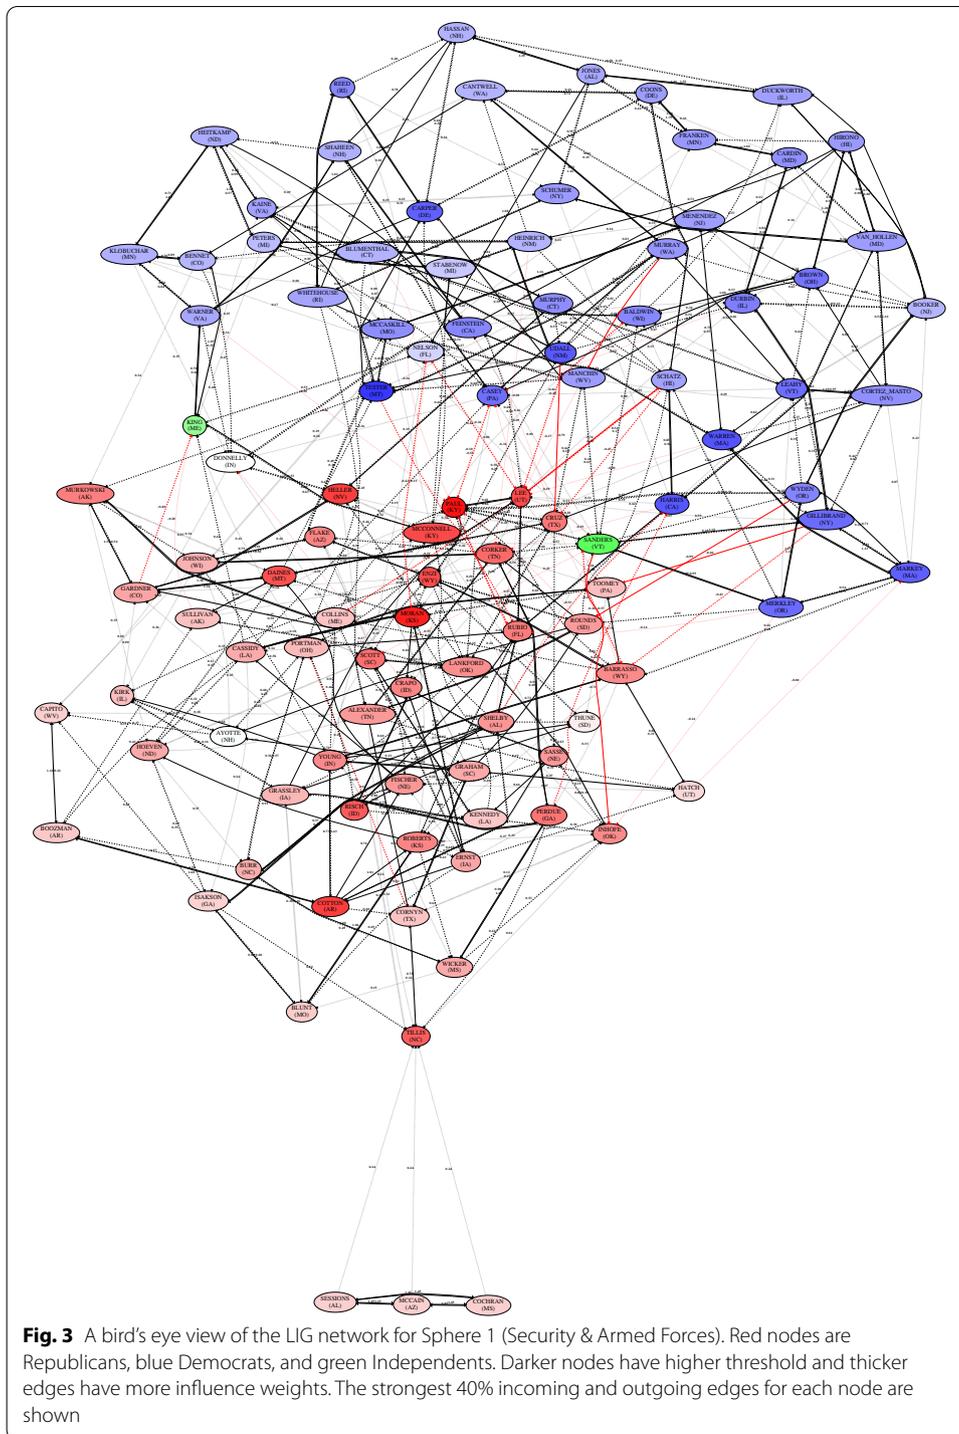


Table 3 Chosen values of ρ for each sphere and the corresponding number of edges and the average best response (BR) error of validation sets

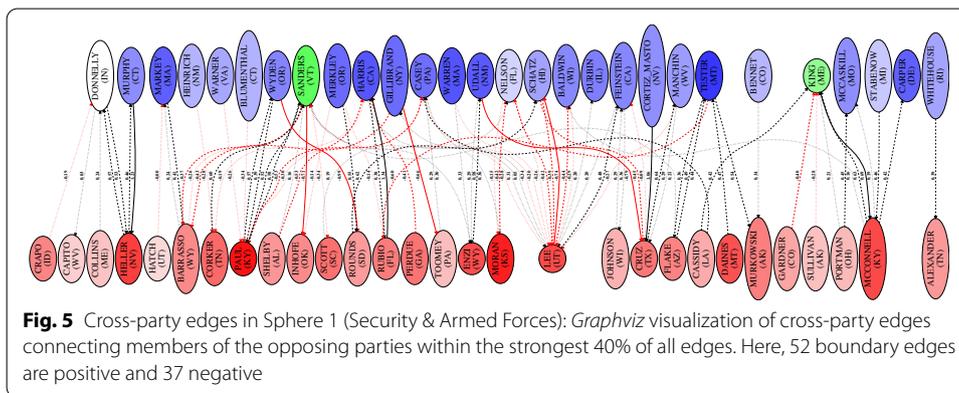
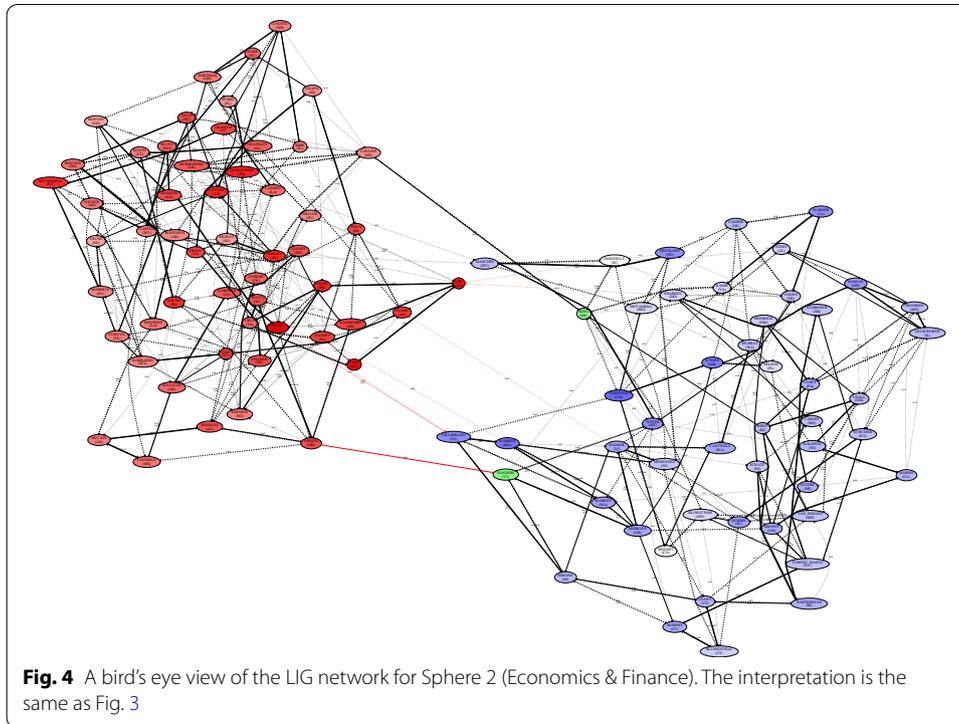
Sphere	ρ	# Edges	BR error (validation)
1	0.002728	1191	7.07%
2	0.003888	1071	5.05%
3	0.003070	1280	5.56%
4	0.003888	1076	5.10%

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

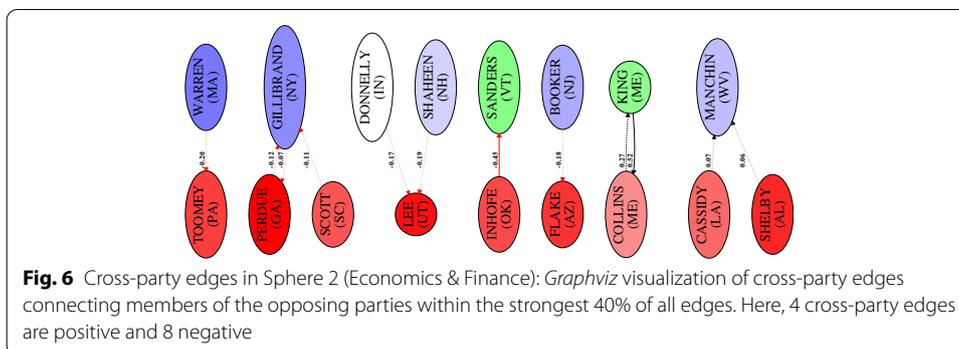


cross-border) edges, influence weights and thresholds, and modularity measures. We begin with cross-party edges.

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 Figs. 5, 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 Fig. 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.

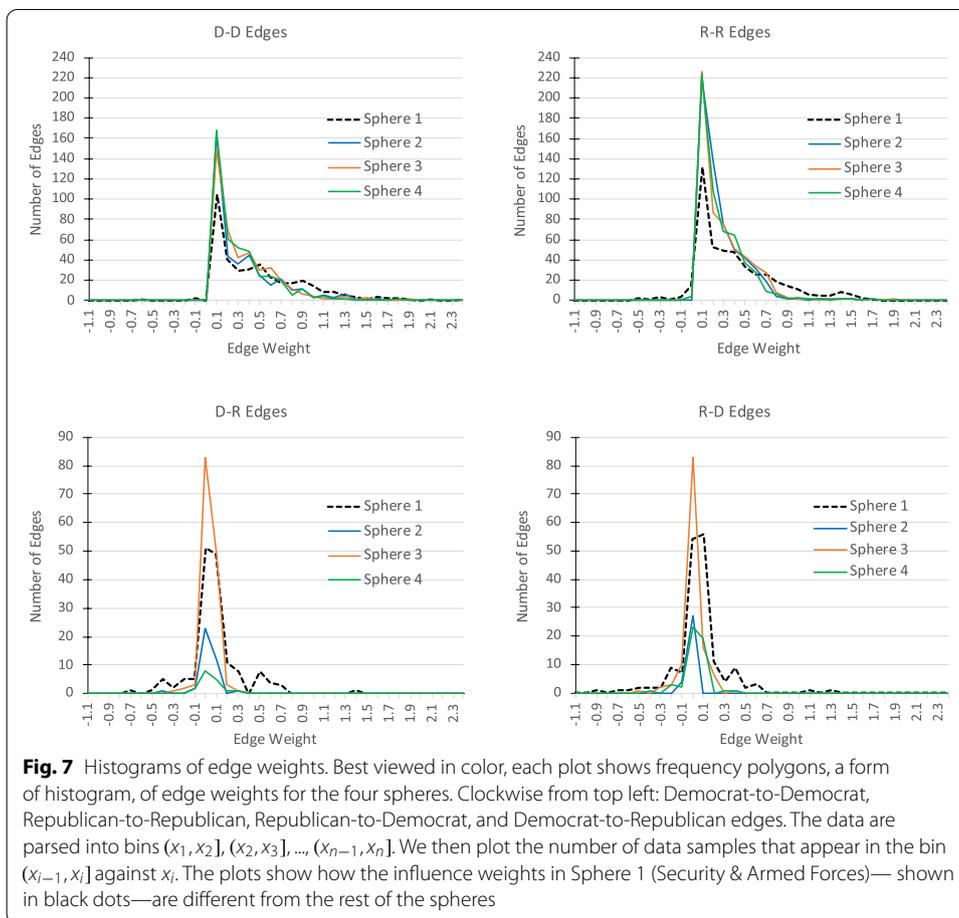
Similarly, examining inter-party edges reveals that Sphere 3 (Energy & Infrastructure) is also very polarized. While there are many edges between both parties in this network, about 70% of them are negative. Positive influences come from a few sources, again including the centrist Senator Collins. Incongruously, 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.

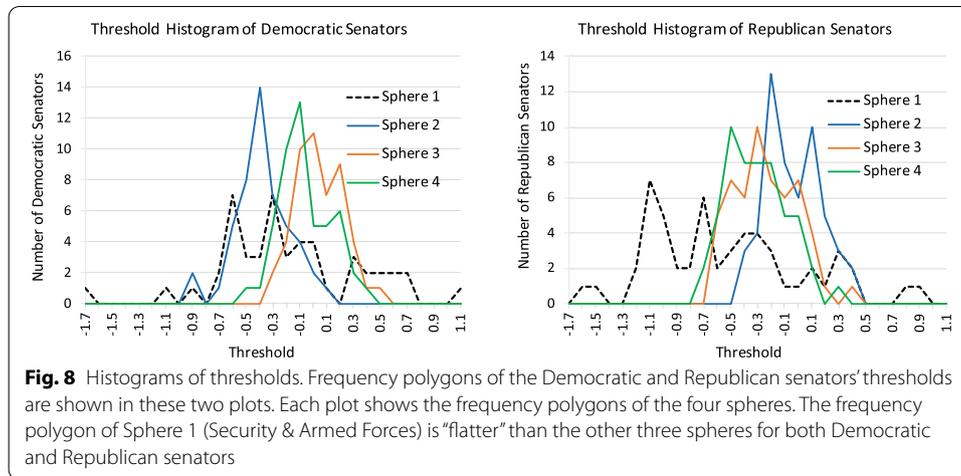
Influence weights and thresholds

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 Fig. 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.

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

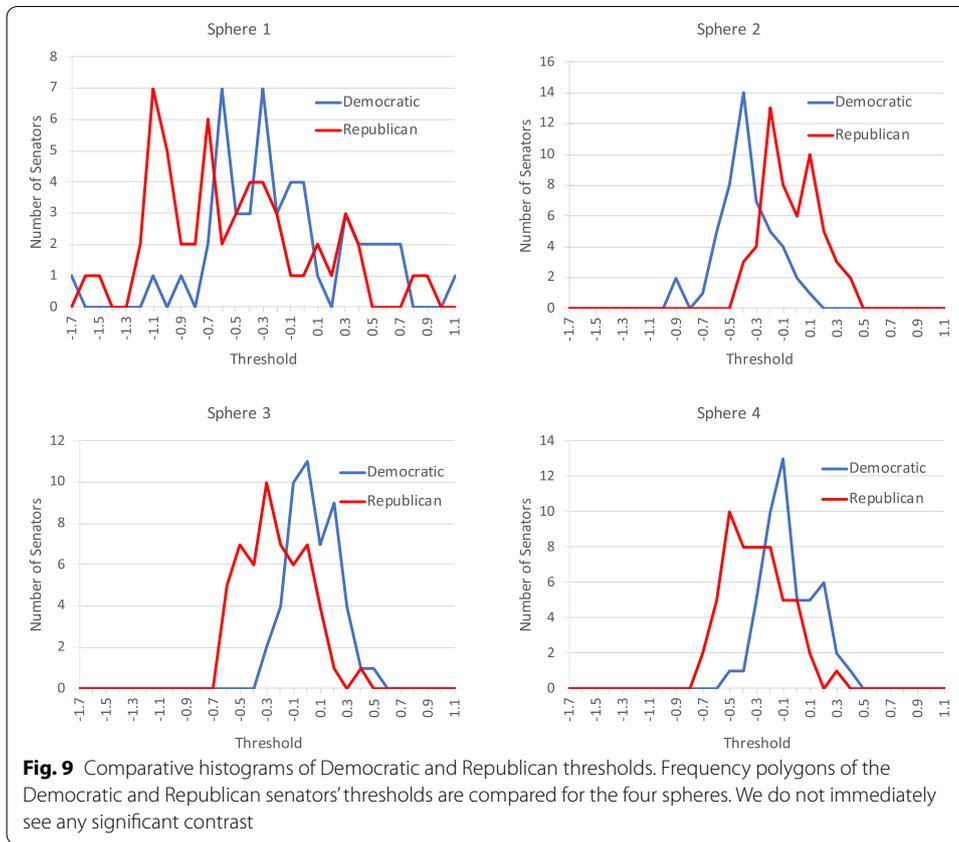


Of course, the influence weights cannot be read alone without considering thresholds because the game-theoretic model accounts for both of these in predicting stable outcomes. Recall that the threshold *magnitude* signifies stubbornness or resistance to influence. More positive threshold values resist 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 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 in 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. While Fig. 8 shows the histogram of each party for different spheres, Fig. 9 makes a comparison of the histograms of the two parties for each sphere separately. The contrast between the two parties is not as remarkable as the contrast among the four spheres for any party.

As a final note, we emphasize that the threshold values on their own lack sufficient predictive power. In fact, the main component of the LIG model is the interdependence among the senators' actions through the influence structure. Having said that, if a sphere is overwhelmingly dominated by bills sponsored by one of the two parties, then it is possible that the machine learning algorithm would assign low threshold values to the senators of that party (that is, those senators would be predisposed to voting *yea*).⁵ Even then, the influence weights would play a role in predicting the stable outcomes. Investigating this issue using sponsorship and co-sponsorship data is an interesting future direction.

⁵ We thank an anonymous reviewer for raising this point.



Modularity

Furthermore, a formal study of polarization rooted in network science produces similar results. Modularity [23, 41, 42] 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 [20]:

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

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.

It is important to note that modularity does not always indicate polarization. As Guerra et al. [22] show, there are networks that exhibit community structure despite

not being polarized. However, in our case, we are not investigating whether Congress is polarized or not. Polarization in Congress is already a settled matter [15]. We are rather investigating to what degree Congress is polarized based on the spheres of legislation. Furthermore, our analysis of cross-party edges ("[Cross-party edges](#)" section) resonates with Guerra et al.'s main idea that in a polarized network, the nodes at the border are on average more connected inside their own community than outside.

Most influential nodes in context

There exists a number of centrality measures that are derived from a structural analysis of networks [31]. 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 [29, 30]. 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 "[Computing most influential nodes](#)" section, it requires computation of all PSNE. We next outline how we compute all PSNE for each sphere of legislation.

PSNE computation

Once the LIG is instantiated by the machine learning algorithm (see "[Machine learning](#)" section), we can compute the set of all PSNE using the algorithm described in [30]. 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

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

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.

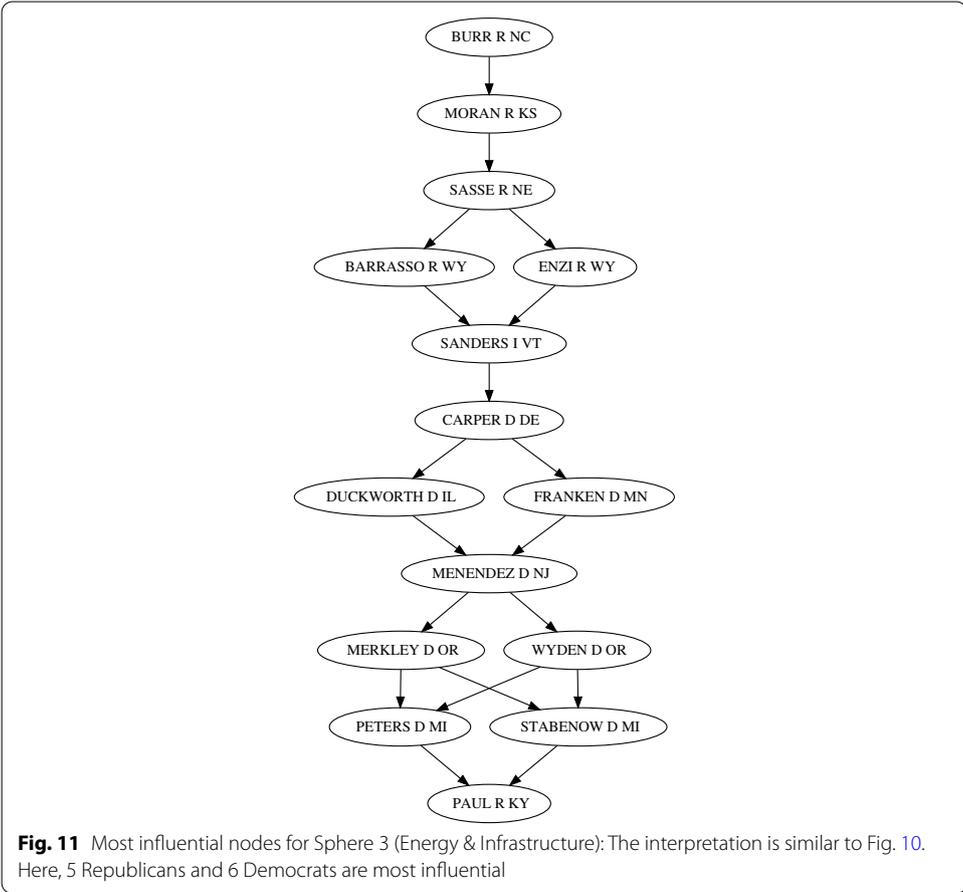
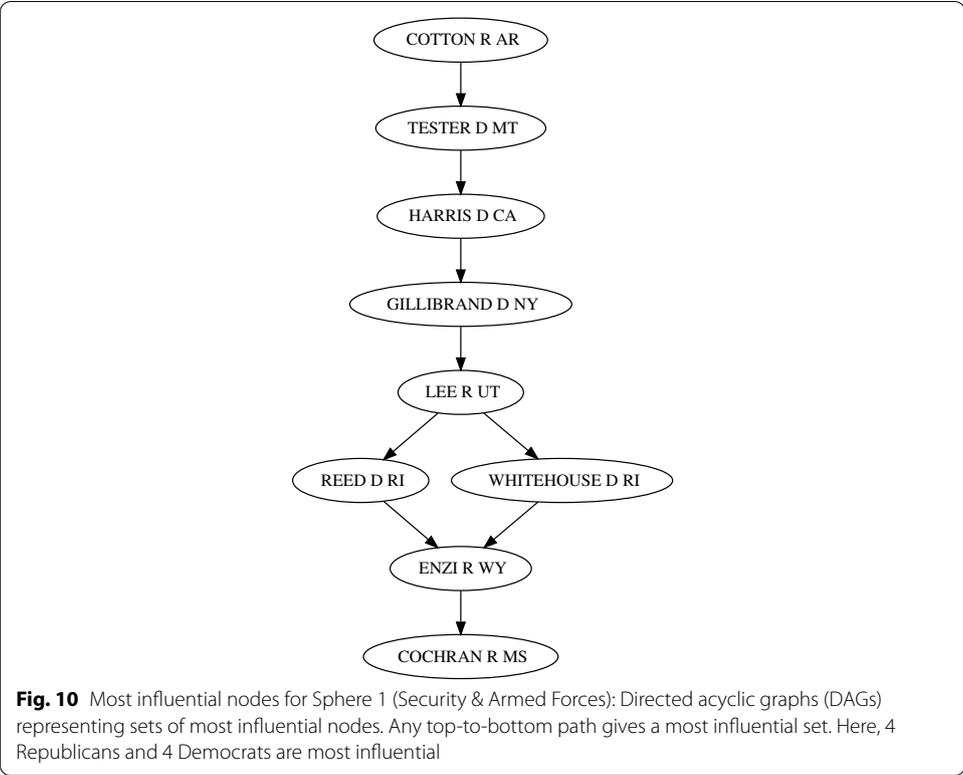
Computing most influential nodes

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 [30], which produces a directed acyclic graph (DAG). The algorithm requires precomputation of all PSNE, which is a provably hard problem [30]. 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 10, 11 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 Figs. 10, 11 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 Fig. 10, 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 Fig. 10 are sufficient to generate the maximum possible support for a bill in Sphere 1. As shown in Fig. 11, 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.

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 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 suggest 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 "[Polarization in context](#)" section.

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 overrepresented, comprising the majority of the top ten most central nodes. *In contrast, the game-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.

Toward richer models: ideal point models with social interactions

We also apply a richer model of influence recently proposed by Irfan and Gordon [28] that extends the LIG model by incorporating ideal points of senators and polarities of bills. 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.

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-GA
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

Table 5 (continued)

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
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

We start with an overview of how Irfan and Gordon’s model [28] builds on the political science literature on ideal point models [8, 13, 32, 45, 46, 48]. Ideal point models are predictive statistical models that assign each senator i an *ideal point* p_i signifying the senator’s legislative position. Usually, more negative values of p_i mean more liberal position and more positive values mean more conservative. Similarly, each bill l is also assigned a *polarity* a_l signifying the position of the bill in the liberal to conservative spectrum. There is a third model parameter called the *popularity* r_l of bill l representing the fraction of senators supporting the bill. The ideal point model in its most basic form defines the probability of senator i supporting bill l using the following logistic function σ :

$$p(x_{i,l} = \text{yea} \mid p_i, a_l, r_l) = \sigma(p_i a_l + r_l).$$

The ideal point model captures the interdependence among the senators using the r_l term. However, this term is an aggregate measure quantified by the number of senators voting *yea* on bill l . In ideal point models with social interactions, Irfan and Gordon expand this aggregate measure by considering how the individual senators are voting and how their votes influence each other [28]. The resulting model is game-theoretic with

the following influence function. Here, other than the new terms l , p_i , and a_l defined in the previous paragraph, the rest of the terms are the same as those in Definition 3.1:

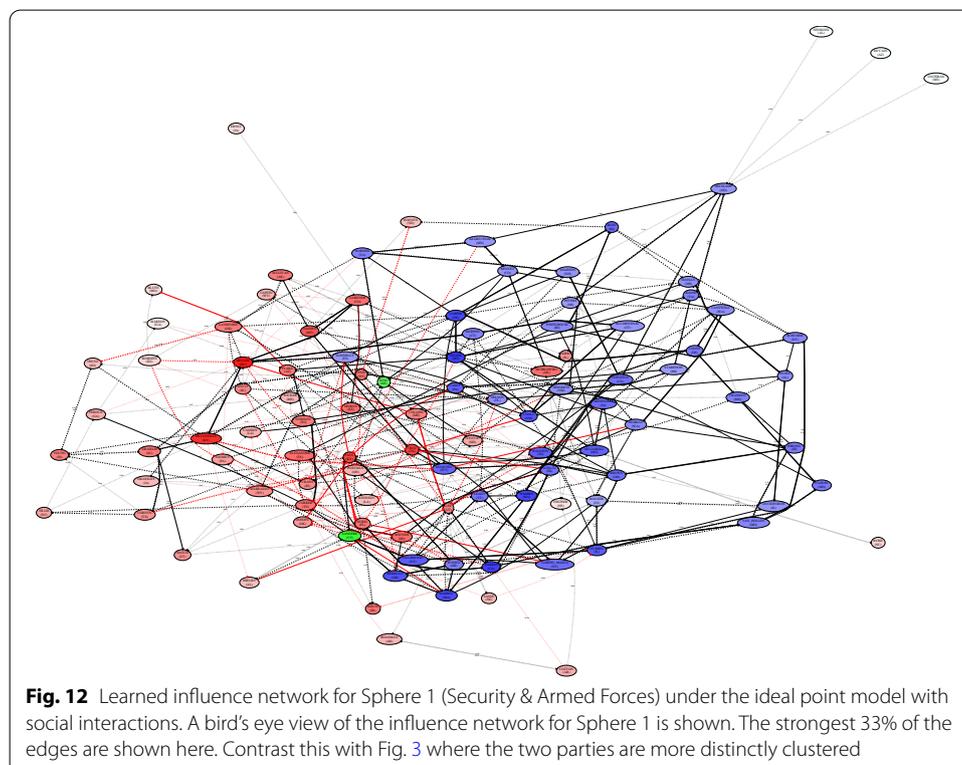
$$f_i(\mathbf{x}_{-i}, l) \equiv \sum_{j \neq i} w_{ij} x_j + (p_i \cdot a_l) - b_i.$$

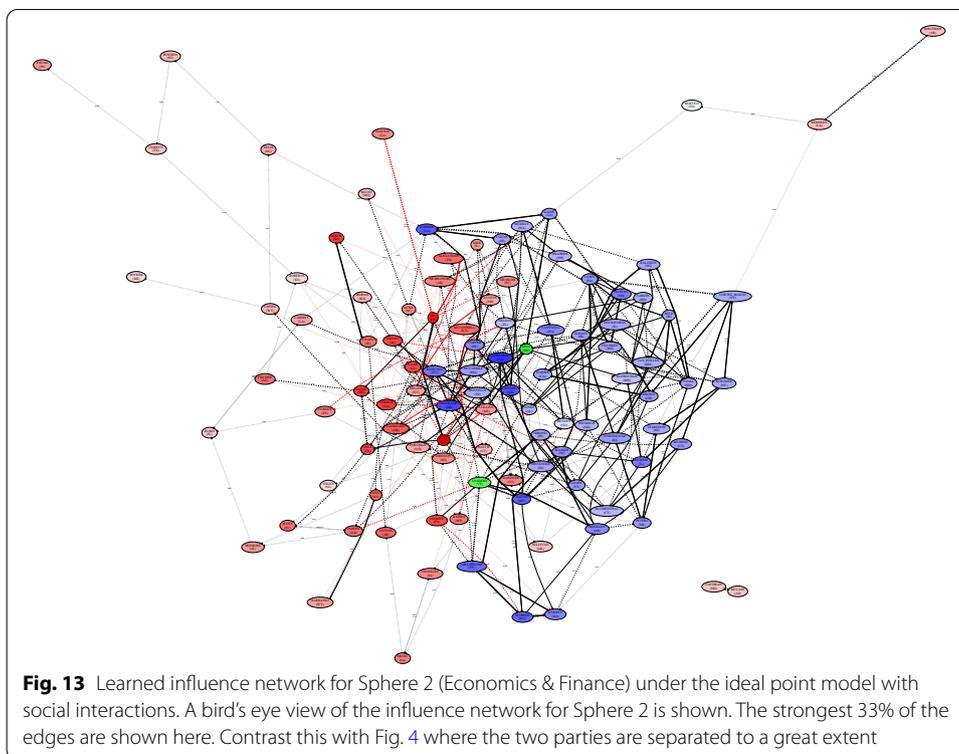
Using the above influence function, the richer game-theoretic model is defined in the same fashion as "The LIG model" section.

As a cautionary note, the way Irfan and Gordon’s model [28] 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. Figs. 12, 13 show the learned networks for Sphere 1 (Security & Armed Forces) and 2 (Economics & Finance) under this richer model (other spheres are in Appendix E.3). First, it is evident that the two parties are not as clustered as they were in the LIG model (compare with Figs. 3, 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 Figs. 14, 15, 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 "Polarization in context" section yields scores of 0.5392, 0.6801, 0.6887, and 0.6229,

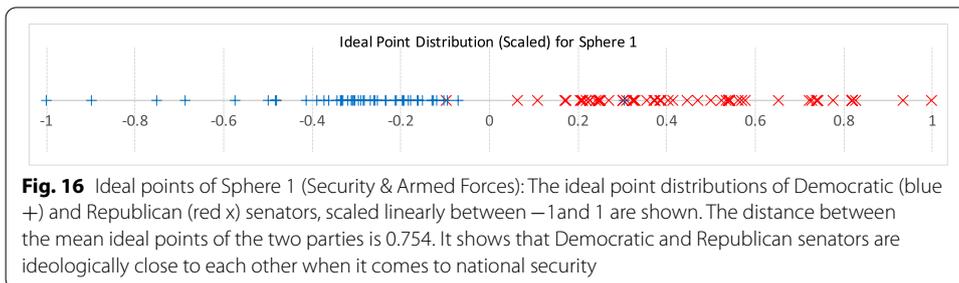
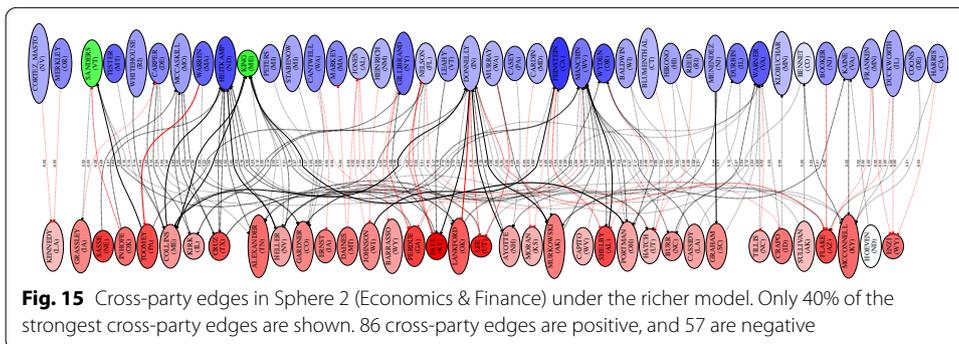
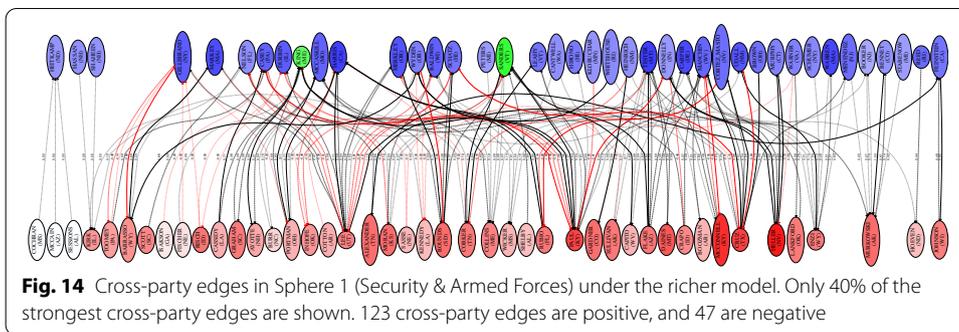




respectively. Both the ideal point metric and modularity scores indicate that Spheres 2 (Economics & Finance) and 3 (Energy & Infrastructure) are most polarizing, whereas Sphere 1 (Security & Armed Forces) is least polarizing. Sphere 4 (Public Welfare) sits in between. These results are somewhat similar 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) [40] to calculate polarization levels across the four spheres. The ideal point distributions for two of the spheres are depicted in Figs. 16, 17 (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 1–4, 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 Figs. 16, 17. 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.

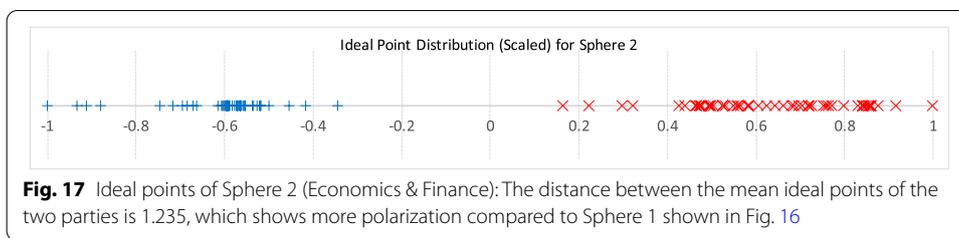
We have also applied a recently proposed measure called *polarization index* [39]. Inspired by the electric dipole moment, the polarization index is measured from an



opinion distribution, where opinions propagate from a set of *elite* entities (e.g., influential politicians and media accounts on Twitter) to *listener* entities (e.g., ordinary individuals on Twitter). The measure is based on opinion distribution (as opposed to dynamics). Here, we apply it to the machine learned ideal point distribution.⁶ We use the following definition of the polarization index μ , where ΔA represents the difference between the fraction of Republicans and Democrats and gc^+ and gc^- represent the gravity centers of the Republican and Democratic senators' ideal points, respectively:

$$\mu = (1 - \Delta A)(gc^+ - gc^-)/2.$$

⁶ We note a semantic difference between the originally proposed polarization index [39] and how we are applying it here. Originally, the polarization index was measured from a distribution of behaviors generated through a network of retweets. In our case, although there are two major parties, the strategic interdependence among the nodes leads to multiple *joint* behaviors represented as PSNE. Both the multiplicity of PSNE and the interdependence of the behaviors in a PSNE make a direct application of the polarization index to the PSNE setting unclear. We are rather applying it to the ideal point distribution.



This definition produces the following polarization indices for the four spheres, respectively: 0.3588, 0.5874, 0.5356, and 0.4227. These are a constant factor off of the ideal point-based polarization measures [40] presented before due to the similarity between the two definitions in our case.

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.

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 [28] for different spheres of legislation. We have briefly touched upon it in "[Towards richer models: ideal point models with social interactions](#)" section. 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 [28] is another interesting direction. In particular, what happens to the balanced, bipartisan composition of most influential sets under the LIG model (see "[Most influential nodes in context](#)" section when we incorporate additional contextual parameters like ideal points and polarities?)
4. On the computational front, Irfan and Gordon [28] 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 [25] in game-theoretic settings is also an interesting direction.

In addition to the above open directions in the context of legislative chambers, the LIG model may also be applied to other settings where network-connected individuals exhibit influence or behavioral interdependence. Some examples in the public health domain are smoking [7] and obesity [6]. Other promising areas include smart electricity grids, vaccination, and the adoption of microfinance.

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

While the study of influence in networks is very broad [16], we focus on models and algorithms for game-theoretic settings. Irfan and Ortiz propose *Linear Influence Games (LIGs)* [30], a type of 2-action graphical game [35]. 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,⁷ and are not required to be symmetric (meaning that 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 the action of i . Each player has an influence function $f_i(\mathbf{x}_{-i}) \equiv \sum_j x_j w_{ij} - b_i$ and a payoff function $u_i(x) \equiv x f_i(\mathbf{x}_{-i})$. 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 that 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 "[The LIG model](#)".

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 [26]. Only voting records are made available to the learning program; no other information is involved. Given these 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 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 "[Machine Learning](#)" section.

The majority of research in analyzing and predicting legislative votes has not been in the game theory space. Rather, roll-call data are 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 [9]. In contrast to prior methods, this MCMC-calculated Bayesian method is computationally efficient at

⁷ The absence of an edge between a pair of nodes is interpreted as a weight of zero.

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. Gerrish and Blei extend the traditional ideal point model, which relies solely on roll-call data, to account for the topics of bills [18]. 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 labeling and interpreting these patterns are left up to the researchers. These bill topics may be, for example, national recognition (“people”, “month”, “recognize”, “history”, “week”, and “woman”) or healthcare (“care”, “applicable”, “coverage”, “hospital”, and “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 [19].

In a recent paper, Irfan and Gordon add context to the LIG models [28]. 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 that 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 affect the content senators write in a bill [51]. Woon argues that, when sponsoring a bill, legislators balance two opposing forces. One pushes them toward writing median language because they want a bill to pass without complications, and the other toward 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, 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 [14]. 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 Labeling Semi-Supervised (GLaSS) method [25], a random-walk-based graph labeling 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 unlabelled 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 considers 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 [23, 41, 42] has been widely used as a measure of polarization in networks, it is often not a definitive measure. Guerra et al. presents a novel metric based on

the edges incident on the boundary nodes [22]. 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 [52]. 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 [40], 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 [52]. Second, the richer model [28] which 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. [53] study polarization in the U.S. Congress, the same setting as ours. However, theirs is based on co-sponsorship 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 [15].

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 [17]. In contrast, ours is a model-based approach where polarization can be considered an inference question.

At a broader level, there have been numerous studies on political polarization. The edited volume by Hopkins and Sides [27] presents a comprehensive treatise from three different perspectives: why American politics is polarized, how it became polarized, and what we can do about it (including whether the alternatives are any better). As a specific example, Conover et al. [10] 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).

Not surprisingly, Twitter provides a trove of data that has been used in several other studies. Notably, Morales et al. [39] give a framework to estimate *polarization index* using a model of opinion generation. Unlike other generative models of opinion propagation [49], their focus is on the distribution of opinions and not the dynamics of opinions. We briefly reviewed their model in “[Toward richer models](#)” section. One major difference between Morales et al.'s work and ours is how we get to the behavioral distribution (or PSNE in our case). In our models, we do not have predefined elite and listener

nodes and do not perform DeGroot-style iterative updates [12]. Furthermore, the complexity of interdependent actions in a PSNE and the multiplicity of PSNE make a direct application of polarization index to our setting challenging (see Footnote 6).

Whereas Morales et al. apply polarization index to a case study of tweets in the aftermath of Venezuelan leader Hugo Chávez's death, their basic idea has been generalized to any Twitter topics by Garimella et al. [21]. Of course, there are methodological differences between the two studies. Garimella et al.'s random walk-based algorithm to measure polarization is promising for large-scale networks. In contrast to these studies, we use machine learning to learn the networks of influence from voting data. Also, our behavioral model is strictly game-theoretic.

There has also been some interesting work on the behavioral choice of individuals in a polarized environment. Bakshy et al. [4] 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.

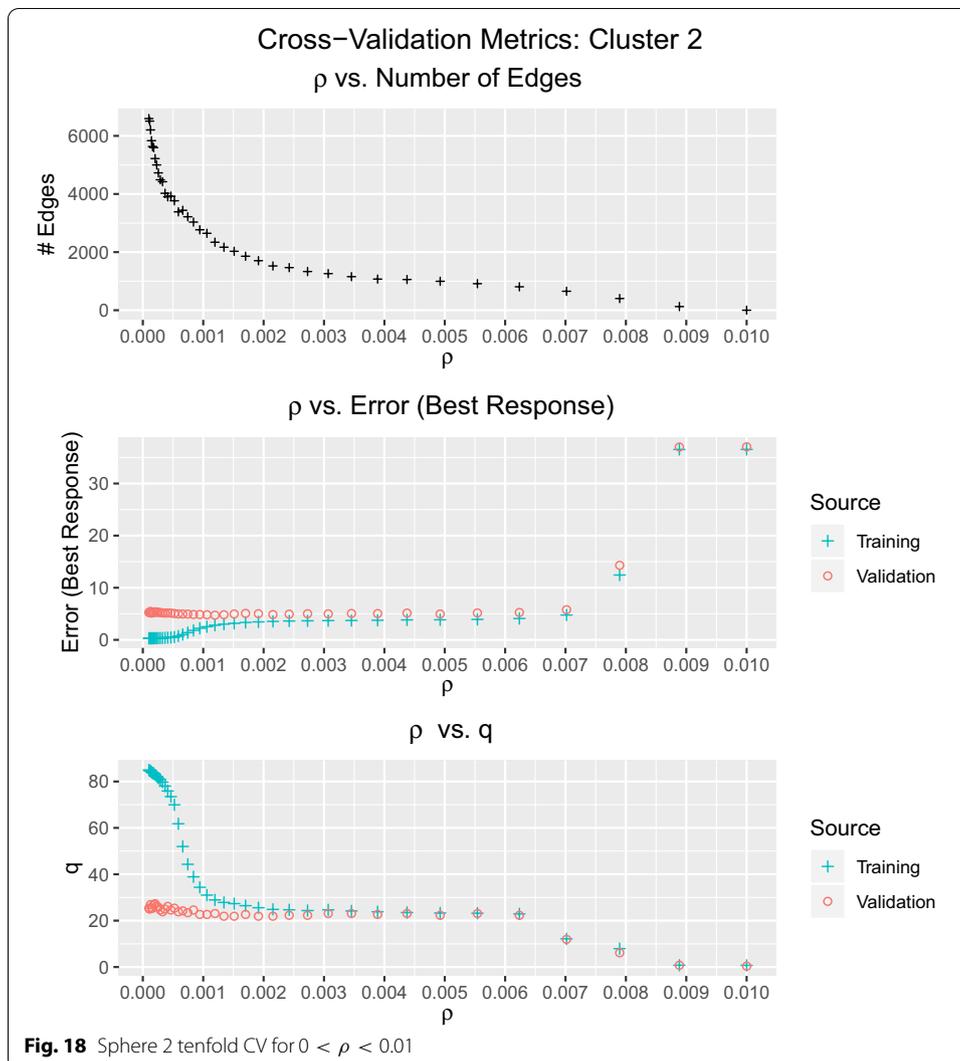
On the computational side, algorithmic approaches to polarization extend beyond modularity. Al Amin et al. [1] give a matrix factorization-based algorithm to uncover polarization in Twitter networks.

Appendix B: Detailed cross-validation results

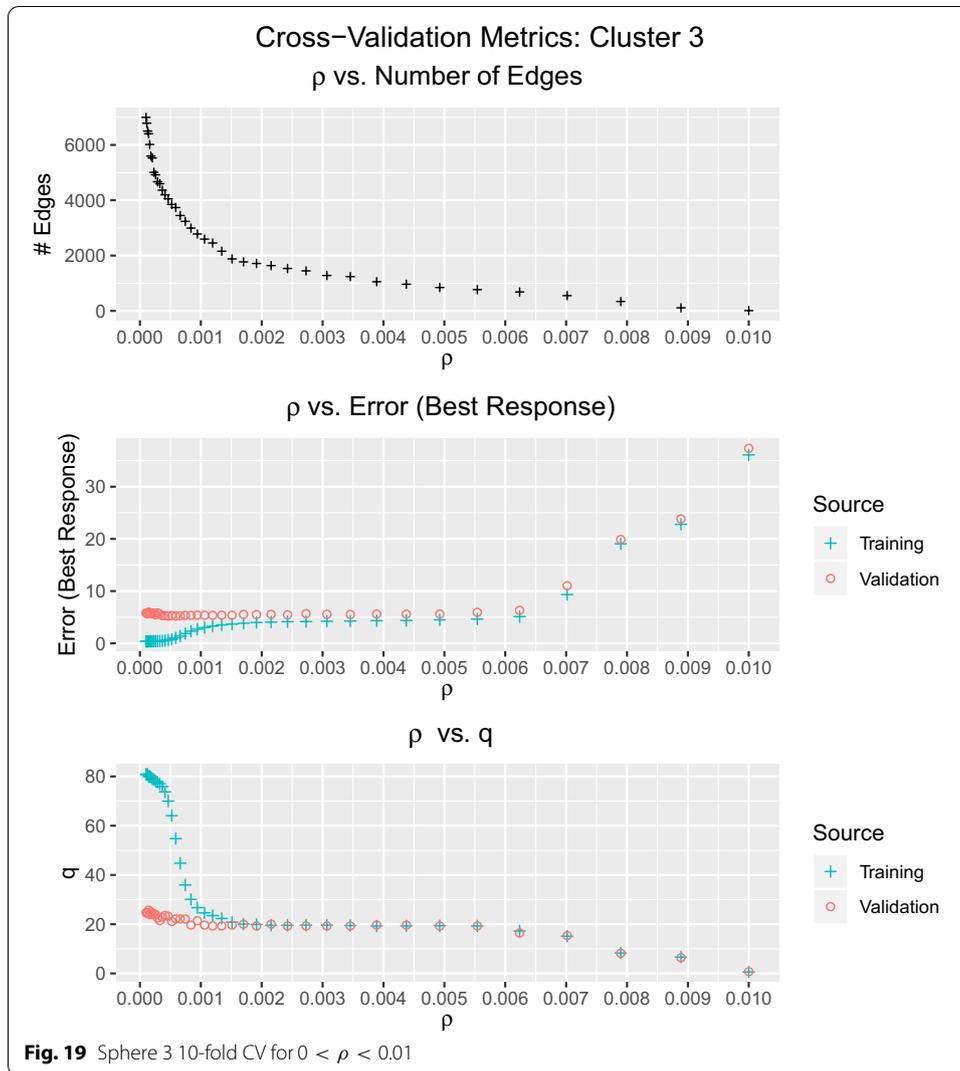
In this section, describe the cross-validation results on learning LIG models for Spheres 2, 3, and 4.

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 overfit, 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 (Fig. 18).

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 overfit 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 (Fig. 19).



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 overfit 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 (Fig. 20).

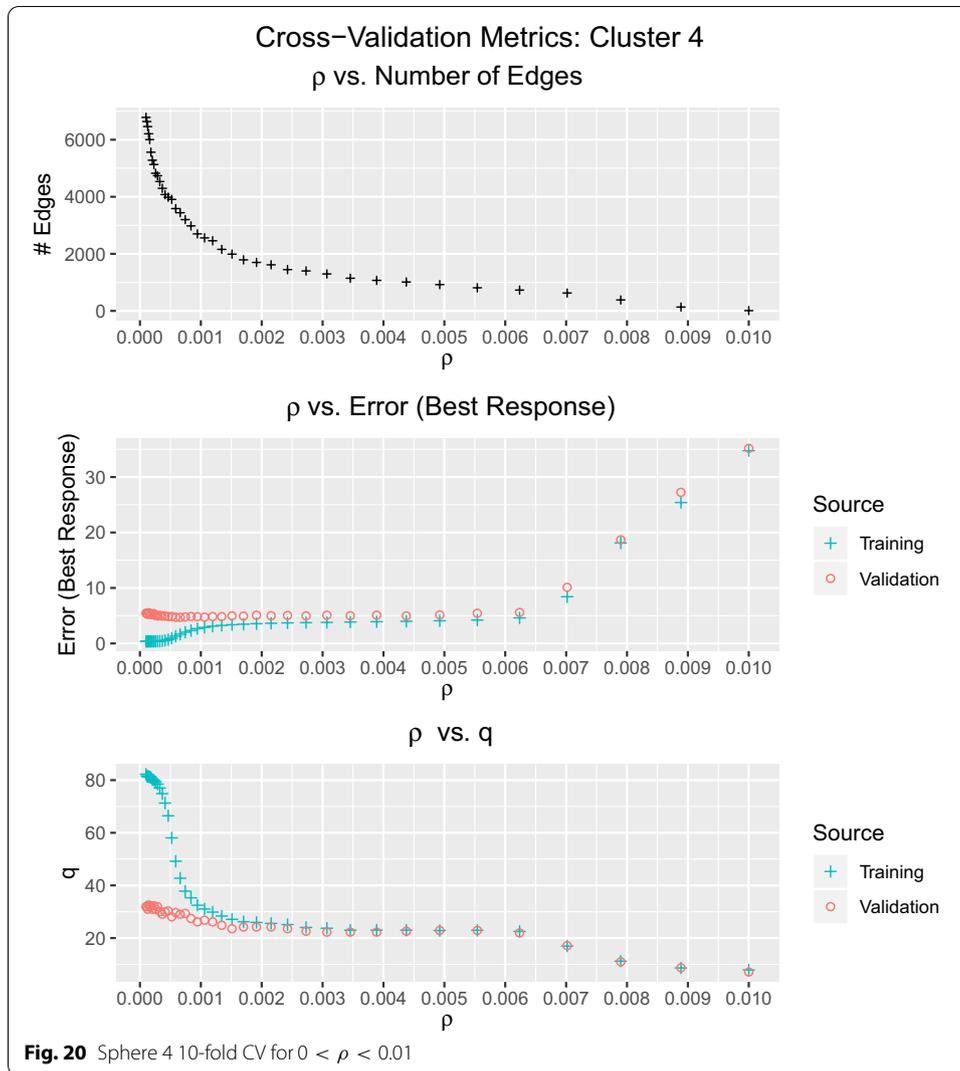


Appendix C: Visualization of LIG networks

Within the main body of the paper, the LIG networks for Spheres 1 and 2 are shown in Figs. 3, 4, respectively. Here, we show the LIG networks for the other two spheres. These are shown in Figs. 21, 22.

Appendix D: Visualization of cross-party edges

The cross-party edges for Spheres 1 and 2 are shown in Figs. 5, 6, respectively. Here, we show the cross-party edges of the remaining two spheres (Figs. 23, 24).



Appendix E: Visualization of most influential nodes

The most influential nodes for Spheres 1 and 3 are shown in Figs. 10, 11, respectively. Here, we show the most influential nodes for the remaining two spheres in Figs. 25, 26.

Appendix F: Application of ideal point models with social interactions

F.1 Cross-validation and model selection

Here, we show tenfold cross-validation results of applying the ideal point model with social interactions [28] 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

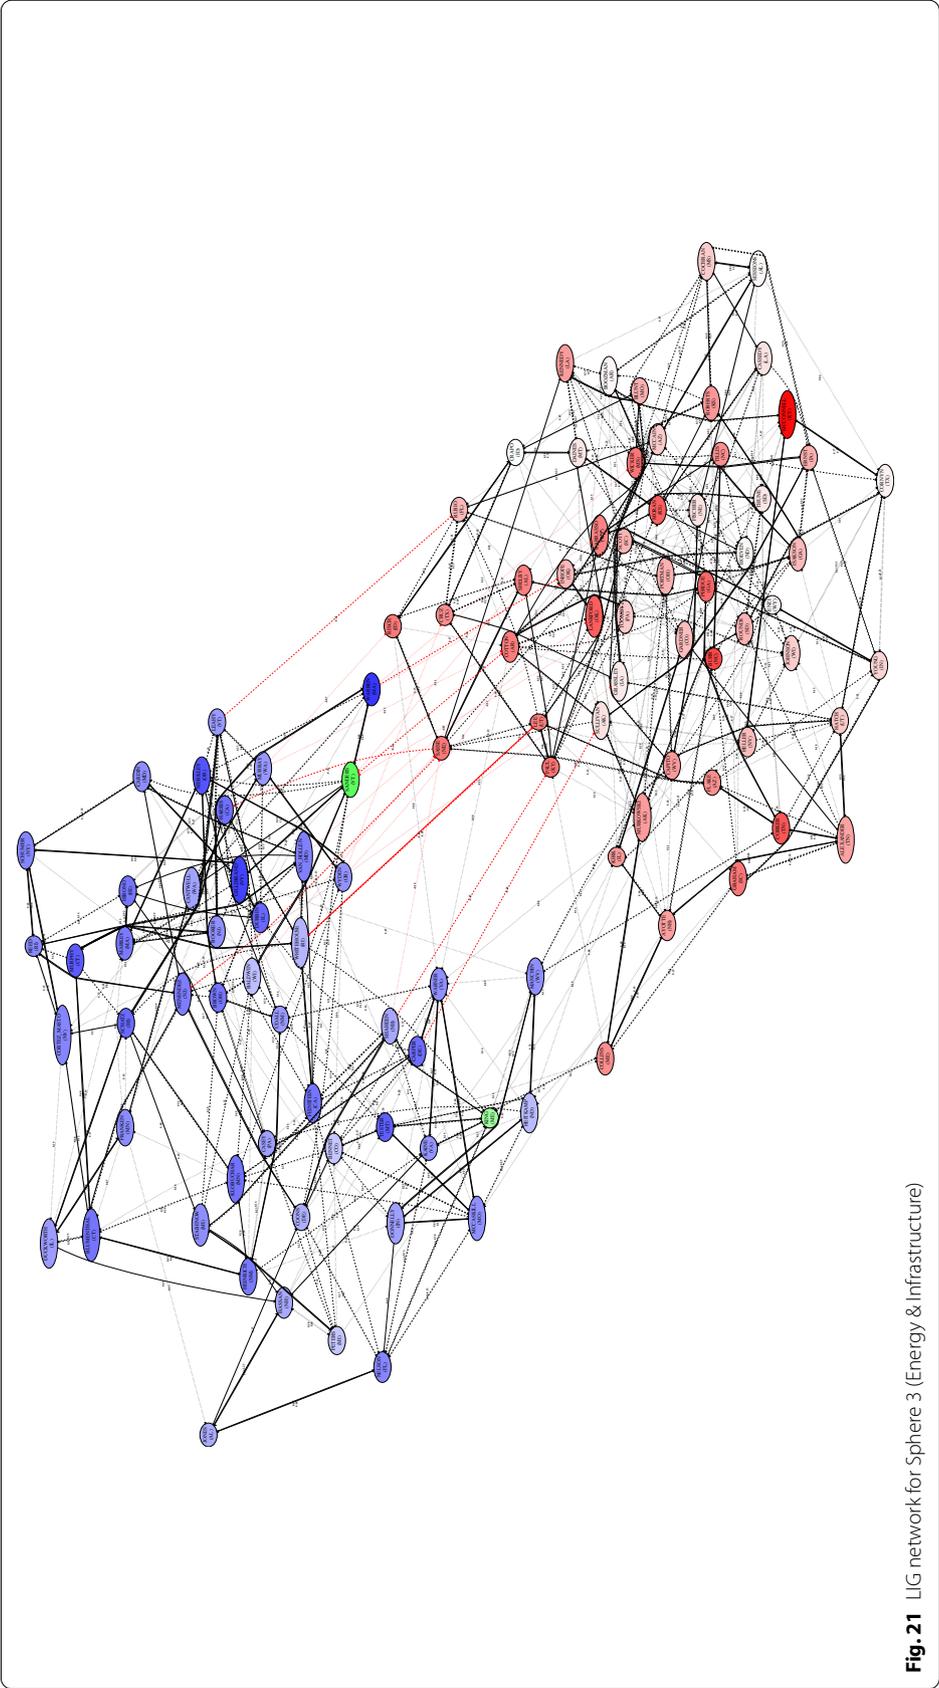


Fig. 21 IIG network for Sphere 3 (Energy & Infrastructure)

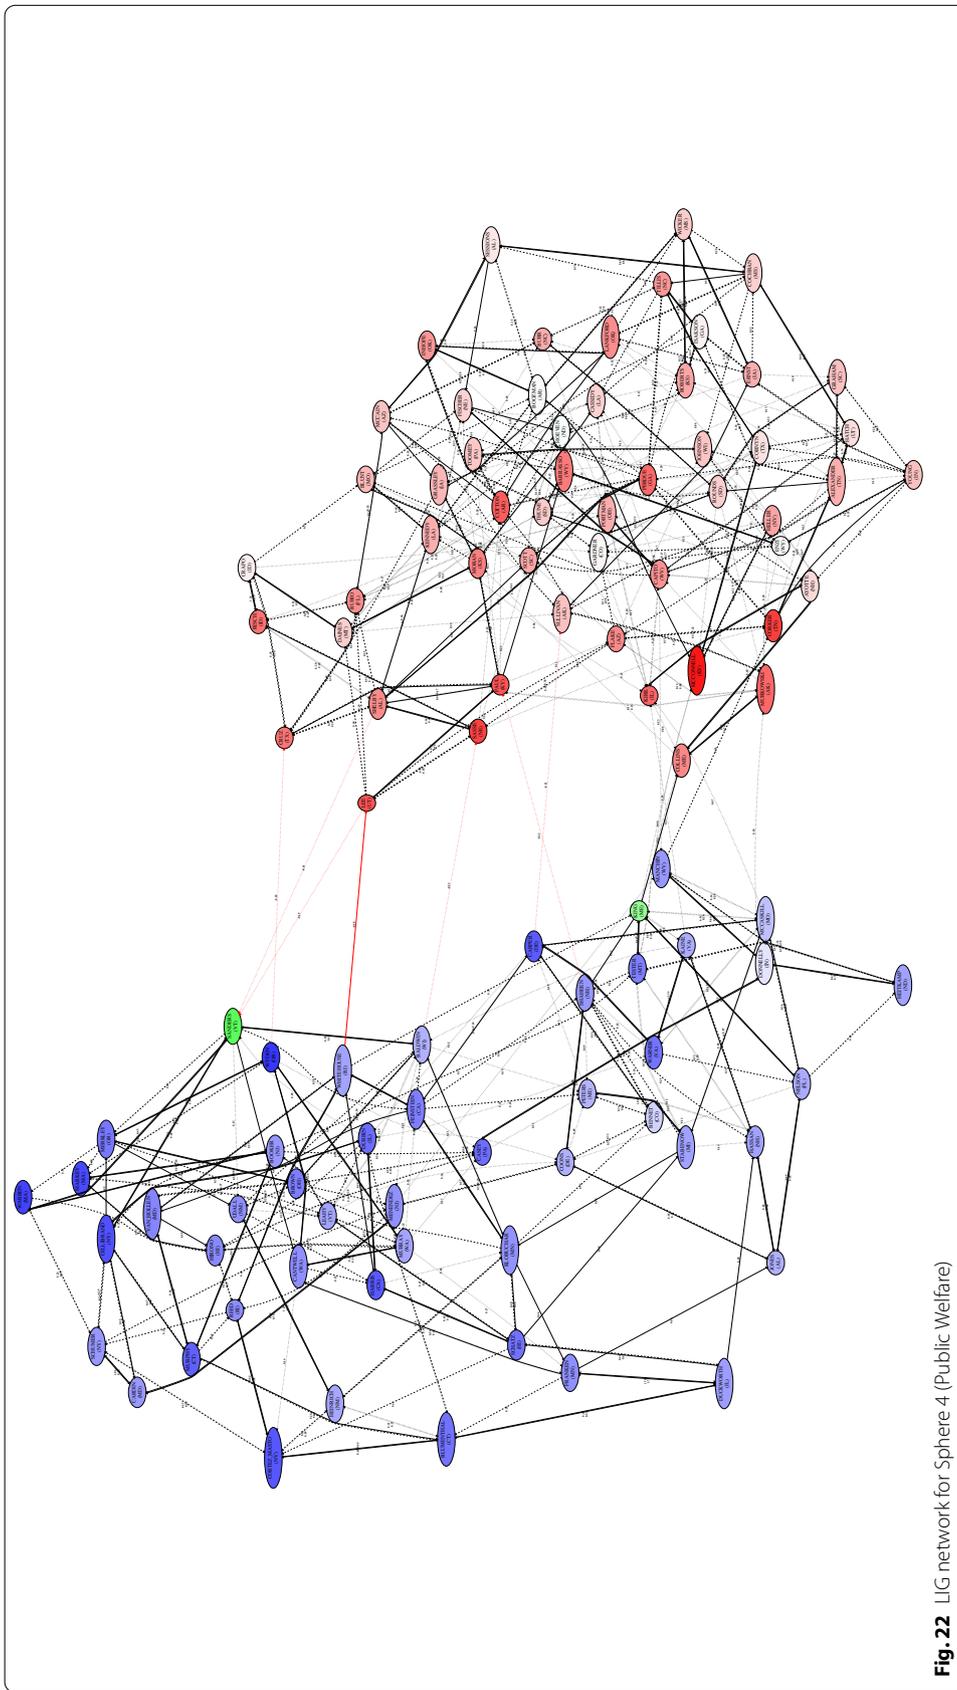


Fig. 22 LIG network for Sphere 4 (Public Welfare)

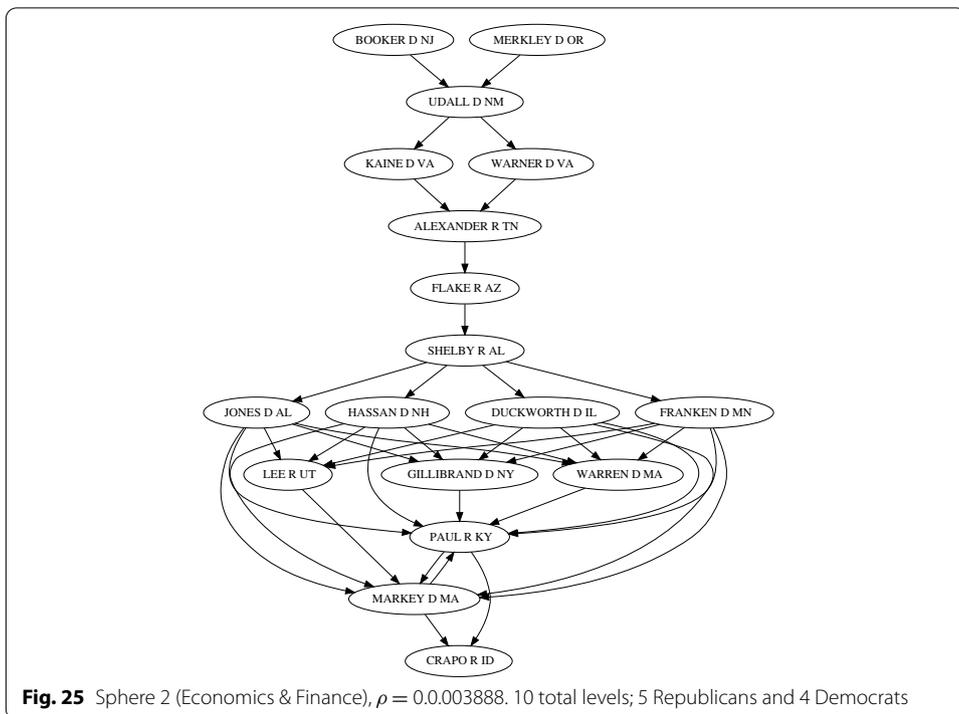
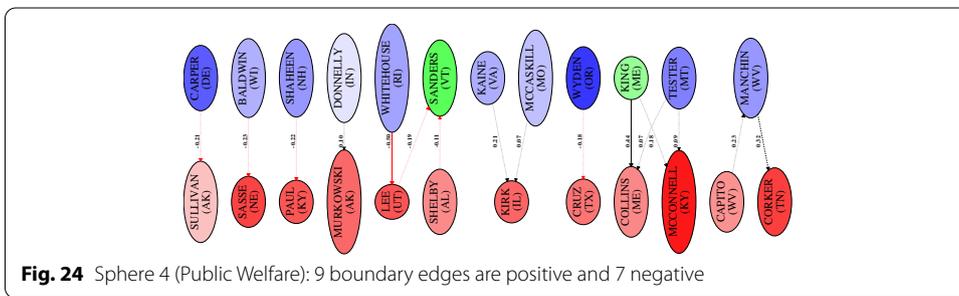
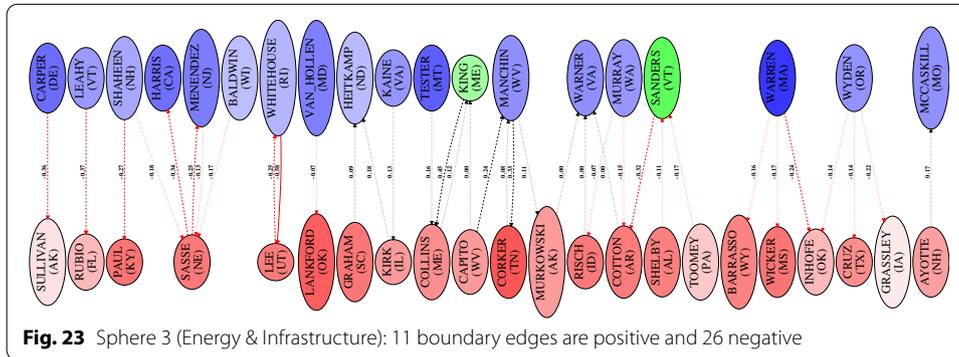




Table 6 Summary of cross-validation results for ideal point model with social interactions

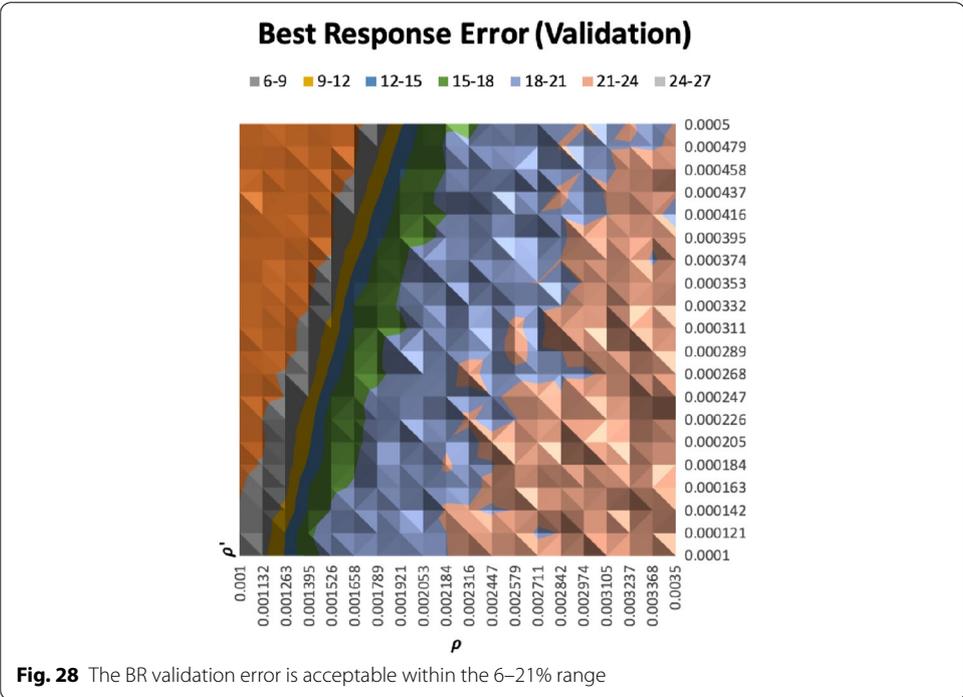
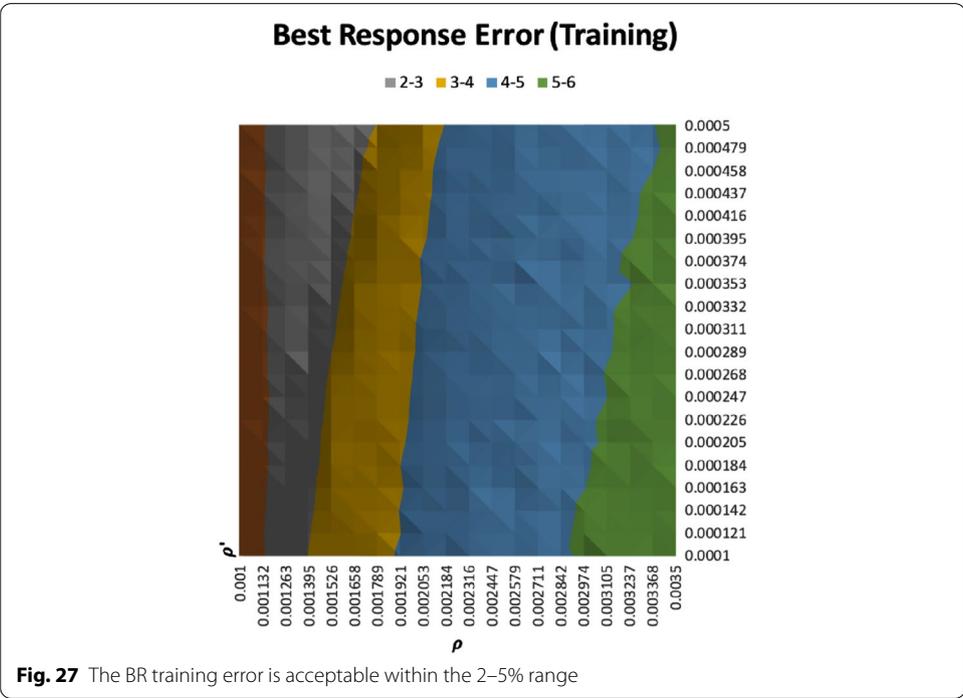
	Sphere 1	Sphere 2	Sphere 3	Sphere 4
ρ	0.002447	0.002447	0.002184	0.002053
ρ'	0.000395	0.000395	0.000374	0.000395
# Edges	1132	1234	1220	1499

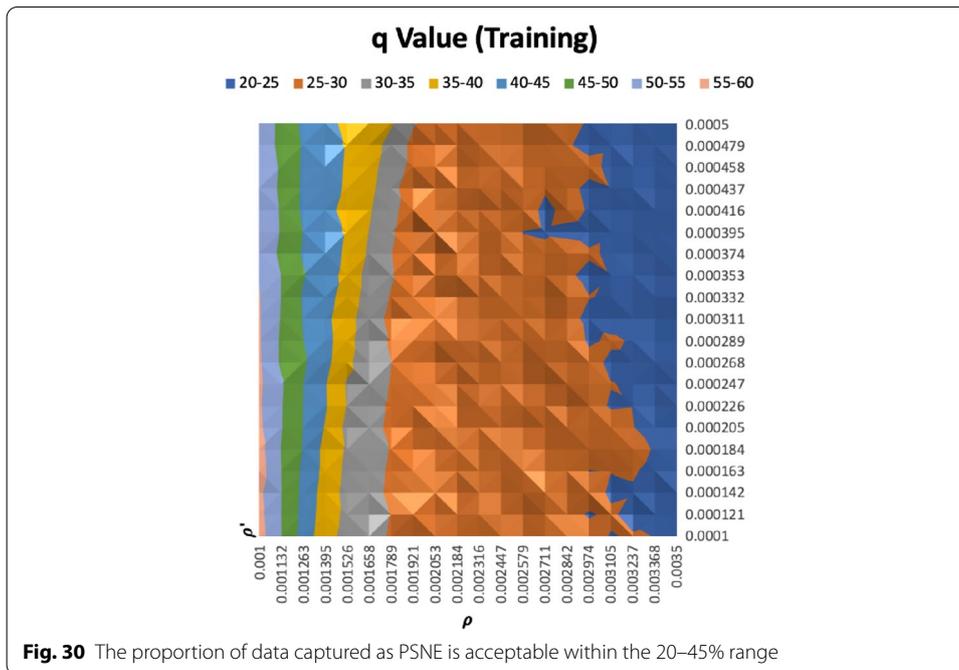
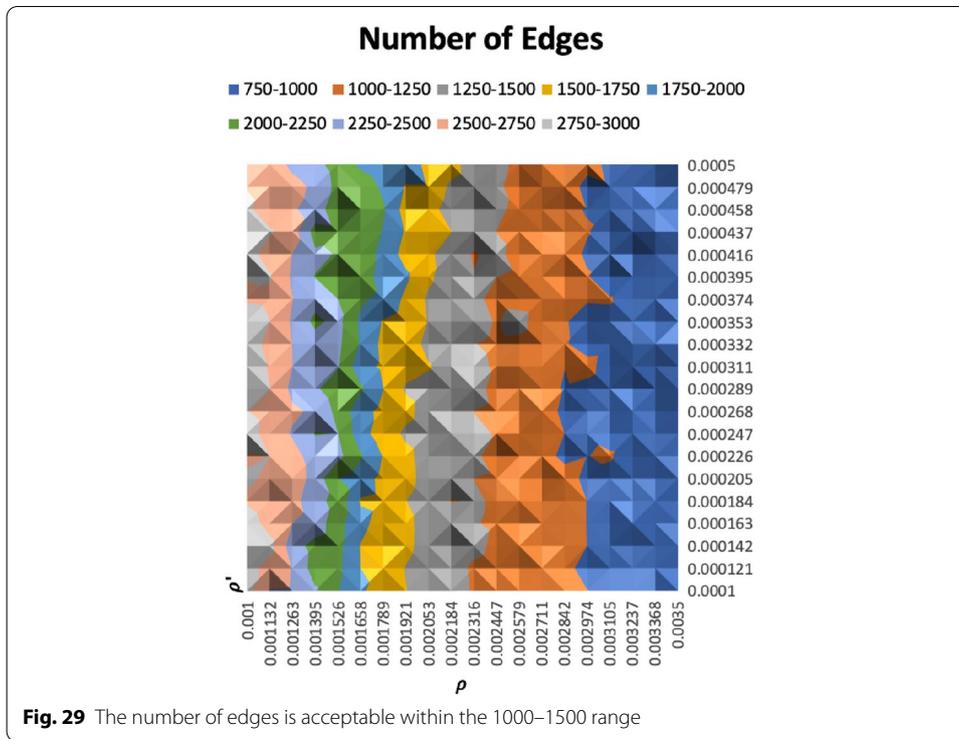
experiments were done over the following range of ρ and ρ' values: $0.001 < \rho < 0.0035$ and $0.0004 < \rho' < 0.0005$. The chosen values of ρ and ρ' are listed in Table 6.

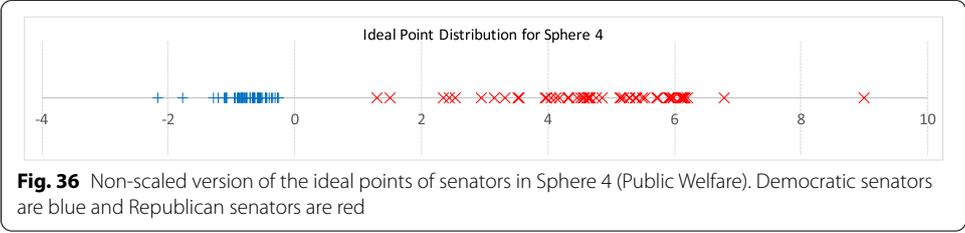
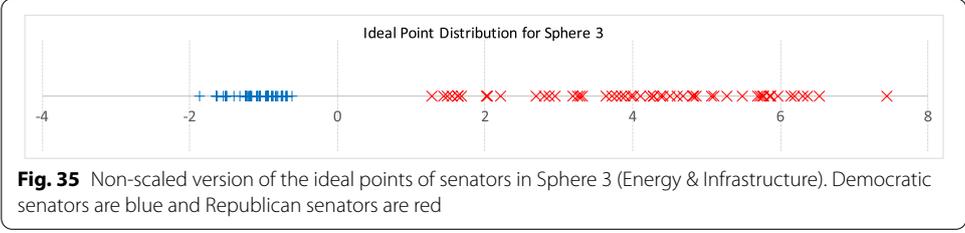
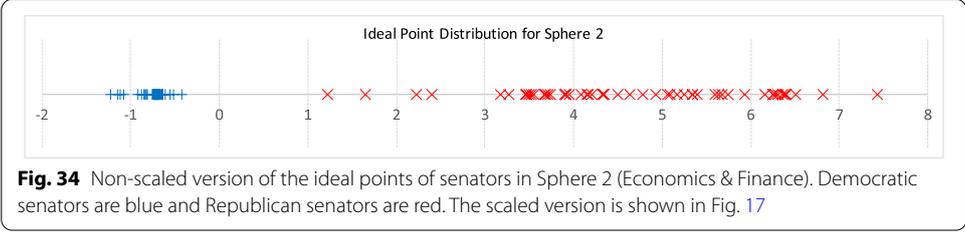
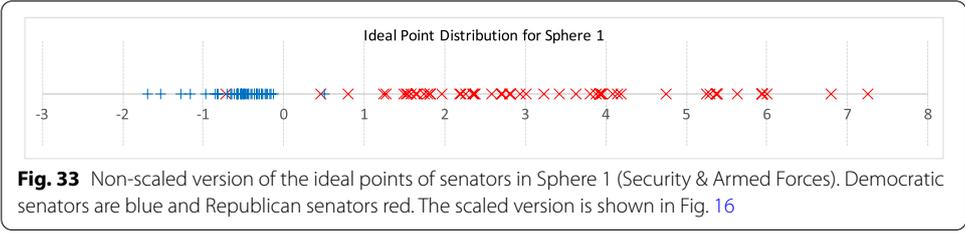
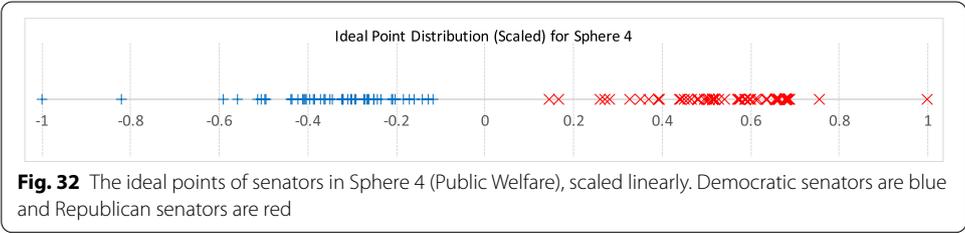
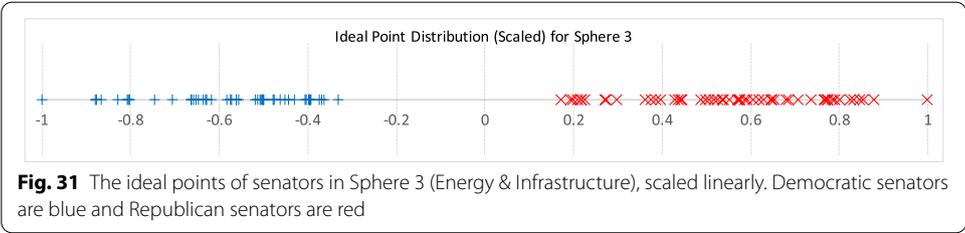
Figures 27, 28, 29, and 30 show the plots for the best response (BR) training error, BR validation error, number of PSNE, and proportion of data captured as PSNE for Sphere 2 (Economics & Finance).

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 Figs. 31, 32, respectively. Note that the ideal point distributions for Spheres 1 and 2 have been shown earlier in Figs. 16, 17, respectively. We also show the non-scaled versions of the ideal point distributions of the four spheres here (Figs. 33, 34, 35, 36).







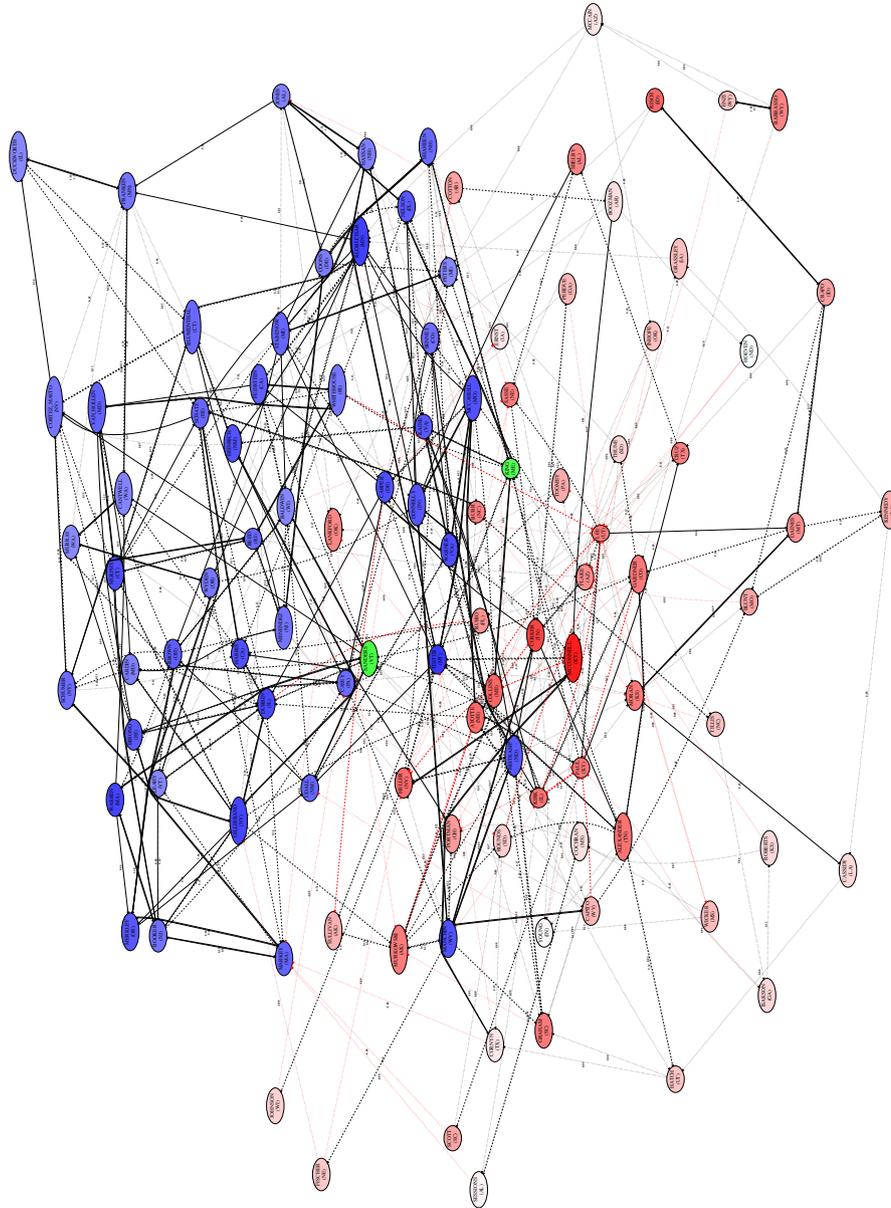


Fig. 37 Network for Sphere 3 (Energy & Infrastructure) under the ideal point model with social interactions. Only the strongest one-third of the edges are shown

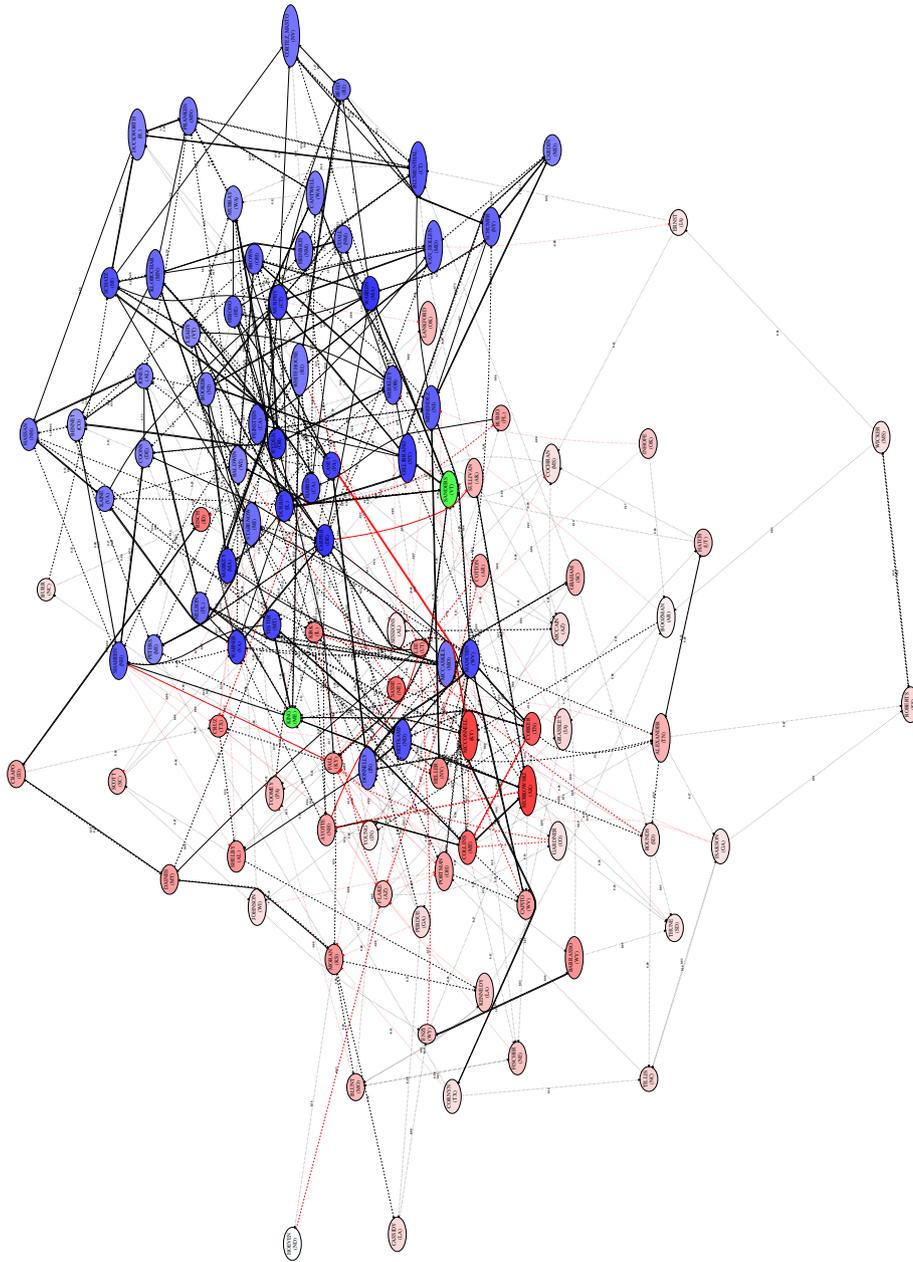
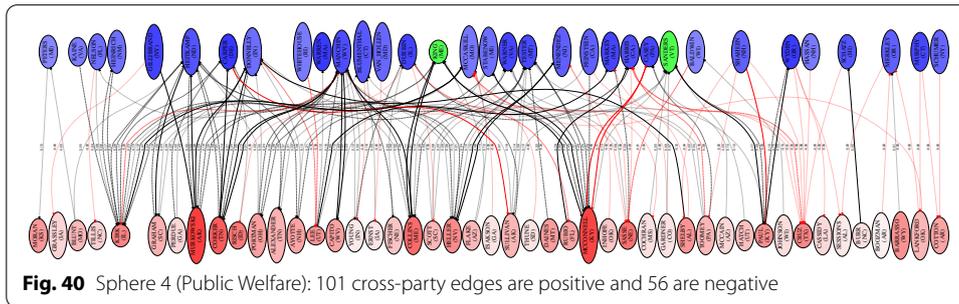
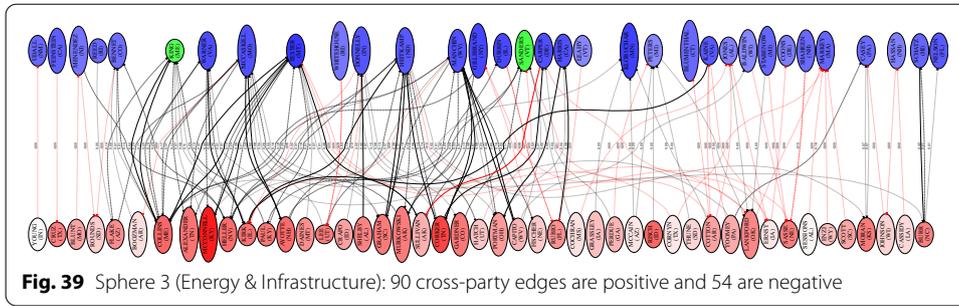


Fig. 38 Network for Sphere 4 (Public Welfare) under the ideal point model with social interactions. Only the strongest one-third of the edges are shown



F.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 (Figs. 37, 38). The networks for the first two spheres have been shown in Figs. 12, 13 within the main body of the paper.

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 (Figs. 39, 40). Figures 14, 15, shown in the main body of the paper, illustrate the cross-party edges for Spheres 1 and 2, respectively.

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 [28]'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

Table 7 (continued)

	Sphere 1	Sphere 2	Sphere 3	Sphere 4
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

Acknowledgements

We are grateful to Drs. Honorio and Ortiz for letting us use their codes [26]. We are also thankful to the anonymous reviewers of COMPLEX NETWORKS 2019 and the COSN journal for many helpful suggestions.

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 [44]. We have significantly extended that paper in this article. Following are the notable additions to the main body of the paper: "Preparing congressional roll-call data" section; LIG example and an example for each definitions in "The LIG model, Cross-validation and model selection for LIG, Influence weights and thresholds, PSNE computation" sections: Pure-Strategy Nash Equilibria (PSNE) Computation; and many figures. In addition, there are light revisions all throughout the conference version, but the following parts went through significant revision: "Learning Algorithm, Game-theoretic vs. structural centrality measures and Towards richer models: ideal point models with social interactions" sections. We also include an Appendix with a detailed literature review and many figures and tables.

Authors' contributions

Irfan and Phillips contributed by writing the manuscript. All three authors were involved with the experiments, with Phillips contributing the most. Irfan synthesized and guided the project and revised all iterations of the manuscript. Ostertag-Hill made valuable contributions to "Polarization in context" and "Towards richer models: ideal point models with social interactions" sections. All authors read and approved the final manuscript.

Funding

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

Availability of data and materials

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

Competing interests

The authors declare that they have no competing interests.

Author details

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

Received: 29 June 2020 Accepted: 1 February 2021

Published online: 21 March 2021

References

- Al Amin MT, 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, 2017; 1–9. IEEE.
- Atienza-Barthelemy J, Martin-Gutierrez S, Losada JC, Benito RM. Relationship between ideology and language in the catalan independence context. *Sci Rep*. 2019;9(1):1–13.
- Bezdek JC. Pattern recognition with fuzzy objective function algorithms. New York: Plenum Press; 1981.
- Bakshy E, Messing S, Adamic LA. Exposure to ideologically diverse news and opinion on facebook. *Science*. 2015;348(6239):1130–2.
- 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, 2011; 379–390.
- Christakis NA, Fowler JH. The Spread of Obesity in a Large Social Network over 32 Years. *N Engl J Med* 2007;357(4):370–379. <https://doi.org/10.1056/NEJMsa066082>. <http://content.nejm.org/cgi/reprint/357/4/370.pdf>.
- Christakis NA, Fowler JH. The Collective Dynamics of Smoking in a Large Social Network. *N Engl J Med* 2008;358(21):2249–2258. <https://doi.org/10.1056/NEJMsa0706154>. <http://content.nejm.org/cgi/reprint/358/21/2249.pdf>.
- Clinton J, Jackman S, Rivers D. The statistical analysis of roll call voting: a unified approach. *Am J Politic Sci*. 2004a;98(2):355–70.
- Clinton J, Jackman S, Rivers D. The statistical analysis of roll call data. *Am Politic Sci Rev*. 2004;98(2):355–70. <https://doi.org/10.1017/S0003055404001194>.
- Conover MD, 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.
- Cover TM, Thomas JA. Elements of Information Theory. Hoboken: John Wiley; 2006.
- DeGroot MH. Reaching a consensus. *J Am Stat Assoc*. 1974;69(345):118–21.
- Davis OA, Hinich MJ, Ordeshook PC. An expository development of a mathematical model of the electoral process. *Am Politic Sci Rev*. 1970;64(2):426–48. <https://doi.org/10.2307/1953842>.
- Desmarais BA, Moscardelli VG, Schaffner BF, Kowal MS. Measuring legislative collaboration: the Senate press events network. *Soc Netw*. 2015;40(1):43–544. <https://doi.org/10.1016/j.socnet.2014.07.006>.
- Farina CR. Congressional polarization: terminal constitutional dysfunction? *Colum L Rev*. 2015;115:1689.
- Friedkin NE, Johnsen EC. Social Influence Network Theory: A Sociological Examination of Small Group Dynamics, vol. 33. New York: Cambridge University Press; 2011.
- Garcia D, Abisheva A, Schweighofer S, Serdült U, Schweitzer F. Ideological and temporal components of network polarization in online political participatory media. *Polic Int*. 2015;7(1):46–79.
- Gerrish SM, Blei DM. Predicting Legislative Roll Calls from Text. *ICML'11 Proceedings of the 28th International Conference on International Conference on Machine Learning*, 489–496 2011.
- Gerrish SM, Blei DM. How they vote: issue-adjusted models of legislative behavior. *Adv Neural Informat Process Syst* 2012;25(1):2762–2770. <https://doi.org/10.1111/j.1751-1097.2010.00727.x>.
- Gómez S, Jensen P, Arenas A. Analysis of community structure in networks of correlated data. *Institut des Systèmes Complexes*. 2009;1(1):2–3.
- Garimella K, Morales GDF, Gionis A, Mathioudakis M. Quantifying controversy on social media. *ACM Transact Soc Comput*. 2018;1(1):1–27.
- Guerra PC, 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.
- Girvan M, Newman ME. Community structure in social and biological networks. *Proceed Nat Acad Sci*. 2002;99(12):7821–6.
- Granovetter M. Threshold models of collective behavior source. *Am J Soc*. 1978;83(6):1420–43.
- Glonek M, Tuke J, Mitchell L, Bean N. Semi-supervised graph labelling reveals increasing partisanship in the United States Congress. *ArXiv*, 1–14 2019.

26. Honorio J, Ortiz L. Learning the structure and parameters of large-population graphical games from behavioral data. *J Mach Learn Res*. 2015;16:1157–210. <https://doi.org/10.1016/j.asoc.2015.04.032>.
27. Hopkins DJ, Sides J. *Political Polarization in American Politics*. New York: Bloomsbury Publishing USA; 2015.
28. Irfan MT, 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.
29. Irfan MT, Ortiz LE. 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>.
30. Irfan MT, Ortiz LE. On influence, stable behavior, and the most influential individuals in networks: a game-theoretic approach. *Artific Intell*. 2014;215:79–119. <https://doi.org/10.1016/j.artint.2014.06.004>.
31. Jackson MO. *Social and Economic Networks*. Princeton, NJ: Princeton University Press; 2010.
32. Jenkins JA. Examining the bonding effects of party: a comparative analysis of roll-call voting in the us and confederate houses. *Am J Politic Sci*. 1999;69:1144–1165.
33. 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.
34. 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.
35. 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.
36. Leskovec J, Adamic LA, Huberman BA. The dynamics of viral marketing. *ACM Transact Web*. 2007;1(1):5.
37. 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.
38. MacQueen J. Some methods for classification and analysis of multivariate observations. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability 1: *Statist*, 1967; 281–297.
39. Morales AJ, Borondo J, Losada JC, Benito RM. Measuring political polarization: Twitter shows the two sides of venezuela. *Chaos Interdisciplin J Nonlin Sci* 2015;25(3):033114.
40. McCarty N, Poole KT, Rosenthal H. *Polarized America: the Dance of Ideology and Unequal Riches*. Cambridge: MIT Press; 2016.
41. Newman ME. Modularity and community structure in networks. *Proceed Nat Acad Sci*. 2006;103(23):8577–82.
42. Newman ME, Girvan M. Finding and evaluating community structure in networks. *Phys Rev E*. 2004;69(2):026113.
43. Olivares G, Cárdenas JP, Losada JC, Borondo J. Opinion polarization during a dichotomous electoral process. *Complexity*. 2019;6:2019.
44. Phillips AC, Irfan MT, Ostertag-Hill L. Spheres of legislation: Polarization and most influential nodes in behavioral context. *Complex Networks and Their Applications VIII*. 2020;882.
45. Poole KT, Rosenthal H. A spatial model for legislative roll call analysis. *Am J Politic Sci*. 1985;29: 357–384.
46. Poole K, Rosenthal H, Hare C. Alpha-nominate applied to the 114th senate. *Voteview* (August 16) 2015.
47. Ross TJ. *Fuzzy Logic with Engineering Applications*, 2nd edn., pp. 362–444. John Wiley & Sons, West Sussex, England 2004. [https://doi.org/10.1016/0042-6989\(71\)90197-0](https://doi.org/10.1016/0042-6989(71)90197-0).
48. Schickler E. Institutional change in the house of representatives, 1867–1998: a test of partisan and ideological power balance models. *Am Politic Sci Rev*. 2000;94(02):269–88.
49. Suchecki K, Eguiluz VM, San Miguel M. Voter model dynamics in complex networks: Role of dimensionality, disorder, and degree distribution. *Physical Rev E*. 2005;72(3):036132.
50. Wasserman S, Faust K. *Social network analysis: methods and applications*. Cambridge: Cambridge University Press; 1994.
51. Woon J. Bill sponsorship in Congress: the moderating effect of agenda positions on legislative proposals. *J Politic*. 2008;70(1):201–16. <https://doi.org/10.1017/S0022381607080140>.
52. Waugh AS, Pei L, Fowler JH, Mucha PJ, Porter MA. Party polarization in congress: A network science approach. arXiv preprint <http://content.nejm.org/cgi/reprint/357/4/370.pdf> 2009.
53. Zhang Y, Friend AJ, Traud AL, Porter MA, Fowler JH, Mucha PJ. Community structure in congressional cosponsorship networks. *Phys a Stat Mechan Appl*. 2008;387(7):1705–12.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.