Cues and Heuristics on Capitol Hill: Network-based Shortcuts in Legislative Decision-Making

Thesis Proposal by Justin H. Gross
Working Committee: S.E. Fienberg, D. Krackhardt,
J.N. Victor, C.R. Shalizi

May 19, 2008

1 The Logic of Shortcuts for Making Decisions under Uncertainty in the U.S. Congress

> The member [of Congress] facing decisions outside his area of expertise is up against the fact that the congressional experts, those he agrees with and those he disagrees with, know a good deal more about the subject than he, they argue on a higher plane, and they debate assertions of fact he must accept or reject largely on faith, lacking the background and time to do thorough research himself. Such resources as he might expend do not take him far, and he is likely in the end to abandon his evaluation of the issue for an easier evaluation of the expert he chooses to follow. [This]... is the basis for the most practical solution to the problems of normal decision-making (Matthews and Stimson [1975], p. 49).

Specialization and division of labor have long been the norm in both the House and Senate, with much of the daily business occurring in committee meetings and subcommittee hearings, and individual legislators attempting to carve out a niche or two and devote much of their attention to developing expertise in those areas. While the vast majority of proposed bills, amendments, and resolutions will never come to a vote, there are still plenty of moments when all will be expected to weigh in publicly on matters outside their realm of knowledge. At times, they will know well in advance that a piece of proposed legislation is scheduled to come to a vote, but at others, they will be caught off guard by a call to the chamber floor and must quickly get caught up on the relevant facts, including not only the most pertinant technical details, but also accounts of the strategic wranglings to which they have not been privy. Imagine yourself in the shoes of some such member of Congress, scurrying from meeting to meeting, back and forth from your home state to Washington. Suddenly you find out that a bill will be coming up for a roll-call vote[1] and although several members have made clear, via cosponsorship, committee participation, and floor debate, their intentions, most will not declare their positions explicitly until moment of the vote. How might you go about reaching a decision?

If members of Congress (MCs) themselves are understood to be agents with limited time and resources in need of efficient ways to determine their own stances, might they not look to colleagues whose voting patterns have matched their own rather than sifting through a messy assortment of far noisier information sources in order to reach a decision? In fact, members themselves have often reported the importance of such strategies, but the dynamics of such relationships have yet to be incorporated in a formal theory of legislative voting.

The current research will develop techniques for large-scale data analysis when interactions among actors, rather than being assumed away, are in fact the very object of study. In fact, there is a small but well-known literature of what I will call "interactional" models of legislative decision-making, but while the creation of data analytic tools such as NOMINATE (Poole and Rosenthal [1997], Poole [2005]) have allowed spatial models to dominate the research over the past two decades, no analogous methods have been offered within this interactional tradition, nor has any effort been made (as far as I am aware) to find comon ground between the two perspectives. Central to my argument that spatial and interactional analyses of legislative voting are not necessarily incompatible is the notion that members of Congress behave in ecologically rational ways (Gigerenzer et al. [1999]). That is, they can adapt to their environment, noticing and taking advantage of everyday tools in order to transcend the cognitive limitations all humans face. Thus, even if they ultimately wish to vote for policies that they expect to lead to outcomes closer to their own ideals than the status quo, it may be that the most efficient (and sometimes the only feasible) way to make such assessments is by exploiting the wealth of information in their network of colleagues and manifested through these legislators' own revealed decisions.

The history of legislative voting theory has been largely shaped by a confluence of trends in economic thought with those in statistical analysis, particularly psychometrics. As legislative scholars turned from mostly narrative accounts toward empirical methods during the mid-twentieth century, they took advantage of a host of new data analytic techniques. At first, this resulted in purely descriptive treatments, but before long, some researchers grew concerned that data accumulation was outpacing theory development (Kingdon [1977]) and sought to remedy this by taking a more disciplined approach to asking questions of data. Among social scientists, economists had been embracing mathematical models with the greatest zeal, and so it is no surprise that it is in the economic literature that the roots of formal models in politics are firmly planted. The spatial (Euclidean) theory of choice in political behavior took the notion of ideal location in physical space in what had been a problem in the analysis of optimal merchant locations (Hotelling [1929]) and abstracted it to social and ideological space, first for

[1] i.e. all voting members will clearly indicate their preference for the official record.

electoral voting (Downs [1957]) and then for the behavior of committees (Black [1948, 1958]).

Around the time that the spatial theory of voting was about to blossom, a quite different approach to the study of legislative voting was being explored. While rational choice theorists were seeking inspiration in classical economics and game theory, a number of other Congressional researchers were finding their own theoretical kindred spirits in a group of groundbreaking computer scientists and organizational theorists. As James Kingdon, one of these researchers explained, lawmakers, "like other decision-makers but perhaps even more than most, must make a large volume of complex decisions, while constrained by limits on time and cognitive capacity to do so without extensive study of each issue." Models that account for real limitations of this sort, Kingdon points out, should be "entirely familiar to readers of Herbert Simon"; they respect the "decisional overload" facing real members of Congress in allowing for "decision-making procedures that cut legislators' information costs and simplify their choices... [in] standard ways... which can be applied vote after vote" (Kingdon [1977]). An emphasis on the transmission of cues was common to the work of legislative scholars embracing this vision of lawmakers as "boundedly rational" actors. Matthews and Stimson, in particular, revealed through extensive interviews with members of the House, that representatives were relying heavily on the cues of trusted colleagues in their "normal decision-making" (Matthews and Stimson [1975]). In work by Cherryholmes and Shapiro, and by Kingdon himself, cue-taking is generalized to include the impact of various forces such as constituent pressure and even one's own previous voting record (Cherryholmes and Shapiro [1969], Kingdon [1973]).

In essence, spatial theories focus on the relative positions of MCs and their ideals, while the interactional theories focus on the connections among these actors and the interdependence among their choices. Common ground may be sought through work on the use of information in the reduction of uncertainty concerning policy outcomes.

2 Motivating Theory

2.1 The Policy Hunt—A Latent Game with Network-Based Shortcuts

In his informational theory of legislative decision-making, Keith Krehbiel distinguishes between policies and outcomes, suggesting that while members of Congress exert direct control over the former via legislative action, they obtain value (utility) from policies only through their associated outcomes, or "effects of policies upon their enactment and implementation" (Krehbiel [1991], p. 66). Thus the positions taken on these policies, whether strategic or sincere, are simply vehicles by which lawmakers attempt to achieve preferred objectives in some outcome space. Krehbiel's crucial insight is that individuals possess only incomplete information regarding how policies will map to outcomes (or probability distributions over outcomes) and their behavior should be modeled as such.

Krehbiel offers an empirical justification for introducing uncertainty:

> Other things being equal, legislators would rather select policies whose consequences are known in advance than policies whose consequences are uncertain. Under conditions of relative certainty, legislators can plan and make the most of credit-claiming (...[or] plan and implement blame-avoidance strategies to minimize losses.) Under conditions of relative uncertainty, however, surprise and the prospect of embarrassment lurk beneath any policy choice (Krehbiel [1991], p. 62).

Krehbiel's Legislative Signaling Game (LSG), developed in joint work with Tom Gilligan (**??**), incorporates uncertainty, asymmetric information among specializing legislators, and the strategic use of privately held information in sequential decision-making.

My own approach is similar in spirit to Krehbiel's, but takes the spatial theory as a point of departure for an interaction-based model of decision making by those members with little direct knowledge of a given proposal nor the the relevant expertise. Members of Congress may be thought of as engaging in a sort of non-zero-sum version of the Hot-Or-Cold[2] game, where the goal is to locate each proposal's image in the outcome space as precisely as possible, by relying on the cues of others, each with his or her own beliefs regarding its location. For simplicity, the outcome space may be thought of as a collection of allowable ordered n-tuples in some n-dimensional Euclidean space, with each component corresponding to a different aspect of the world that may be affected by a given policy. I am not convinced that the Euclidean approach is appropriate (for reasons that will be explored in the thesis), but it is the default in the extant literature, and serves at the very least as a rough representation of the outcome space. The goal for each player may be to make "correct" decisions as often as possible, in the sense of selecting the alternative (Yea or Nay) that brings her more value at each stage. Or it may be to maximize overall accrued value in the course of a session, so that supporting a proposal mapping to an outcome less preferred than the status quo causes greater regret the more lopsided the error. Various assumptions about the incentive structure may be tried out in order to examine their implications. For example, if one takes Krehbiel literally in the quote above, one may model cue-taking legislators as engaging in a game of minimizing the risk of embarrassment (a type of minimax strategy, where those not possessing much information have as a top priority minimizing the maximum backlash they will receive for the session's decisions). Consider how the vote on the 2002 joint resolution authorizing the use of force in Iraq may have been approached. Or, as one former senator told me, a key imperative in voting is "watching your backside relative

[2]In this children's game, one player must search for an object that has been hidden by the others. As the seeker moves about, the others shout out "getting warmer", "getting colder", "very hot", etc. The metaphor is loose; in the policy hunt, the cue-givers are fixed in their positions, with the cue-taker having a rough sense of their whereabouts, and they shout out whether they themselves are hot or cold, depending on whether they are closer in their ideals to the projected outcome of a proposal or to the current status quo. Based on their shouts, the cue-taker may attempt to detect the proposal's location in order to determine her own best vote.

to assuming every vote can attract millions in negative advertising, and wanting to be proven right on a vote over the long haul".

Krehbiel notes that the work on cue-giving developed by Matthews and Stimson, Cherryholmes and Shapiro, and John Kingdon is based on a premise he deems "consistent with that of LSGs". In the representation offered by these authors, "legislators devise strategies to cope with uncertainty", and in that regard, cue-taking and signal receiving are "conceptually the same", although "game theoretic characterizations tend to be more explicit about the strategic aspects of information transmission" (Krehbiel [1991], p.170). My goal is to formalize this earlier work in a manner conducive to both exploration, via computer simulation, of deductive implications, and empirical analysis using real legislative data.

If one starts with a spatial theory of voting and takes it to be not simply a way of summarizing voting patterns, but in fact a meaningful representation of each agent's internal decision-making process, there is little justification for assuming an objective outcome space shared by all. Indeed there is reason to believe that since each individual visualizes the world differently, there is not a unique outcome space, but rather as many outcome spaces as there are actors in the system. Each may be taken to have her own mental map of more and less favorable states of the world and her own beliefs about where others' ideal points lie in that perceived space. From this point of view, it makes little sense to ask about *the* location of a bill or MC, except to the extent that we may find a way to synthesize these individual maps into a sort of average schema.

In each Member of Congress L_i's perception of the outcome space, Ω_i, only one point is known with certainty, the agent's own ideal point, $\mathbf{0}_i = (0, 0, ..., 0)_i$, the origin with respect to which all other vectors are measured in her own mind (see figure 1). At the start of the game, legislator L_i may have fairly imprecise beliefs about where colleagues' ideal points are located relative to her own, but the culmination of each successive vote provides an opportunity for better learning their locations and thus updating her beliefs about their preferences. To this end, she may consider the probability distribution she places over the current (kth) proposal $p_k's$ image in her outcome space, as well as the image of the status quo q_k. The latter will tend to be more precise, as it is easier to assess the current state of the world on a matter outside one's own expertise than it is to predict the likely outcome of a policy about which one is ill-informed. Nobody, not even the initial proposer of a measure, has perfect information pinpointing its image in the outcome space. Rather, some legislators have more precise beliefs than others about how certain proposals map to outcomes. That is, for any given policy, each member has her own prior probability distribution over its location in the outcome space, that may be updated as information is accumulated. A rational L_i, with unlimited cognitive power, may be thought to start out with some (possibly vague) prior $\pi_i(p_k, q_k) = \pi(O_i(p_k), O_i(q_k))$, then observe the votes of a subset of colleagues, $v_{\mathbf{j}k}$, $\mathbf{j} \subseteq \{1, ...N\}$ on a single proposal p_k and use this information to update her beliefs via Bayes' Rule:

$$\pi_i((p_k, q_k)|v_{\mathbf{j}k}) \propto \pi_i(p_k, q_k)L(v_{\mathbf{j}k}|(p_k, q_k)),$$

where the second factor on the right is the likelihood of the observed votes as a function of $(O_i(p_k), O_i(q_k))$, the location of the image of proposal and status quo in her outcome space.

(Figures 1–2 about here)

While it would be cumbersome to work with hundreds of separate mental maps of outcome spaces within a formal model, few should disagree with the claim that this provides a more realistic picture of the actual decision-making process. Not only are predictions of outcomes and perceptions of one anothers' ideal points bound to be heterogeneous; the very nature of individuals' outcome spaces are no doubt quite varied and would appear quite foreign even among political allies. To the extent that I will work directly with this mental map model, it shall only be in order to investigate the conditions under which ignorance may be rational; that is, if the account above is taken to be a reasonable approximation of the decision-making process by those members of Congress not involved in the formative stages of crafting a policy, it will be far more efficient to follow a simple decision rule based on received cues, than to attempt to gather extensive information oneself or actually go through the exercise of repeatedly updating beliefs via Bayes' Rule in an attempt to carry out the Policy Hunt.

2.2 Constructing a Formal Model

On each roll call vote, those without direct knowledge or preexisting interest will rely on their voting network and social network of colleagues (plus in some cases, an external network consisting of constituents, lobbyists, media, and advisors) to draw helpful cues on how to vote. For now, I will use a generic decision function, as more thought will be necessary to determine what would make sense here. (For example, (1) whether agents should use some weighted combination of the information, say via logistic regression, or some hierarchical decision-tree that considers certain components sequentially, and (2) whether a continuous value (utility) function or simple binary reward (1 for correct, 0 for incorrect) shall be assumed, or instead the value function be dispensed with altogether in favor of an automated adaptive decision-making rule.)

Put simply, the cue-taker's task is to transform a binary code with missing bits (votes of colleagues), together with attributes of these colleagues and the proposal, into a 1 or -1 policy decision. Thus the action space $D = \{-1, 1\}$, where -1 corresponds to a "Nay" vote, and 1 to "Yea". A decision function δ maps an information set into $D, \delta : I \to D$. Specifically, the available information on proposal k may be written as $I_k = (\mathbf{v}_k^*, \mathbf{A}, \alpha_k, \mathbf{H}_{ik})$, where $\mathbf{v}_k^* \in \mathcal{V} = \{(v_{1k}, v_{2k}, ...v_{Nk}) : v_{jk} \in \{-1, 0, 1\}\}$ is the incomplete list of how MCs will vote on roll call k, with -1 or 1 for every potential cue-giver, and 0 for every cue-taker. $\mathbf{A} = (\mathbf{a}_1, \mathbf{a}_2, ...\mathbf{a}_N)$ is an attribute matrix, the concatenation of attribute vectors for each of the N legislators (eg., $\mathbf{a}_i = (i'$s party, i's state, i's years of service, i's leadership role if any, i's committees$))$, α_k are proposal-specific attributes, and $\mathbf{H}_{ik}$ is the history of all previous roll call votes up to, but not including the kth, as L_i recalls it, including her current assessment of how bills signed into law have translated into outcomes. We might generalize the information set by adding private knowledge of L_i about proposal p_k and its probable mapping to the outcome space; this would allow cue-takers to have the option of following their own signal, i.e. not rely purely on cues of colleagues.

3 The Data

The data are publicly available via a number of online sources (Thomas-website, Voteview-website), as well as cleaned versions of public databases, generously shared by James Fowler, a political scientist at the University of California, San Diego. I have gathered some information, such as topics and debaters myself. Analysis will focus on the 2003-2008 U.S. Senate (108th – 110th), and perhaps a session of the House of Representatives.

Variables include:

1. Numbered roll call votes (date, time, principal proposal, nature of the motion, votes of each member, President's position if taken, committees of jurisdiction, sponsor and cosponsors, topic words)

2. All bills, amendments and resolutions, most of which never reach a vote (sponsors and cosponsors with date of signing, dated records of major actions taken, related actions in the other chamber)

3. Debaters on each motion (collected from CSPAN archives (CSPAN-website))

4. Committee information (members, chair, ranking minority member)

5. Caucuses (primarily in the House)

4 Preliminary Results: The Predictive Success of Simple Heuristics

Once we set to the task of predicting votes on each proposed measure from a small subset of votes deemed to be early revealed preferences, we have many options for how to proceed. For every roll call k, we may partition the MCs into two sets, L_k and L'_k, where L includes the indices of all those legislators taken to be early position-takers (e.g., sponsor, cosponsors, members of the committee of jurisdiction, and debaters on the floor), and L' indexes everyone else.

Possible classifiers for the votes of those in L' are as varied as the decision functions they reflect. As a first step, I propose two extremely simple ones, which I call the Buddy System and nearest-neighboring-Vote (nnV). The two are inspired by notions of fast-and-frugal heuristics (Gigerenzer et al. [1999]) and exemplify what are known as naive classifiers in machine learning. Despite their simplicity, the rules classify extremely well given, I would argue, fairly little information in comparison with other vote prediction algorithms. Both are fit via cross-validation in order to establish out-of-sample validity and avoid overfitting.

4.1 The Buddy System (BaB)

Under the Buddy System, each MC's "best available buddy" (BaB) is the colleague whose vote is the best classifier of the MC's own vote (i.e., their agreement rate is highest.) There are generally a number of very good buddies for each legislator, and the clustering of these good classifiers is something to be explored in the future.

Using all the roll calls in the training set, I order each MC's buddies from best to worst. (For now, only those who vote similarly are considered buddies, but of course those who tend to vote exactly opposite most of the time will also tend to be good classifiers of one another's votes.) I assume that each cue-taking MC will vote the way that his or her best buddy among potential cue-givers votes. Eventually, I will use actual early position-takers as likely cue-givers, but for now, I simply select a subset at random as pseudo-cue-givers. The number of cue-givers may be varied, as well as the number of roll calls held out for testing. Results are included for leave-one-out cross-validation based on 25 pseudo-cue-givers for each bill. Compare these with the 80-85% classification rate typical if we predict each member to vote with the party majority.

(Figure 3 about here)

4.2 Nearest Neighboring Vote (nnV)

The Buddy System, predicts an MC's vote with the corresponding vote of the most similar colleague whose vote is known. Another approach would be to look for similarities among roll calls. The logic here is that for each roll call in the test set, we can identify one in the training set that most resembles it. Here too, the identified cue-givers are the means by which similarity is established. One can imagine MCs learning the preferences of a number of colleagues and noting (consciously or unconsciously), "Oh, it's that sort of bill."

Here, the entire set of cue-givers (or pseudo-cue-givers for this first-cut attempt) can be be used, rather than just a single "best buddy" cue-giver. On the other hand, only a single "best" roll-call from the training set will be used to predict the missing votes on each roll call in the test set; this is chosen to be the one on which the cue-givers voted most similarly to the roll call being predicted, with ties broken at random

Both techniques are of the nearest-neighbor type. A single best clue is being used, rather than a weighting of many clues. This sort of predictor tends to do very well and, despite its simplicity, there is reason to believe it may actually come close to the type of decision-making mechanism at play. Gigerenzer and Goldstein refer to a "Take the Best" heuristic, by which actors may simply "bet on one good reason" in choosing an action (Gigerenzer et al. [1999]).

In the figure below, I show the results of classifications given varying size of the pseudo-cue-giver set. Each simulation proceeds by setting aside one roll call at a time, randomly selecting the indicated number of "cue-givers", and then choosing the nnV from among all other roll calls constituting the current test set. The "cue-takers" are predicted to vote as they do on the nnV and the percent correctly classified is recorded. This is repeated for all 675 roll calls and the histogram indicates the resulting distribution of classification rates.

(Figure 4 about here)

5 Work Ahead

5.1 Methodological Challenges

5.1.1 Incorporating Context: Heterogeneity over Types of Motions, Committees of Jurisdiction, and Topics

Major large scale analyses of legislative voting share the simplifying assumption that proposals are homogeneous and independent of one another. Bills, amendments, judicial appointment confirmations, and motions to end debate are all treated as if subject to identical decision-making processeses. Omnibus bills become interchangeable with the dozens of proposals to amend them. Models ignore clustering of votes based on the specific topic of legislation under consideration. This is not surprising, given the additional complexity of the alternative, but such analyses sacrifice much information that is likely relevant to actual decision-making.

The tradeoff between pooling all votes on one hand and creating separate models for voting under different conditions, on the other, is a familiar type of quandary in statistical modeling; a balanced approach that takes advantage of the large number of observations in the former and the specificity of the latter can be found in multilevel (aka hierarchical) linear models. This basic outlook will be useful here; legislators may be consdered to have some underlying propensity to take cues from one another under arbitrary conditions, with a random effect based upon the particular setting. Adapting the approach to use in network inference remains a challenge to be addressed in the dissertation. I envision the underlying use of one's network of colleagues as something like the following:

$$V_{ik} = f\left(\mathbf{v}_k^*, \mathbf{Z}_{cue-taking}, \gamma_{ik}\left(\mathbf{Z}_{context(k)}^{setting}\right), \epsilon_{ik}\right), E(\epsilon_{ik}) = 0$$

with legislator $i's$ vote on bill k a function of the latent baseline cue-taking network $\mathbf{Z}_{cue-taking}$ (itself based on proximities in the social, trust of judgment, and ideological networks), and a proposal-specific random effect γ_{ik} based on context-specific matrices $\mathbf{Z}^{setting}$ of predicted departures from the baseline propensities to vote together. For example, if we take the settings to be determined entirely by committees of jurisdiction on a bill, the random effect would allow a legislator to be more or less influenced by certain colleagues' opinions depending on the relevant committee. The random setting-specific effects, as well as the idiosyncratic error ϵ_{ik}, will, depending on the model (likely logit or probit), be assumed to have some appropriate simple distribution. In crafting an classifier $\hat{V}$ or estimator $\hat{p}(V) = Pr(V = 1|\text{information set})$ we will have to settle for the observed co-voting network X_{cv}, co-voting within particular settings, together with possible proxies for ideological and social proximity, such as the tendency to cosponsor one another's substantive and non-controversial bills, respectively.

5.1.2 Drawing Inferences Based on Affiliation Networks: Bipartite, Tripartite and Beyond

Although it is convenient to refer to an underlying co-voting network, this is really something of a fiction (in more than just the usual sense in which all models are fictions). For each ordered pair (i, j) of legislators, there may be some underlying latent inclination of i to follow j's lead on an arbitrary vote if the former has little information on the proposal in question. The corresponding entry in the hidden cue-taking adjacency matrix would contain the unconditional probability that i votes with j. An appropriate estimate of this matrix could be taken from a complete survey of legislators, preferably after every single vote, to gauge the extent to which they consult one another. While we might question their recollection or the veracity of their responses, the survey instrument would at least be well-matched to the network under study. Instead we must settle for a full record of the votes. As a first cut at modeling, I use the symmetric matrix $\mathbf{X}_{CV}$, where each entry $\mathbf{X}_{CVij}$ gives the fraction of roll calls on which i and j vote together. This co-voting matrix is not a bad place to start and is the basis for the simple yet accurate Buddy System and Nearest Neighboring Vote predictions mentioned earlier. There are a number of problems with using it in modeling possible mechanisms of cue transmission, however. Several drawbacks are typical of models that oversimplify network dynamics: the adjacency matrix need not be symmetric (the degree to which it will be based on reciprocity and may in fact produce an identifiable parameter), the co-voting matrix does not give any indication of dependence among dyads on individual votes, and so on. However, there is a more serious problem; in contrast to something like cosponsorship, which may convincingly be treated as indicating support from one legislator to another, voting records provide an exceptionally weak signal about dyads. Virtually all legislators can be expected to vote Yea or Nay on a roll call, but this may only indirectly tell us something about pairs of members finding themselves on the same side of a vote.

In fact the co-voting adjacency matrix actually arises from collapsing what is known in the social network literature as an *affiliation network* (Wasserman and Faust [1994], ch. 8). Just as people may be connected via parties jointly attended, organizations in which they are members, or movies in which they appeared, legislators may be viewed as affiliated via their set of shared votes. Treated in this way, the network should be represented as a bipartite graph, G, with two distinct types of nodes, legislators L, and proposals P, such that $G = (L \uplus P, E)$ and edges only between sets rather than within, so each $e = (l, p), l \in L, p \in P$, for each $e \in E$. This particular affiliation network is peculiar, though; if one simply takes members to be affiliated through mutual support for a motion, a lot of information is thrown away. If, on the other hand, we also allow affiliation through joint opposition to a motion, this induces a special pattern in the set of motions. Rather than following the example of attendance at common events, it is as if the actors choose to attend nearly every party, but are connected to one another via the particular room they hang out in once (the Yea room or the Nay room).

The affiliation network is complicated further by the fact that proposals cluster together within main threads of legislative action. Some motions are independent, while others come in a flurry of amendments, counteramendments, and procedural motions all related to some principal bill. For the most part, one can assume that those in the know about one motion in the thread will be familiar with the issues involved in any others, at least if the amendment is germane.

(Figure 5 about here)

If we are to include settings such as the committees through which these threads of legislation pass, then the bipartite graph becomes a tripartite graph. Little work has been done on making inferences on such higher order multipartite networks, and this will be an area in which I shall seek to contribute.

(Figure 6 about here)

5.2 Substantive Research Questions to be Addressed via the Relational (Network) Model

5.2.1 Congressional Mavericks and Renegades: In search of the most independent Senator

In the popular press, certain members of Congress are widely depicted as being independent thinkers (current presidential candidate John McCain, for instance). This reputation may have been earned by being consistently unpredictable over years of service, or purely on the basis of a few surprising stands taken on prominent bills, or may simply be a misrepresentation. It may be illuminating to define independence of decision-making in terms of statistical independence of voting behavior; that is, someone whose own vote is (by some metric to be decided) not particularly easy to classify by their colleagues' votes should be viewed as a true independent, whereas someone whose vote is easy to predict by colleagues' votes should not be. To do this convincingly, I will need to account for direction of cue transmission. A maverick may make up his or her own mind on virtually every bill, but when acting as a cue-giver, may attract a large number of votes—after all, a respected renegade may inspire imitation, and in so doing, become a leader.

To get a rough initial sense of how this will work, consider vectors of observed co-voting ratios for senators, based on roll calls during 2003-04, the vector on which my Best Available Buddy classifier was based. For senators who tend to be fairly unpredictable with respect to colleagues' votes, the highest co-voting ratios may be expected to be lower than those for senators who tend to vote in the same bloc most of the time. The variance of these ratios should be also be lower for independents and though we expect two modes for intra- and inter-party co-voting, the modes will be closer together for those not particularly sensitive to the pressure for party discipline. Glancing at histograms for several senators expected to be partisan and a few known as more independent, we see that none particularly stand out, with the possible exception of Zell Miller, included here for his reputation of voting consistently with the Republicans, despite being a Democrat. The modes of Miller's co-voting ratio histogram are relatively close together, indicating that he likely still votes with the Democrats on some partisan motions.

One senator, the self-proclaimed Rockefeller Republican Lincoln Chafee, has a qualitatively distinct histogram. His histogram is unimodal and while most Best Buddies vote together over 90% of the time, Chafee's Best Buddy only votes with him on 81% of roll calls. Even more surprisingly, while senators tend to have their lowest co-voting rates between 20-30%, Chafee's "Worst" Buddy votes with him 58% of the time!

(Figure 7 & 8 about here)

One might expect that someone in Chafee's situation, without any natural cue-givers on a wide selection of issues, would simply be forced to learn more about proposals outside his areas of specialization, or rely more on his staff for trustworthy assessments. By his own account, this was the case for Chafee.

> Yes, my particular status was such that we really had to scramble more to get good research on a vote. Usually there was another senator or two we could share information with and try to cast a sound vote. Often these "soulmates" would change depending on the issue; Lamar Alexander, John Sununu, Gordon Smith... and also often included Democrats with whom we were friendly and weren't automatic votes for their leadership.
>
> From my own experience in the Senate the quality of my office staff is the most important factor in dealing with a fast paced flurry of votes. Generally, there is at least enough time to call the office from the cloakroom and quickly brainstorm before voting. The quality of the advice from the office must include a mix of veteran experience, a knowledge of home state politics, and a good moral compass. And yes, there are always at least a few fellow senators whose votes we watched carefully to use as a guide.
>
> (L. Chafee, in personal email correspondence 5/6/08, 5/8/08)

5.2.2 Pivotal Blocs

Keith Krehbiel, in his book Pivotal Politics: A Theory of U.S. Lawmaking (Krehbiel [1998]), suggests a major source of so-called legislative gridlock to be the high threshold required to avoid the threats of filibuster, Presidential veto, and other maneuvers that make super-majorities so valuable in Congress. Krehbiel's model is set within a unidimensional spatial framework and replaces the celebrated median voter with the more general pivotal voter whose ideological inclinations make her support of proposals crucial to their success. The premise of his argument makes a lot of sense, and an analogous interactional model should seek to identify pivotal blocs under various scenarios. One advantage of the clustering perspective is that by recognizing covoting patterns, including nested blocs, we have the opportunity to develop a yet more nuanced view of pivots; in some situations, a large coalition may be pivotal, but then we can ask which bloc within the coalition may be pivotal in swinging the coalition one way or another. An interaction-based pivot would be classified on the basis of how well it predicts the votes of those in neighboring blocs, rather than by estimated location on an axis in the abstract policy space. A natural question will be who has potential influence based on their membership within pivotal blocs.

6 Appendix: Papers to be Spawned from this Line of Research (* indicates priority item)

1. **Predicting Success from a Weighted Measure of Sponsor Role, Number and Diversity of Cosponsors
2. Decomposition of Social and Ideological Predictors of Cosponsorship (with tendency to cosponsor noncontroversial, low salience resolutions as proxy for social proximity)
3. **Nearest Neighbor Bills & the Buddy System of Cue-taking: Simple Heuristics that Allow Surprisingly Accurate Predictions from Limited Information (including formal model of the rationale for cue-taking under very weak assumptions about constraints on time and attention).
4. **Congressional Mavericks and Renegades: In search of the most Independent Senator
5. *Latent Clusters in Congress: Using Network Science to Identify Pivotal Blocs of Lawmakers
6. Evolution of the Congressional Covoting Network Through Time
7. A Fully Generative Model of the Journey from Bill to Law

8. Predicting Covoting Patterns in the European Parliament from Staffer Interactions (will address missingness (due to nonresponse) of independent dyadic variables)

References

Duncan Black. On the rationale of group decision making. *Journal of Political Economy*, 56:23–34, 1948.

Duncan Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

Cleo H. Cherryholmes and Michael Shapiro. *Representatives and Roll Calls: A Computer Simulation of Voting in the Eighty-eighth Congress*. The Bobbs-Merril Company, Inc., 1969.

CSPAN-website. http://www.c-spanarchives.org/congress.

Gerd Gigerenzer, Peter M. Todd, and The ABC Research Group. *Simple Heuristics That Make Us Smart*. Oxford University Press, 1999.

Harold Hotelling. Stability in competition. *The Economic Journal*, 39:41–57, 1929.

John W. Kingdon. *Congressmen's Voting Decisions*. University Press of Kentucky, 1973.

John W. Kingdon. Models of legislative voting. *The Journal of Politics*, 39:563–595, 1977.

Keith Krehbiel. *Information and Legislative Organization*. The University of Michigan Press, 1991.

Keith Krehbiel. *Pivotal Politics: A Theory of U.S. Lawmaking*. University of Chicago Press, 1998.

Donald R. Matthews and James A. Stimson. *Yeas and Nays: Normal Decision-Making in the U.S. House of Representatives*. John Wiley & Sons, 1975.

Keith T. Poole. *Spatial Models of Parliamentary Voting*. Analytical Methods for Social Research. Cambridge University Press, Cambridge, 2005.

Keith T. Poole and Howard Rosenthal. *Congress: A Political-Economic History of Roll Call Voting*. Oxford University Press, Oxford, 1997.

Thomas-website. http://thomas.loc.gov/home/rollcallvotes.

Voteview-website. http://www.voteview.com/dwnl.htm.

Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications.* Cambridge University Press, 1994.

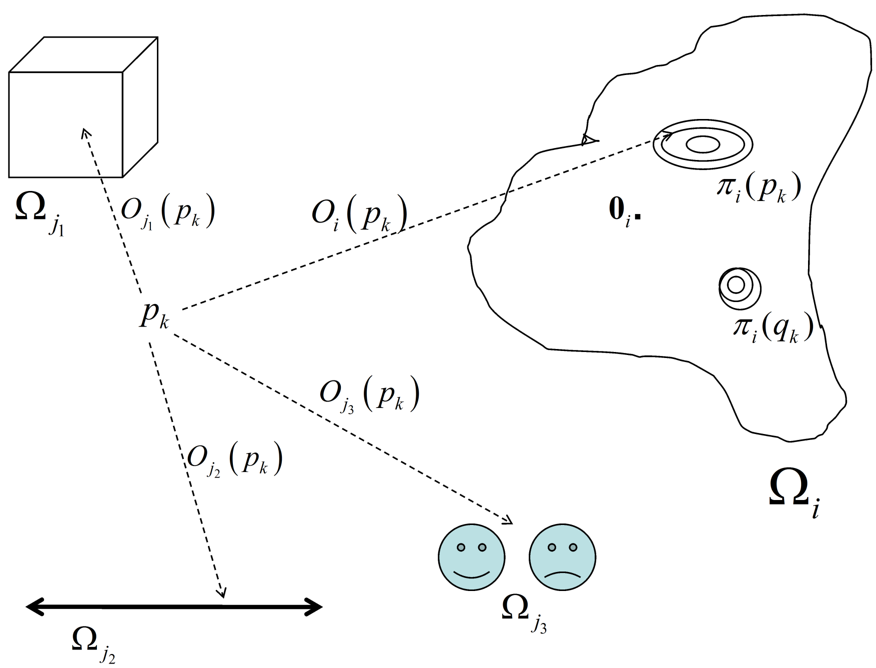

Figure 1: O_i maps proposal p_k into legislator L_i's outcome space Ω_i, but the true value of $O_i(p_k)$ is hidden from even L_i herself, due to her incomplete information about the proposed policy and the uncertainty as to how the policy would be implemented and enforced, and received by the public. L_i has some prior belief about where the policy would map to in Ω_i, represented by distribution $\pi_i(p_k)$, as well as some presumeably more precise beliefs concerning the current status quo, $\pi_i(q_k)$, but knows only her own ideal point 0_i with certainty. Meanwhile, other members try to gauge the value of the proposal in terms of their own, potentially quite distinctive, outcome spaces, Ω_{j_1} and Ω_{j_2}. The outcome spaces need not be Euclidean, though it is certainly convenient to assume them to be. Imagine the extreme case of legislator L_{j_3}, who sees the world in terms of black and white, good and evil. Any nuances are automatically filtered away (or subject to an unconscious threshold) and a kneejerk classification of proposals as good (smiley) or bad (frown) results. Witnessing colleagues' tendencies to vote for "bad" or "good" bills, they can then group the individuals themselves in this manner (with perhaps a third classification for those confused souls who sometimes seem to vote for good bills, and other times for bad ones).

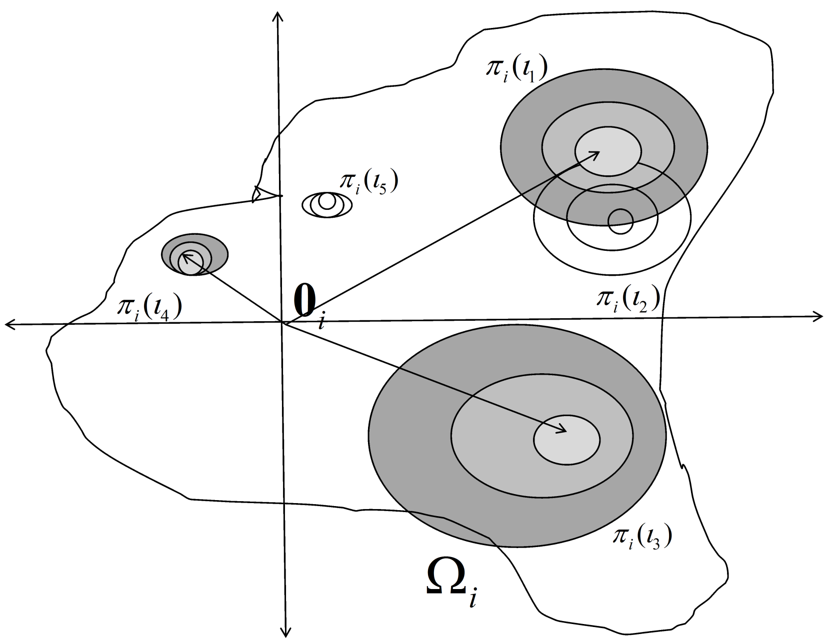

Figure 2: L_i has some belief about where colleagues' ideal points fall in his outcome space. When trying to locate a proposal, he may rely on information on the weight of his ties with current cuegivers (here L_1, L_3, L_4), together with their votes, to reach a decision directly rather than trying to triangulate the location of $O_i(p_k)$, and then calculate the probability that it is closer to his own ideal than $O(q_k)$. We may think about how each agent might use his network to span the outcome space in order to make certain decisions on "autopilot".

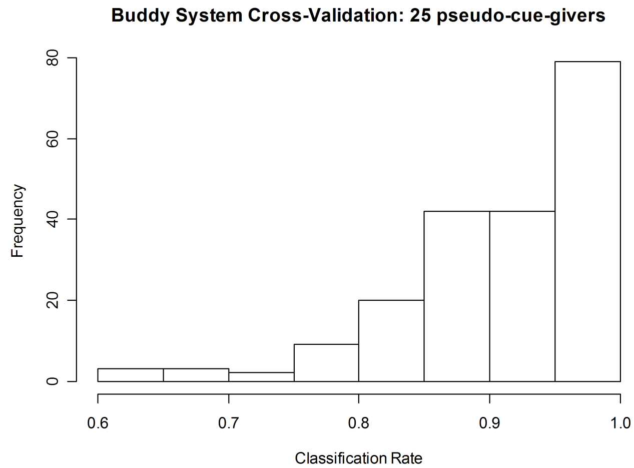

Figure 3: Classification rates using Best Available Buddy to Predict Own Vote

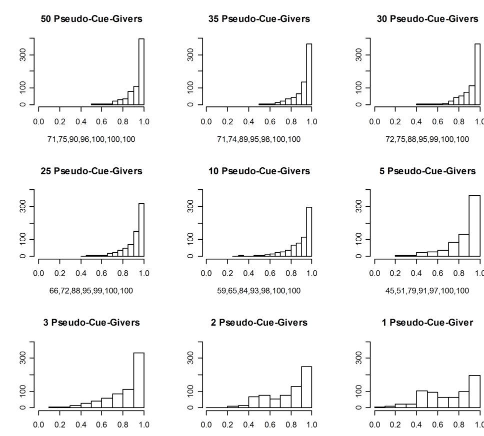

Figure 4: Classification rates using Nearest Neihboring Vote on all 675 roll calls during the 2003-04 (108th) Senate, varying the number of cue-givers allowed (Percentiles along the horizontal axes: 2.5%, 5%, 25%, 50%, 75%, 95%, 97.5%)

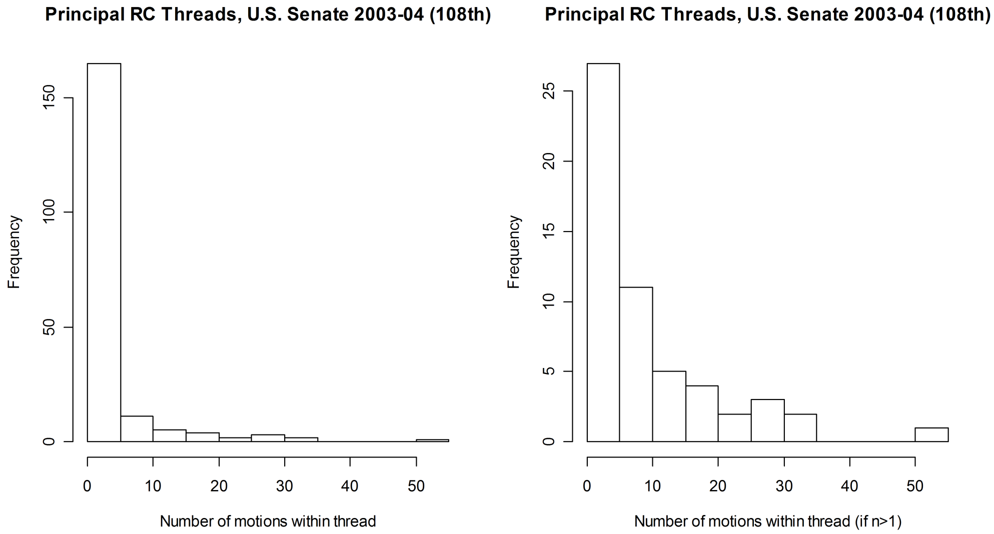

Figure 5: The 675 roll call votes in the 2003-04 U.S. Senate cluster into 193 threads of legislation, including 138 isolated votes, 38 threads with between two and ten votes, and one thread containing 51 separate votes.

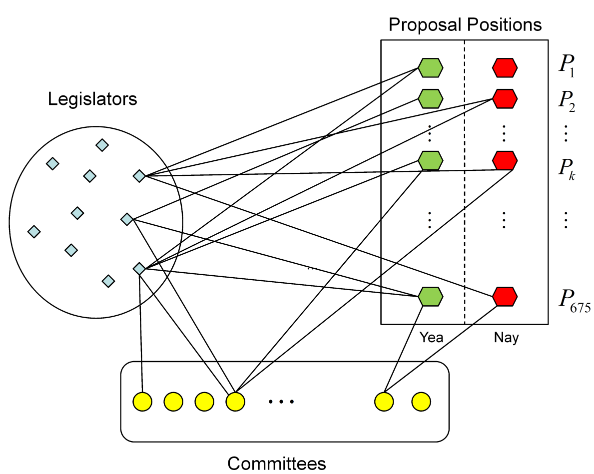

Figure 6: Tripartite Graph of the Network Among Legislators, Positions on Proposals, and Committees (Settings)

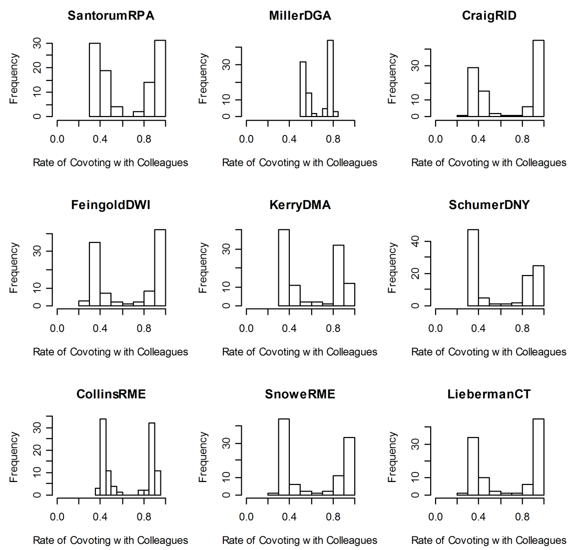

Figure 7: Partisan Conservatives, Partisan Democrats, and Moderate Independents

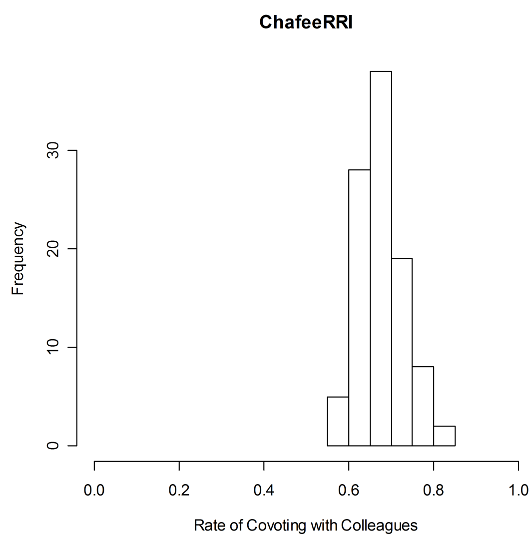

Figure 8: All of Chafee's co-voting rates lie between 0.58 and 0.81.