

# The Echo Chamber: Strategic Voting and Homophily in Social Networks

Alan Tsang  
Cheriton School of Computer Science  
University of Waterloo  
Waterloo, Ontario, Canada  
akhtsang@uwaterloo.ca

Kate Larson  
Cheriton School of Computer Science  
University of Waterloo  
Waterloo, Ontario, Canada  
klarson@uwaterloo.ca

## ABSTRACT

We propose a model where voters are embedded in a social network. Each voter observes the ballots of her neighbors in the network, from which she infers the likely outcome of the election. Each voter may then revise her vote strategically, to maximize her expected utility. Our work focuses on plurality voting, where strategic voting is a major concern. We show that in practice, strategization increases with voter knowledge, yet can improve the social welfare for the population. Real world social networks exhibit a property called homophily; sometimes called “The Echo Chamber Effect”, which is the tendency for friends to have similar ideologies. We find that homophily dampens the benefits of strategization, and correspondingly, lowers the frequency of its occurrence. This effect may contribute to the low number of strategic voters observed in real world elections. Additionally, strategization may lead to the elimination of less popular candidates, as voters revise their votes to less preferred but more hopeful candidates. This phenomenon is known as Duverger’s Law in political science, and we show that it does not hold in certain network structures.

## Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*; J.4 [Computer Applications]: Social and Behavioral Sciences—*Sociology*

## General Terms

Economics, Experimentation

## Keywords

Behavioral game theory, Social choice theory, Social simulation, Emergent behavior, Iterative voting

## 1. INTRODUCTION

The last decade has seen tremendous growth in the popularity of social networks in both popular media and research communities. These networks represent a complex web of interactions be-

tween both individuals and institutions. They capture relationships and social structures that define communities both niche and vast. The relationships within these communities hold the key to how information flows within the network, and ultimately, how individuals’ actions may be influenced by each other and by the institutions whom they respect.

Voting is a method of social choice where a community elicits the personal preferences of individuals to conduct collective decision making. A major concern in voting systems is manipulation via strategic voting. This happens when voters benefit from casting a ballot that does not reflect their true preferences; while this may be beneficial for the voter, it misinforms the community on the needs of its constituents. In order for voters to manipulate successfully, they must have some knowledge regarding the outcome of the election. One reasonable model is to view the election as a series of rounds, where voters put forth tentative ballots that may be continually revised; this is called Iterative Voting, which assumes voters have complete information on the ballots of all other voters [17]. In a social network, however, voters are restricted to observing only the actions of their neighbors. Each voter must form a model of the likely outcome of the election based on this incomplete information, and use this model to inform their actions. This assumption may appear unrealistic at first glance. Since, after all, one does not simply make decisions based on a sampling of opinions from Facebook friends. However, our use of the term social network extends beyond relationships in online social media platforms, and also include experts and associates, media outlets, and any other source of opinion and information that may contribute to the decision making process.

Real world social networks exhibit a number of interesting properties that may impact the strategic behavior of its voters, and should be considered in any realistic model. Of particular interest to our voting model is a property called *homophily*: the tendency for people to connect and socialize with those sharing similar characteristics, beliefs and values. This concept dates as far back as Plato, who wrote in *Phaedrus* that “similarity begets friendship”. In their seminal work, McPherson, Smith-Lovin and Cook offer a survey of evidence that adults, in particular, preferentially associate with those of similar political persuasions [15]. This effect is not only limited to individuals. Hargittai, Gallo and Kane examined the link relationships between sites of top conservative and liberal bloggers discussing political issues, and found homophily to be prevalent; i.e. sites were much more likely to discuss and reference each other when they shared political views. Even more importantly, upon examining the context of links between conservative and liberal blogs, they found that fully half of them were embedded with “straw-man” arguments that reinforced the political position of the author by distorting the opposition’s position [12]. This is especially relevant

**Appears in:** *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.  
Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

to our model because voters derive information about the election from their neighbors in the social network. A homophily of opinions can lead to the so called *Echo Chamber Effect*, where a voter is surrounded by associates that share similar beliefs, reinforcing its validity regardless of its merit.

In this paper, we present a behavioral model of voters embedded in a social network. Voting occurs in successive rounds during which voters may alter their ballots. Voters can observe only the ballots of their friends in the social network. Each voter assumes her friends are representative of the wider population, and will vote strategically to maximize her own expected utility. We explore the behavior of this model on a variety of random graph networks, including ones that exhibit homophily. We focus on using plurality as our voting system. While strategization is a major concern in plurality, we find that it improves the social welfare compared to truthful voting. We also find evidence of the Echo Chamber Effect in our data: interestingly, it lowers social welfare by decreasing the amount of strategization that occurs. This may explain the relatively low percentage of voters that strategize in many real world elections (for example, in [3] and [11]). Finally, while our model converges quickly in practice, we show a counterexample where voters never converge to a stable state.

## 2. MODEL

Let  $V = \{1, 2, \dots, n\}$  represent our set of voters. They are embedded in a social network, represented as a simple, directed graph  $G = (V, E)$ . We adopt the convention that a directed edge  $(i, j) \in E$  denotes that voter  $i$  observes voter  $j$  and as such,  $j$ 's actions may influence  $i$ . An edge may represent communication between friends, a leader's influence on followers, or patronage of media and news platforms. Let  $\mathcal{N}(i)$  denote the set of voters observed by  $i$ ; i.e.  $\mathcal{N}(i)$  is the out-neighbors of  $i$ .

Let  $C = \{c_1, c_2, \dots, c_m\}$  represent the set of available candidates. Let  $\mathcal{F}$  be the voting function used to aggregate those ballots to choose a single winner; it may or may not be deterministic. The choice of  $\mathcal{F}$  will define a set of valid ballots that can be submitted by voters; let us denote this set as  $\mathcal{B}$ .

The voting process proceeds in rounds. In round  $t$ , each voter  $i \in V$  submits a ballot  $b_i^{(t)} \in \mathcal{B}$ . The voter formulates this ballot as a response  $R_i : \mathcal{B}^{|\mathcal{N}(i)|} \rightarrow \mathcal{B}$  based on her observations of her friends – i.e. the previous ballots of her out-neighbors. These rounds may represent a series of preliminary polls leading up to the final election. We assume all voters begin with the truthful ballot.<sup>1</sup> Voting continues until no agents choose to revise their ballots, whereupon the winner is decided by the voting function  $\mathcal{F}$ . When no voters wish to deviate from their current ballot, the system has converged to an equilibrium. If it reaches this state, we say the system is stable.

### 2.1 Model of Voters

Models of voters in multiagent systems literature are divided between those utilizing ordinal preferences (where only the ordering of outcomes matter) and cardinal preferences (where outcomes are associated with utility values).<sup>2</sup> While each model has its own merits, we choose the latter model because our voters infer and weight the probabilities of the different outcomes, and act rationally to maximize expected utility.

Formally, voters derive utility based on the candidate that is elected by  $\mathcal{F}$ . Each candidate  $c_i \in C$  advocates a position  $p(c_i)$  in some

<sup>1</sup>Or a truthful ballot, depending on the voting system

<sup>2</sup>Cardinal utility models are used commonly in the literature, for example in Random Utility Theory [1].

domain  $\mathcal{D}$  that is common knowledge. Each voter  $i$  favors a position  $p_i \in \mathcal{D}$  known only to herself. If  $\hat{c}$  is the winning candidate elected by  $\mathcal{F}$ , then a utility function  $u_i(p_i, p(\hat{c})) : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$  determines the value of this outcome.

For the purpose of this paper,  $\mathcal{D}$  are the integers from 0 to 100 (inclusive), and preferences are single-peaked. This allows us to benchmark our result to previous work (e.g. [7, 8]); this one dimensional scale is also commonly used in political science literature to represent the left-right political spectrum [2, 13]. We assume the utility a voter derives from the outcome decreases with the square of the distance between her favored position  $p_i$  and the winner's advocated position  $\hat{p}$ :

$$u_i(p_i, \hat{p}) = -|p_i - \hat{p}|^2.$$

For brevity, we write  $u_i$  to imply  $u_i(p_i, \hat{p})$  where the position of the candidate  $c_i$  and the position favored by the agent is clear from the context. Throughout this paper, we will refer to the social welfare of the elected outcome. If  $\hat{p}$  is the position of the elected candidate, the social welfare  $SW(V)$  is the sum of the utilities for all voters for that outcome:

$$SW(V) = \sum_i u_i(p_i, \hat{p}).$$

### 2.2 Response Model

Each voter assumes her friends are representative of the wider population. If a ballot  $b$  is observed in a fraction  $f$  of her friends, then she assumes any voter within the network will submit ballot  $b$  with probability  $f$ .

We formally specify the response model for plurality voting for simplicity, but it can be adapted to any voting system with finite  $|\mathcal{B}|$ . This means each ballot is an individual candidate, and  $\mathcal{B} = C$ . Let  $(s_1, s_2, \dots, s_m)$  represent the number of voters in  $\mathcal{N}(i)$  voting for candidates  $(c_1, c_2, \dots, c_m)$ . Voter  $i$  will then assume each voter (other than herself) in the network will support candidate  $c_x$  with probability  $\frac{s_x+1}{S}$ , where  $S$  is a normalizing constant to make the probabilities sum to 1. The +1 is a Laplace smoothing, and is necessary to ensure that all ballots remain possible. This means the ballots from the rest of the electorate follow a multinomial distribution with support  $\mathbf{s} = (\frac{s_1+1}{S}, \frac{s_2+1}{S}, \dots, \frac{s_m+1}{S})$ ,  $S = |\mathcal{N}(i)| + m$ .

We can calculate the probability of any outcome of the election by using the multinomial distribution. Let the vector  $\mathbf{b} = (b_1, b_2, \dots, b_m)$  denote the outcome where the remaining  $n-1$  voters in the network contribute  $b_i$  ballots supporting candidate  $c_i$ . The probability of this outcome is calculated as follows:

$$Pr(\mathbf{b}; n-1; \mathbf{s}) = \frac{(n-1)!}{b_1! b_2! \dots b_m!} \frac{\prod_{i=1}^m (s_i+1)^{b_i}}{S^{n-1}}$$

With complete information, a rational voter only profits from casting a ballot when it is pivotal. With incomplete information, however, our voter must calculate the probability of each winning tie, and cast a ballot that, in expectation, will break ties to maximize her utility. For simplicity, we assume that winning ties between 3 or more candidates are such remote possibilities that they functionally have probability zero. Then, let  $T(y, x)$  be the probability of a winning tie between candidates  $x$  and  $y$ , calculated by enumerating all possible such ties and summing their probabilities. Additionally, we also consider all near-ties, where the addition of one vote to candidate  $x$  will cause a winning tie with  $y$ ; let  $\tilde{T}(y, x)$  be the probability of this outcome.

Finally, voter  $i$  revises her ballot to support the candidate  $x$  with the maximal marginal gain in expected utility  $\mathcal{C}_x$ , calculated below. If a voter observes no other ballots (i.e.  $\mathcal{N}(i) = \emptyset$ ), her ballot remains fixed.

We consider two tie-breaking rules: probabilistic and lexicographic tie-breaking. Below is a modification of prospective ratings introduced by Myerson and Weber [18], for unbiased probabilistic tie-breaking and risk-neutral voters:

$$\mathcal{C}_x = \sum_{y=1}^m \left( \frac{1}{2} T(y, x)(u_x - u_y) + \frac{1}{2} \tilde{T}(y, x)(u_x - u_y) \right)$$

An analogous modification exists for lexicographic tie-breaking, where  $\mathbf{1}_{x < y}$  is an indicator variable with  $\mathbf{1}_{x < y} = 1$  when  $x$  lexicographically precedes  $y$ , and 0 otherwise:

$$\mathcal{C}_x = \sum_{y=1}^m \left( \mathbf{1}_{x > y} T(y, x)(u_x - u_y) + \mathbf{1}_{x < y} \tilde{T}(y, x)(u_x - u_y) \right)$$

### 2.3 Sequential vs Simultaneous Updates

We consider two methods for scheduling when opinion updates take place: sequential and simultaneous. In sequential updates, voters are updated one at a time in a fixed order in each round, and they observe the most up-to-date ballots of their neighbors (which may be updated earlier in the current round, or in the previous round). In contrast, in simultaneous updates, all voters respond simultaneously to observed ballots from the previous round.

### 2.4 Graph Models

We will study the behavior of strategic voters within randomly generated networks. Two important structural characteristics of real world social networks are that they are *small-world* and *scale-free*.

In small-world networks, the average distance between any two vertices in the graph grows as a logarithm of the number of vertices. We expect information to travel quickly through small-world networks, which may have an effect on the aggregate strategic behavior of the population.

Real world networks are often scale-free, which means they are comprised of a handful of highly-connected hubs and many sparsely connected vertices. Highly-connected hubs may represent popular public figures or mass media outlets. In strategic voting, they may wield considerable influence within the network.

We consider 4 graph models in our paper: the Erdős-Renyi (ER) and the Barabási-Albert (BA) random graph models, as well as modifications of these models to incorporate homophily.

Erdős-Renyi is a random graph model that incorporates minimal assumptions. Given density parameter  $pr$ , a directed edge connects any vertices  $i$  and  $j$  with probability  $pr$ . Edge  $(i, j)$  is added with probability independent of the addition of  $(j, i)$ . Erdős-Renyi random graphs are small-world, but not scale-free.

We modify the Erdős-Renyi model to incorporate homophily (hER) by multiplying the probability of adding edge  $(i, j)$  by the homophily factor  $h = 1 - |p_i - p_j|/100$ . Two voters having the same private position have the largest probability of being connected, while voters having diverging positions are decreasingly likely to be connected. Note that the edge density of the resulting graph is decreased as a result of this change.

Barabási-Albert is a preferential attachment model that generates scale-free networks. These networks have many properties similar to human generated social networks. Given attachment parameter  $d$ , each new vertex is added to the graph connected to  $d$  existing vertices. These vertices are selected randomly, with probability proportional to the out-degree of the vertex. In this model, when a new vertex  $i$  is connected to  $j$ , we add *both* the edges  $(i, j)$  and  $(j, i)$  to  $E$ . This ensures information has the opportunity to flow

throughout the network. Barabási-Albert random graphs are both small-world and scale-free.

We incorporate homophily into the Barabási-Albert model (hBA) by multiplying the likelihood of an existing vertex by the same homophily factor  $h$  described above. Note that the edge density of the resulting graph is *unchanged*.

Figure 1 shows an undirected example from each (non-homophilic) random graph model. Both graphs have 40 vertices and are parameterized so that each node has average degree 3.




Figure 1: Example of an ER random graph (top) and a BA random graph (bottom).

## 3. EXPERIMENTAL DESIGN

Our investigation will focus only on the plurality voting rule. We first investigate the effects of the two tie-breaking schemes and update methods. As with Clough's investigation [7], we initialize a population of 169 voters in the baseline graph models: ER and BA. For tractability, we limit ourselves to 3- and 4-candidate scenarios. The positions of candidates and voters are drawn independently, uniformly at random from the interval  $[0,100]$ . The parameters of the graph models are chosen so that the resulting conditions have average out-degree approximately 8, 12, 16, 20, 24, and 28.

In our second set of experiments, we investigate the effects of graph structure and homophily on the behavior of voters and the social welfare of the selected outcome. We focus the experiment on sequential updates and lexicographic tie-breaking, but extend the conditions to include all four graph models. Once again, parameters are chosen to produce the same set of average out-degrees. Note that the density parameter  $pr$  of hER graphs must be doubled to ensure sufficient edge density.

The simulation is written in the D programming language, and compiled using using DMD32 D Compiler v2.067.1 on a 64-bit Windows 7 machine. We limit each election to a maximum of 25

Update/Tie	% Strat	Updates	Avg PoH	Avg PoS
<b>Erdős-Renyi Random Graph</b>				
seq / lex	0.268	57.9	1.2312	1.0795
seq / prob	0.267	57.2	1.2287	1.0776
sim / lex	0.277	78.6	1.2523	1.0689
<b>Barabási-Albert Random Graph</b>				
seq / lex	0.243	44.9	1.2128	1.0866
seq / prob	0.235	43.3	1.2133	1.1002
sim / lex	0.252	68.6	1.2149	1.0885

**Table 1: Effects of update and tie-breaking methods (ER and BA graphs with  $m=4$ ). The metrics measured are the percentage of agents casting strategic ballots, the number of updates before convergence, the Price of Honesty and the Price of Stability.**

rounds, though this limit is never reached. Each data point in the first set of experiments is the average of 400 replications; each data point in the second set is the average of 800 replications.

## 4. RESULTS

We define several metrics measured across our experiments. The **Price of Honesty** (PoH) is defined as the ratio of social welfare of the truthful outcome to that of the strategic outcome.<sup>3</sup> Since both utility values are negative, the larger the PoH, the more costly the truthful outcome is, relative to the strategic outcome. Likewise, we define the **Price of Stability** (PoS) is the ratio of social welfare of the strategic outcome to that of the optimal outcome.<sup>4</sup> We also measure the percentage of voters that engage in strategic play – i.e. the fraction of voters who converge to a ballot that is not truthful – as well as the average number of updates required to reach stability.

Table 1 summarizes these four metrics measured on ER and BA graphs ( $m = 4$ ). Within each graph type, there is little change in the amount of strategization, PoH nor PoS across the three conditions. Despite reaching a similar amount of strategization, simultaneous updates requires a larger number of updates to reach stability. By comparison, the differences between strategization, PoH and PoS is much larger between the two graph types. The same pattern appears in each of the other conditions. We conclude that neither the update methods nor tie-breaking mechanism has a significant impact on the behavior of the voters or the result of the voting process.

Next, we move to the second series of experiments, and the central findings of the paper. We compare the four aforementioned metrics across the four graph models. Strategization is a major concern in elections using the plurality system. However, we show in our experiments that it actually improves the overall social welfare of the elected outcome. Throughout our experiment ( $> 4800$  total trials), we found consistently that the average PoH for each condition is greater than 1; that is, in expectation, the candidate selected by strategic voting achieves a higher social welfare than that selected by truthful voting.

As one might expect, the amount of strategic play increases as voters gain access to more information as connectivity increases

<sup>3</sup>There are various names given for this metric: for example, “improvement in social welfare over truthful” in [16], and “dynamic price of anarchy” in [4]

<sup>4</sup>Since the voter response is deterministic, we may view the outcome of the strategic voting process as unique, and this definition parallels the usual definition of Price of Stability or Price of Anarchy. If viewed as an online algorithm, this measure is analogous to the competitive ratio.

(see Figure 2). However, this gain is asymptotic and the ceiling of strategic play is reached relatively quickly. Interestingly, the ceiling is lower in graphs with homophily than those without. The rate at which strategic play increases (with edge density) is dependent on the graph type, with ER graphs reaching saturation more quickly than BA graphs.

Figure 3 shows the Price of Honesty and the Price of Stability under the different graph models.<sup>5</sup> We include only  $m = 4$  plots, but the same qualitative trends occur for  $m = 3$ . Here we see a possible explanation for the lower strategic ceiling observed in homophilic graphs: it is simply less profitable. The PoH is consistently lower than PoS in these graphs, though they begin to converge at higher edge densities. That is, in these graphs, the social welfare of the strategic outcome is closer to that of the truthful outcome than the optimal outcome.

As strategization occurs in plurality elections, voters begin to abandon less promising candidates for the likely winners, even if they are less preferable. The net result of this behavior is that a multi-party system using the plurality rule will eventually devolve into a race between the two front running candidates. This tendency of plurality favoring 2-party systems is observable in electoral systems around the world, and is known in political science as Duverger’s Law [10].<sup>6</sup>


The consistency of Duverger’s Law is measured by the SF Ratio: the ratio of support for the third and second place candidates [9].<sup>7</sup> Complete agreement with Duverger’s Law would mean no voters will “waste” their votes on lower ranking candidates, and will only cast their ballots in favor of the two leading candidates. This would be reflected by an SR Ratio of 0. Figure 4 shows the distribution of SF Ratios under different graph models, at the condition with the lowest edge density conditions (average out-degree 8). Duverger’s Law would predict that the distribution of SF Ratios be concentrated as a sharp peak near 0. In both 3- and 4- candidate elections, there is little agreement to Duverger’s Law in most graphs, with fewer than 50% of the instances exhibiting an SR Ratio of less than 0.1 (i.e. the third place candidate enjoy less than 10% of the support of the second place candidate). If hER graphs are excluded, at least 50% account for those instances with SF Ratio of at least 0.2. It is interesting to note that in both 3- and 4- candidate elections, hER graphs stand out as showing the most agreement to Duverger’s Law. Notably, the dominant feature of these graphs is homophily, suggesting it helps voters enact Duverger’s Law, even when little information is available to an individual voter.

Figure 5 is a histogram showing the distributions of SF Ratios for ER and hER models, for the three lowest connectivity settings. The bars in blue represents the same data as presented in Figure 4, which is gathered at the lowest connectivity setting (with average out-degree 8). The orange bars shows the distribution of SF Ratios in graphs with average out-degree 12. Here, it is clear that the distribution peaks at 0, and Duverger’s Law is rapidly being restored due to an increase of information available to individual voters. In approximately 65% of the hER instances, the SF Ratio is below 0.1; in the ER graphs, the percentage increases to 80%. The trend continues as we increase the connectivity, as shown in the average out-degree 16 condition (shown as gray bars).

<sup>5</sup>Mann-Whitney  $U < 303,000$ ,  $n_1 = n_2 = 800$ ,  $P < 0.01$ , one-tailed, for all conditions in Figure 3, with two exceptions: ER (avg out-degree 8), and BA (avg out-degree 12). We obtain similar results of statistical significance on  $m = 3$  conditions.

<sup>6</sup>Canada and India are notable exceptions to this rule.

<sup>7</sup>The term SF Ratio refers to the second and first runner-up candidates.



**Figure 2: Fraction of agents strategizing (3- and 4-candidates). Note the different scales in the vertical axis.**

## 5. CONVERGENCE

In our empirical simulation, all trials converge to stability, and do so quickly. It is natural to ask whether the response model is guaranteed to reach an equilibrium in either the sequential or the simultaneous settings. Figure 6 sketches an undirected social network with preferences such that the voter responses result in a cycling of ballots. In this network, there are three candidates, denoted  $A$ ,  $B$ , and  $C$ . The vertices of the graph are divided into four groups, labelled  $V_1$ ,  $V_2$ ,  $A$  and  $B$ .  $A$  and  $B$  are cliques on  $n'$  vertices; all voters in  $A$  have candidate  $A$  as their top preference, and correspondingly with  $B$ , for candidate  $B$ .  $V_1$  contains  $n'$  vertices; each has preference  $A \succ B \succ C$ , and is connected to every vertex in  $B$  and  $V_2$ , but not to each other. Similarly,  $V_2$  contains  $n'$  vertices; each has preference  $C \succ A \succ B$ , and is connected to every vertex in  $A$  and  $V_1$ , but not to each other.  $n'$  may be some large number, such as 10.

It is easy to see that there exist positions for the candidates such that none of the vertices in  $A$  or  $B$  will change their ballots. Each sees strong support for her favorite candidate, which ensures the most likely winning ties will involve that candidate.

Let us consider the sequential update process that updates the vertices of  $V_1$  before  $V_2$ . Each agent votes truthfully in the first round. In the second round, each vertex in  $V_1$  sees  $n'$  supporters for  $B$  and  $C$ , and infers that the outcome will be a likely tie between those two candidates; each vertex switches support to their second-choice  $B$ . Each vertex in  $V_2$  then observes a tie between  $A$  and  $B$ , and also switches to their second-choice:  $A$ . One can then verify that these fickle changes are reversed in the third round, with all voters in  $V_1$  and  $V_2$  reverting back to their truthful choices; thus, the cycle perpetuates.<sup>8</sup>

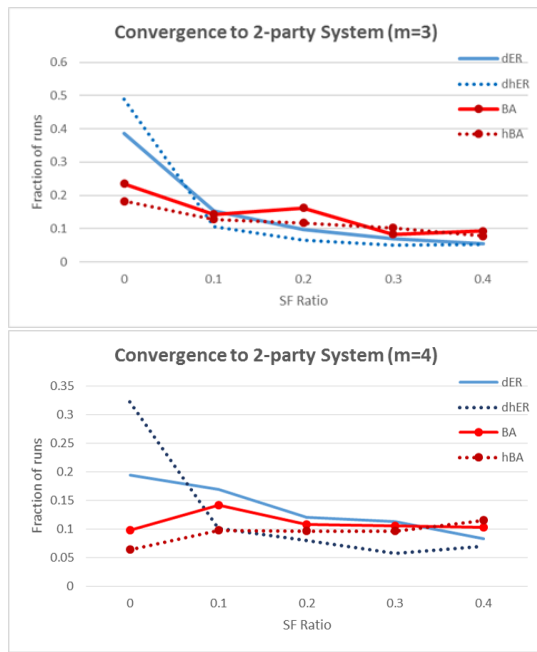
<sup>8</sup>A number of positions for our candidates and voter blocs will produce this behavior. For example, consider candidates  $A$ ,  $B$ , and  $C$  having positions 10, 9, 12 respectively. Let blocs  $B$  and  $V_1$  prefer position 10 (therefore prefers candidates  $A \succ B \succ C$ ), and  $A$  and  $V_2$  prefer position 12 (prefers candidates  $C \succ A \succ B$ ).



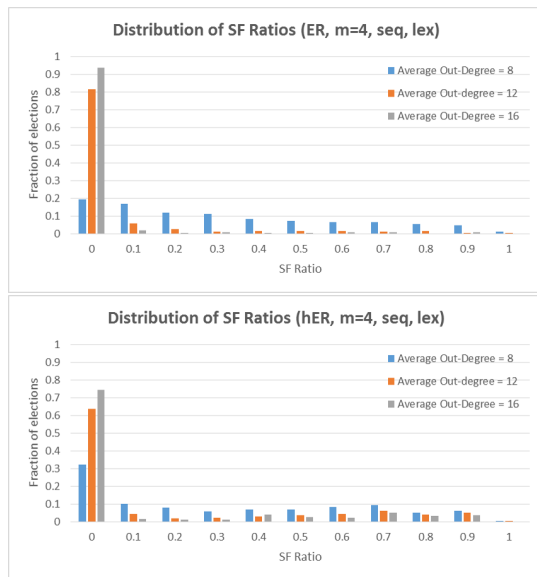
**Figure 3: Price of Honesty and Stability.**

The same counterexample works for the simultaneous update process, with  $V_1$  and  $V_2$  changing in alternate rounds.

Contrast this result with convergence results in the related model of Iterative Voting. By comparison, Iterative Voting occurs in the absence of a social network, where all ballots are common knowledge. Voters iteratively revise their ballots only if it alters the outcome to their benefit. Meir et al. showed that Iterative Voting converges under plurality when voters respond one-at-a-time, but not when they update simultaneously. [17] Lev and Rosenschein demonstrated a similar result for veto, and showed that there is no guarantee of convergence in other scoring rules. [14]




**Figure 4:** Degree of convergence to a 2-candidate system, measured as the SF ratio (3- and 4-candidates). Note the different scales in the vertical axis.



**Figure 5:** Distribution of SF Ratios show the degree with which results from each graph model conform to Duverger's Law.

## 6. DISCUSSION

While we have obtained empirical results for our model, the question remains as to how well it generalizes to real world scenarios. As was alluded to in the Introduction, the social network we depict with our model is a general social network. The neighbors in the network describe not merely "Facebook friends", but include all sources of information that may be considered by a voter in deciding on her ballot. This may include close friends, trusted confidants and knowledgeable associates, but will also news feeds,



**Figure 6:** Voters need not converge to stability.

political blogs, and subscriptions to any number of popular media outlets. Such institutions acts as highly-connected nodes in the social network, much like hubs in Barabási-Albert random graphs. Further, as is shown in Hargittai, Gallo and Kane [12], even such social institutions are not immune to the same homophily exhibited in people.

The successive rounds of voter revision in our model represents the preliminary period preceding an election where voters may discuss and revise their opinions. In the real world, this is often accompanied by a series of preliminary polls leading up to the main election. These polls can be a major factor in strategic voting. Such polls are comprised of (tentative) ballots sampled from a random subset of the population. This is exactly the relationship captured by the (non-homophilic) Erdős-Renyi random graphs, where each voter may view the ballots of a number of other voters sampled uniformly, independently at random from the population.

With homophily being such an intrinsic property of real world networks, it is interesting to note that the graph depicted in Figure 6 shows a very low degree of homophily (for vertices in  $V_1$  and  $V_2$ ). This lack of homophily is necessary for the counterexample to function. Voters that are connected to likeminded voters are less likely to change their votes away from their truthful ballot. They observe many other voters declaring the same ballot, and therefore their favorite candidate is very likely to participate in winning ties. In fact, a careful analysis of the graph structure of Figure 6 reveals what is needed to cause a faithful voter to vote strategically: they view their own position as hopeless, and must be convinced to "pitch in" to resolve a close race between two less-favored candidates. This is in agreement with observations of political elections, such as the empirical study conducted in Cain [5].

This observation may give insights as to why there is less strategic voting in the presence of homophily, and also why the strategic outcome is (comparatively) less profitable. When voters are surrounded with those of similar opinions, it creates an Echo Chamber Effect where they view their own position as being more widely supported than it is. It causes them to be further entrenched in their current position, and they require a larger amount of conflicting evidence to change their minds. The effect causes a voter to have a harder time discerning whether their own position is in the minority, and prevents them from shifting to a more strategic choice. This, as it turns out, has a net negative effect on the social welfare of the elected outcome. Moreover, this effect may explain the relatively small number of strategic voters observed in real world elections (for example, in [3] and [11]): it is not that few voters are strategic, it may be that many voters fail to recognize the strategic opportunity due to their Echo Chamber.

## 7. RELATED WORKS

A cornerstone of our model is the Knowledge Graph model [6]. In their paper, the authors propose a general framework for limiting voting knowledge, restricting each voters' observations to their neighbors in the knowledge graph. However, they do not define any response behavior for individual voters, nor explore the aggregate behavior of the population.

The behavioral nature of our voter model can be attributed to Myerson and Weber [18]. In this paper, the authors propose the concept of a Voting Equilibrium, a refinement of the Nash Equilibrium. Motivated by the process of political elections, they consider that strategic voters may reason on the result of a preliminary poll by considering the tie-probabilities between the various candidates. Then, if strategic voting produces results consistent with the original poll, we are said to be at a Voting Equilibrium.

Clough presented an early political science exploration of strategic voting in social networks [7, 8]. However, the network she uses is a simple  $13 \times 13$  grid-based undirected graph on 169 nodes, which is neither small-world nor scale-free. Each voter responds by considering only tie-probabilities, while our model considers all pivotal cases under different tie-breaking rules. Her work focuses solely on investigating Duverger's Law. Her finds parallel ours: SF Ratios drop dramatically when going from 28 to 8 neighbors. Unfortunately, her model does not offer any finer levels of granularity for investigating this behavior.

A more recent line of inquiry inspired by Myerson and Weber is Iterative Voting [17]. As has been mentioned several times already, Iterative Voting proceeds in rounds. In each round, voters best respond to the previous ballots, either simultaneously or sequentially. Voters have complete information on all ballots, and revise their ballots only when it will change the outcome. Unlike our model, Iterative Voting is guaranteed to converge from a truthful state under plurality and veto [17, 14]. Whereas the update method and the tie-breaking rule were important to their results, we find that our model is robust against changes in these criteria. Brânzei et al. have also investigated social welfare of iterative voting under different voting systems (Plurality, Veto and Borda) [4]. They define the Dynamic Price of Anarchy (DPoA) to be the worst case ratio between the social welfare of the winner elected under truthful voting versus strategic voting. This is similar to our definition for Price of Honesty. Since their model does not operate under a social network, they are able to compute analytical bounds for DPoA under different voting rules. Similar to us, they show that strategic voting improves the elected outcome under Plurality.

Iterative voting has been applied to social networks only very recently in Sina et al. [20], which focuses on manipulation by a chair under plurality voting. Our model differs from Sina et al.'s in that our voters individually infer the likely outcomes of the election based on their limited information, and always act upon this information (to maximize their expected utility based on tie-probabilities). By contrast, in the Iterative Voting model applied by Sina et al., agents only choose to revise their vote when they observe an exact pivotal condition in their neighborhood.

Reijngoud and Endriss have also modeled how voters might respond to information from a series of polls [19]. In their paper, they consider different mechanisms (poll information functions, or PIFs) for summarizing the information present in the current ballot profile. They analyze the susceptibility and immunity to manipulation of different voting rules and PIFs. They also propose models for strategic agents with ordinal preferences, and analyze the performance of different voting rules in the presence of these agents.

## 8. CONCLUSION AND FUTURE WORK

In this paper, we proposed a model of strategic voting on social networks, based on a natural assumption on the part of voters that their friends are representative of the population. We show that strategization leads to improved social welfare of the elected outcome in all conditions. Network structure has an effect on the social welfare of the elected outcome. However, as edge density increases, the amount of information available to each voter also increases, and the number of strategic voters quickly saturate at a ceiling. The ceiling is independent of graph structure, but highly dependent on homophily.

It is this network homophily that causes the Echo Chamber Effect. This may offer insight on why a relatively low number of voters are strategic in real world elections. When surrounded by others with similar opinions, voters do not see an opportunity or even a need to strategize, even when their position holds little merit. This ends up hurting social welfare of the elected result.

As Figure 6 demonstrates, our model is not guaranteed to converge to stability. However, stability is reached relatively quickly in practice. In our simulations, no instances used more than 10 rounds to reach stability. It is unclear why this is the case, and may be a direction for future work. Are such cyclic instances rare? Under what conditions can we guarantee stability? Are such conditions natural to human networks?

Another natural question to ask is, how susceptible to manipulation are voters on a social network? Will voter strategization hinder or amplify the effects of manipulations? If candidates have knowledge of the social network, what strategies may they take to improve their own odds?

Finally, it would be interesting to extend this framework to other, more interesting voting systems. Duverger's Law applies only to plurality, so we expect to see less convergence to 2-party systems when using other voting rules. What effects will this have on strategization and social welfare? Tie-probability modeling for other systems remains an exciting open question.

## REFERENCES

- [1] H. Azari, D. Parks, and L. Xia. Random utility theory for social choice. In *Advances in Neural Information Processing Systems*, pages 126–134, 2012.
- [2] D. Black. On the rationale of group decision-making. *The Journal of Political Economy*, pages 23–34, 1948.
- [3] A. Blais. Why is there so little strategic voting in Canadian plurality rule elections? *Political Studies*, 50(3):445–454, 2002.
- [4] S. Brânzei, I. Caragiannis, J. Morgenstern, and A. Procaccia. How bad is selfish voting? In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [5] B. E. Cain. Strategic voting in Britain. *American Journal of Political Science*, pages 639–655, 1978.
- [6] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *Logics in Artificial Intelligence*, pages 18–30. Springer, 2004.
- [7] E. Clough. Strategic voting under conditions of uncertainty: A re-evaluation of Duverger's law. *British Journal of Political Science*, 37(02):313–332, 2007.
- [8] E. Clough. Talking locally and voting globally Duverger's law and homogeneous discussion networks. *Political Research Quarterly*, 60(3):531–540, 2007.
- [9] G. W. Cox. *Making votes count: Strategic coordination in the world's electoral systems*, volume 7. Cambridge Univ Press, 1997.

- [10] M. Duverger. Factors in a two-party and multiparty system. *Party Politics and Pressure Groups*, pages 23–32, 1972.
- [11] D. S. Felsenthal and A. Brichta. Sincere and strategic voters: An Israeli study. *Political Behavior*, 7(4):311–324, 1985.
- [12] E. Hargittai, J. Gallo, and M. Kane. Cross-ideological discussions among conservative and liberal bloggers. *Public Choice*, 134(1-2):67–86, 2008.
- [13] M. J. Hinich and M. C. Munger. *Analytical politics*. Cambridge University Press, 1997.
- [14] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 611–618. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [15] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, pages 415–444, 2001.
- [16] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 313–330. ACM, 2014.
- [17] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *Proc. of 24th Conference on Artificial Intelligence (AAAI-10)*, pages 823–828, 2010.
- [18] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *American Political Science Review*, 87(01):102–114, 1993.
- [19] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 635–644. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [20] S. Sina, N. Hazon, A. Hassidim, and S. Kraus. Adapting the social network to affect elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 705–713. International Foundation for Autonomous Agents and Multiagent Systems, 2015.