Approval in the Echo Chamber

Ben Armstrong and Kate Larson
Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
{ben.armstrong, kate.larson}@uwaterloo.ca

ABSTRACT

Recently there has been interest in iterative voting, where voters are able to update their votes based on voting information from previous rounds. In this paper we conduct a series of empirical studies in order to understand the strategic issues which arise when agents, voting to approve a set of k candidates, can base their voting or approval decisions on information from their neighbours in a social network. We illustrate that the k-approval voting rule often results in cyclic voting behaviour, that social network structure matters in terms of strategization, and that homophily in the network decreases strategization for the k-approval voting rule.

Keywords

Social choice, Social simulation, Emergent behaviour, Strategic Voting, Iterative Voting, Approval Voting, Homophily

1. INTRODUCTION

Major elections in recent years have seen uncommonly high levels of divisiveness across the electorate. This tendency to associate with only those considered similar to oneself is called homophily and can cause individuals to be surrounded primarily by others with similar opinions, leading all groups to be convinced that they have the majority of support. A possible cause for the recent levels of homophily exhibited in the world is social networks such as Facebook or Twitter [5].

The recent rapid growth in prevalence of real-world social networks makes them excellent sources of information for understanding effects such as homophily. Social networks as a general tool are used to model human interactions by representing personal information about existence and strength of relationships, and the distance between two people while also allowing insight into societal trends.

An active area of research studies how social networks affect elections by modeling networks of voters with the connections between voters representing the relationships the voters have. The political opinions of a voter's neighbours can have an impact on the voter's decisions throughout the election, also referred to as the "strategy" of the voter.

Appears at: 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017). Held as part of the Workshops at the 16th International Conference on Autonomous Agents and Multiagent Systems. May 8th-9th, 2017. Sao Paulo, Brazil.

Strategic voting occurs when a voter submits a ballot that is not entirely honest. In plurality voting systems this most often manifests as voting for a candidate that you prefer over the candidate you believe is going to win even though neither of those is your most preferred candidate.

In addition to plurality there are voting systems such as approval voting. In approval voting, the voters may approve of as many candidates as they wish; in some cases strategization might involve approving the 2 or 3 most preferred candidates. Approval voting has many advantages over plurality voting and other voting rules. It allows a voter to safely vote for their favourite candidate while also allowing strategization. Approval voting has been used in a variety of situations including papal elections for over 300 years and by the American Mathematical Society. Many voting rules are quite complex while approval voting is praised for its simplicity [15]. In this paper we focus on k-approval voting, in which voters must approve of exactly k candidates.

This paper studies the intersection of homophily, social networks, strategic voting, and approval voting. Through simulations we explore the effect that homophily has on a variety of differently structured social networks in which the voters employ strategic k-approval voting. Our findings suggest that homophily leads to a much lower social welfare, having a larger number of candidates leads to a better outcome, and that strategization increases with the number of candidates.

2. RELATED WORK

This work builds upon work done by Tsang and Larson [16] which uses a similar model and focuses on the plurality scoring rule, which is equivalent to $1 - approval$. It was shown that under the plurality rule strategization leads to an improvement in social welfare over the truthful outcome, and that the presence of homophily decreased the occurrence of strategization. Overall, similar results are shown in this paper.

The model used in this paper has been inspired by previous work, particularly that of Chopra et al.[3] which introduces a *knowledge graph* containing voters and an edge (i,j) when voter i is able to observe the current preference of voter j. The difference is that Chopra et al. are focused purely on the strategic behaviour of each voter rather than the behaviour of the entire system.

The model of voter decision-making is based on the work done in [10]. The authors study voting equilibria which occur when all voters in a population cause an outcome they have no incentive to switch away from. It is shown that approval voting in a three candidate election leads to a winner located in the median of the voter positions. Certainly, in our model when voters must approve 2 of the 3 candidates one of their honest approvals will be for the median candidate, making that candidate almost certainly the winner.

Clough studies Duverger's Law using a model of iterative voting very similar to ours in [4]. In her model voters are on a grid network, which is neither small-world or scale-free, and use information from their neighbours in order to participate in iterative plurality voting. The focus in her work is on Duverger's Law which is primarily of interest under plurality voting rules and not studied here though her results do suggest there may be less strategization when voters have less information which is not seen in our results. The primary differences between her work and our own are the network structure, use of plurality voting, and a response function that considers only ties rather than ties and near ties.

Several papers show that iterative voting does not necessarily converge under most common voting rules [9, 7, 12, 13]. In particular, approval voting is shown to have no guarantee of convergence. However, the models used in these papers contained voters with complete information that modified their ballots only when they believed the modification would change the outcome of the election.

On the subject of strategic voting, much work has been done. In particular, Smith provided results from a Bayesian regret analysis of approximately 2.2 million simulations showing approval voting to be better than most common voting rules in the presence of strategic voters [15]. Interestingly, his data also showed that in all common voting rules strategic voting leads to more unhappiness than truthful voting which is in contrast to the findings of our model. The difference between these results could be due to the presence of iterated voting in our model and the lack thereof in Smith's. Slinko and White provide a study of an effect observed in our model [14]: Often, in the first round of updating, a majority of voters will approve of a candidate that has very little honest support. This is done in an attempt to remove support for candidates that are not the voter's favourite candidate but nonetheless have a large amount of support. The result is that often after a single round of updates the least popular candidate has the most approvals. Slinko and White refer to this as "strategic overshooting" and provide results indicating that there is often a "safe" ballot that will be better than an unsafe ballot.

Approval voting has been studied in several contexts, Poundstone performed a Bayesian regret analysis on a variety of voting methods and concluded that unrestricted approval voting was much simpler and led to more satisfaction than most other voting methods [11]. Our work restricts the number of approvals voters are permitted.

3. MODEL

In this section we describe the voting problem analysed in this paper. Let a set, V, of n voters be situated in some social network, $G = (V, E)$, where G is a directed graph such that $(i, j) \in E$ means that voter i observes voter j, and thus may be influenced by the voting behavior or opinions of voter j. We define the out-neighbours of i to be the set $\mathcal{N}(i) = \{j|(i, j) \in E\}$, and thus this is the set of voters who may influence voter i.

Let $C = \{1, \ldots, m\}$ be the set of candidates or alternatives over which the voters in V may cast votes, and let each candidate $c_j \in C$ be associated with some *position* $p(c_j) \in [0, 100]$. Furthermore, our voters, V, have single-peaked preferences over the candidate set. Each voter, $i \in V$, has a preferred position p_i, and thus its utility if some candidate is selected with position $\hat{p}$ is

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

Each voter casts a ballot, b, from a set of admissible ballots $\mathcal{B}$. A social choice function, $\mathcal{F} : \mathcal{F} \mapsto \mathcal{P}(C)$ is used to aggregate the voters' ballots and select a subset of candidates as winners. We are interested in situations where voting is iterative and progresses in rounds. In round t, each voter i simultaneously casts a ballot $b_i^{(t)} \in \mathcal{B}$ which is chosen in response to the previous ballot $b_j^{(t-1)}$ of each out-neighbour $j \in \mathcal{N}(i)$. If all voters refrain from updating their ballots in a given round, voting stops (and the system is considered stable). Otherwise, voting continues until r rounds have passed. Each round could be considered as a formal poll, or as a more informal update of decisions by voters that happens naturally over time. After voting stops, the winning set of candidates is decided using the voting function, $\mathcal{F}$.

3.1 Ballot Formation

In this paper we are particularly interested in how voters form their ballots as part of the iterative voting procedure and we make the observation that voter i can base its ballot decisions on the previous ballots of members of $\mathcal{N}(i)$. We argue that each voter believes that their neighbours are representative of the rest of the network, and thus if a fraction f of their neighbours approve of a particular candidate the voter assumes that the same fraction of voters in V approve that candidate and so may strategically cast a vote accordingly.

More formally, in this paper we study the k-approval voting process in order to understand how iterative voting may lead to strategization amongst networked agents. K-approval voting is a member of a larger class of voting rules, scoring rules, in which ballots are vectors representing a score given to each candidate by that voter. A scoring rule uses a vector of the form $(\alpha_1, \alpha_2, ..., \alpha_m)$ where $\alpha_i \geq \alpha_{i+1}$. The candidate ranked first by the voter is given α_1 points, the candidate ranked second given α_2 points, and so on. The candidate with the highest aggregate score wins. In k-approval voting, a ballot has the form $\{\underbrace{1, ..., 1}_{\text{k 1's}}, 0, ..., 0\}$.

If s_j is the total score for candidate c_j in $\mathcal{N}(i) \cup i$ then i believes that the fraction of support in the entire network for c_j is $\frac{s_j+1}{S}$ where $S = \mathcal{N}(i)+m$. We use Laplace smoothing to ensure that all candidates have a non-zero chance of winning and thus the vector $\mathbf{s} = (\frac{s_1+1}{S}, \frac{s_2+1}{S}, ..., \frac{s_m+1}{S})$ represents the level of support for each candidate.

In order to decide upon which candidates to approve, a voter assigns a prospect rating to each candidate x by enumerating all possible ties between candidates x and y, calculating the likelihood of that outcome and multiplying by the utility gained if x wins. Specifically, the likelihood of a particular outcome from the other $n-1$ voters $\mathbf{b} = (b_1, b_2, ..., b_m)$, where b_i is the number of approvals of candidate i, is given by:

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

Let $T(y,x)$ be the chance of a tie between x and y where x and y are both in a winning position, referred to as a *winning tie*, and $\widetilde{T}(x,y)$ be the chance of an outcome where candidate x has one less vote than candidate y and candidate y is winning. Then voters assign each candidate x a prospect rating using lexicographic tie-breaking, given by:

$$C_x = \sum_{y=1}^{m}(\mathbf{1}_{y<x}T(y,x)(u_x - u_y) + \mathbf{1}_{x<y}\widetilde{T}(y,x)(u_x - u_y))$$

$\mathbf{1}_{x<y}$ is 1 when x lexicographically precedes y and 0 otherwise. $(u_x - u_y)$ gives the marginal utility gain from voting for x over voting for y. In each round voters calculate C_x for each candidate and approve the k candidates with the highest values. An alternative approach for ballot decision-making is discussed in Section 7.

3.2 Network Structure and Properties

We will be interested in understanding how network structure and properties influence strategic choices of voters in the context of k-approval iterative voting. We study two different types of random network structure in this paper: Erdös-Renyi (ER) and Barabási-Albert. These graph structures are used as they have one important property in common, both being *small-world* graphs, and differ in another property, being *scale-free*. Further, we study these network both with and without the presence of homophily, the tendency of similar individuals to associate with one another more often than those with differing views.

Graphs are *small-world* if the average distance between any two nodes in the graph grows proportionally to the logarithm of the number of nodes in the graph [6]. This results in nodes typically being connected to other nodes by a very short path. *Scale-free* graphs have a degree distribution which follows a power law, resulting in many nodes with many edges fewer than average and many nodes with a greater than average number of edges [6]. Many real world networks exhibit small-world [17] and scale-free [1] properties.

Erdös-Renyi (ER) graphs are generated by a parameter, pr, representing the probability of attachment. Any two nodes i, j are connected by a directed edge from i to j with some provided probability pr. ER graphs are small-world but not scale-free. When studying ER graphs with homophily, we multiply pr by a *homophily factor*

$$h = 1 - \frac{|p_i - p_j|}{100}$$

in order to increase the probability of a voter being connected to similar voters. Adding homophily has the effect of reducing the edge density by approximately $\frac{2}{3}$ for the same value of pr.

Barabási-Albert (BA) graphs use preferential attachment to generate a larger variance in average degree. An attachment parameter d is decided upon, the graph begins with d vertices, all connected to each other with edges in both directions. The remaining $n-d$ vertices are added one at a time, attaching each one to d existing vertices, selected randomly with probability proportional to the degree of each existing vertex. When an edge is added from i to j, the edge from j to i is also added. These graphs are small-world and scale-free. Homophily can be incorporated by multiplying the probability of adding any connection by homophily factor h. This modification has no effect on the edge density of the network.

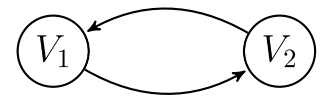

Figure 1: A very basic example of a social network. V_1 is influenced by V_2 and V_2 is also influenced by V_1.

4. CONVERGENCE OF ITERATIVE K-APPROVAL VOTING

The first property we are interested in is whether iterative k-approval voting converges, that is, whether a state is reached where no voter wishes to update their ballot. While it was known that in a number of situations iterative voting was not guaranteed to converge [7], recent experiments using plurality voting showed that non-convergence was rarely an issue in practice [16]. Unfortunately, as we show in this section and later support with experimental findings, iterative k-approval voting is likely to not converge. We illustrate the problem through a simple example.

Under k-approval, each voter must approve of exactly k candidates and does so based upon a prospect rating assigned to each candidate (discussed in Section 3). Consider the simple network shown in Figure 1 under 2-approval. Let V_1 have preferred position 93 and V_2 have preferred position 24. Assume, furthermore, that there are 3 candidates A, B, and C with positions 0, 43, 35 respectively. We can then induce the following preference orderings over candidates for each voter:

$$V_1 : B \succ C \succ A$$
$$V_2 : C \succ B \succ A$$

If the voters submit truthful ballots on the first iteration, then candidates B and C are each awarded a score of 2, and lexicographic tie-breaking results in B winning. At first glance, this seems like a reasonable outcome - both voters agreed on their ballots, and so one might expect that no updates would occur in further iterations of the voting process. However, that is not the case.

After observing each others' initial ballots (due to the structure of the network in Figure 1), each voter computes prospect ratings, using the equations in Section 3, for each candidate. These prospect ratings are shown in Table 1. Voters choose the k candidates with the highest prospect ratings so V_1 will not change its ballot, however, V_2 will switch to a ballot approving A and C. Thus, in round 2, candidates A and B have one approval, while C has two, resulting in candidate C (V_2's preferred candidate) being declared the winner. Prospect ratings are generated for the candidates after this second round of voting (Table 2), and again resulting in a change in ballots. In particular, voter V_2 would prefer to approve candidates B and C, as it did originally, thus beginning a never-ending cycle.

While at first glance, the cyclic voting behaviour of voter V_2 may seem counter-intuitive, it does have a rational un-

	A	B	C
V_1	-257	313	-56
V_2	-15	-43	58

Table 1: Prospect ratings for each candidate after the election at time $t = 0$.

	A	B	C
V_1	-721	323	398
V_2	-50	-13	63

Table 2: Prospect ratings for each candidate after the election at time $t = 1$.

derpinning. After the first ballot, V_2 observes equal support for candidates B and C and is aware that the tie-breaking rule favours B. Thus, by reducing support for B, and approving A and C, voter V_2 is able to ensure that C is the winning candidate. However, in the second round of voting, there is now support for all three candidates (2 approvals for C, one approval each for A and B. By shifting its approvals from A and C to B and C, V_2 is able to ensure that its least preferred candidate A will certainly not be a winner since its receives no approvals. Thus, it reverts back to its original ballot.

If the above explanation correctly personifies the "thought process" of V_2 it reveals both an interesting emergent strength and weakness of the model. First, V_2 does correctly identify that its preferred candidate will not win, despite having a large amount of support, and changes an approval from a more preferred candidate to a less preferred candidate, a much more intelligent action than could have been expected. Second, the voter does not seem to realize that while reverting to its original ballot will accomplish the immediate goal of removing A as a contender in the election, it will also cause a return to the original situation in which B, the most preferred candidate of V_2 loses. Thus, a more advanced model might look ahead and see what effect a change in ballot might have or look to history to avoid cyclical situations. It may also be useful to consider a weaker definition of convergence, where the system is considered stable after a candidate wins for a particular number of consecutive rounds.

5. EXPERIMENTAL SETUP

We are particularly interested in deepening our understanding of the relationship between strategization and iterative voting under k-approval on a social network. To this end, we conducted a series of experiments, varying different aspects of the underlying social network of voters and the number of candidates.

For all experiments we set $k = 2$ and set the number of iterations to be at most 20. Unless otherwise noted, we set the number of voters to be 150, and varied m from 3 to 5. For $m = 3$, our findings are the average of 200 trials, while for $m = 4$ and 5 our findings are the average over 100 trials.

In the first set of experiments, we studied what happened as we varied the underlying social-network structure. In particular, for each class of graph (ER, homophily+ER (hER), BA, homophily+BA (hBA)) we set the parameters so that voters had an average out-degree of approximately 12, 20, and 28 for $m = 3$ and 4. We measured and report several

m	Homophily?	PoH	M:T	PoS	M:O	% Str
Erdös-Renyi Graph						
3	No	1.234	1.518	1.476	1.809	0.203
	Yes	1.169	1.493	1.438	1.806	0.177
4	No	0.992	1.083	1.074	1.168	0.385
	Yes	1.259	1.609	1.370	1.737	0.407
5	No	0.950	0.954	1.053	1.056	0.397
	Yes	1.073	1.225	1.194	1.350	0.384
Barabási-Albert Graph						
3	No	1.213	1.488	1.460	1.784	0.193
	Yes	1.144	1.511	1.393	1.811	0.145
4	No	1.008	1.137	1.085	1.222	0.377
	Yes	1.195	1.530	1.298	1.651	0.378
5	No	0.942	0.966	1.041	1.067	0.398
	Yes	1.036	1.190	1.166	1.326	0.379

Table 3: Summary of results for experiments with 3-5 candidates comparing, for graphs with and without homophily, the average Price of Honesty, Mean:Truthful ratio, Price of Stability, Mean:Optimal ratio, and proportion of voters engaging in strategic behaviour.

	ER	hER	BA	hBA
$m = 3$	14.66	16.64	14.23	16.38
$m = 4$	11.86	15.27	12.12	15.09

Table 4: Average number of updates per agent, as a function of social-network structure.

metrics including the prevalence of strategization and the effect of connectivity on social welfare.

A second set of experiments was run on a smaller population of 60 voters with $m = 5$ candidates and $k = 2$. Average voter out-degrees were varied over 4, 12, and 20. These experiments begin to provide hints as to how the number of candidates affects the social welfare of the system. However, due to the limited population size and wide variance in average degree relative to population size these simulations are intended as only a starting point for a study on the effects of the number of available candidates.

6. RESULTS

In this section we report our findings. We are interested in understanding the frequency with which voters update their ballots, the amount of strategization that occurs as a function of the underlying social network, and the degree to which strategization is either beneficial or harmful to the system in terms of social welfare. We initially report our findings from experiments with $m = 3$ and 4 candidates, and then provide a short discussion of our preliminary findings with 5 candidates.

6.1 Updating of Ballots

One measure of interest is the frequency in which voters change their votes over a certain period of time. This provides us with insight into both the level of strategization

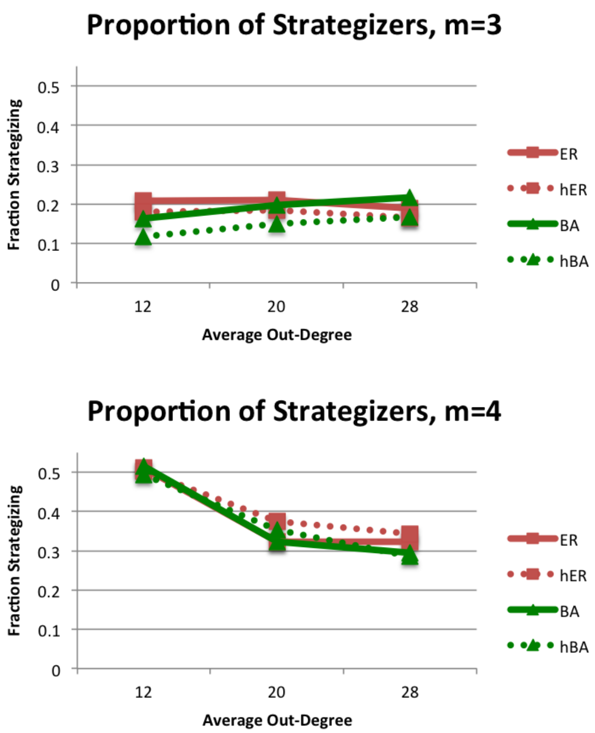

Figure 2: Proportion of voters strategizing for each network type.

occurring in the system, as well as the cognitive overhead required by voters as they decide which ballots to submit in each round. Table 4 reports the average number of ballot updates for each voter over a period of 20 iterations. We make several observations. First, the voting rarely converged, as was discussed earlier in the paper, and so voters were still best off updating their ballots after 20 iterations. Second, the number of candidates seemed to influence the number of updates slightly. With 4 candidates, there were consistently fewer updates across the system. Finally, while graph structure (i.e. ER vs BA) did not seem to be a significant influence, the presence of homophily in the network was important. This was somewhat unexpected, as we had thought that being surrounded by voters with similar views should make a voter more confident in their ballot.

6.2 Degree of Strategization

A voter is considered to be voting strategically if its ballot is anything but entirely honest. Figure 2 shows the effect of homophily on strategization. In each graph with homophily, the fraction of voters strategizing is consistently (albeit, very slightly in the case of $m = 4$) lower than the non-homophily version of that graph.

The fraction of strategizing voters tends to decrease with increasing degree for Erdös-Renyi graphs. In Barabási-Albert graphs, when $m = 4$, that trend continues however it seems as though when $m = 3$ the fraction of strategizers increases with the degree. The reason for this is unclear but it does represent a difference from plurality voting in which strategization always increased with edge density (to a plateau) [16].

Also interesting to note is that of every single strategizing voter, exactly one of their approvals was strategic and the other honest. This is unavoidable when m is 3 but at $m = 4$ agents are capable of approving their two least favourite candidates but seemingly never consider it useful to do so. This is consistent with the idea that approval voting should always allows you to vote for your favourite candidate while also voting strategically for a "lesser of two evils" of candidates more likely to win than your favourite. It has been shown that when voters are allowed to decide the number of candidates they approve it is always useful to approve one's favourite candidate [15].

6.3 Benefits of Strategization

In this section we report on our findings as to how beneficial strategization is for the entire system. We define the *social welfare* of some candidate $\hat{c}$ with position $p(\hat{c})$ being chosen as

$$SW(V) = \sum_{i \in V} u_i(p_i, p(\hat{c}))$$

where p_i is the preferred position of voter i.

We use several other 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 [16, 2, 9]. Since both utility values are negative, the larger the PoH, the more costly the truthful outcome is, relative to the strategic outcome. We also define the Price of Stability (PoS) to be the ratio of social welfare of the strategic outcome to that of the optimal outcome [16]. A smaller PoS shows that strategization is more beneficial than honesty; the lowest possible value occurs at 1 when the strategic outcome is the optimal outcome. A PoH larger than 1 indicates that strategic behaviour is more beneficial to the population while a PoH less than one indicates truth-telling is more beneficial. These, or similar, metrics have been used in many settings for evaluating the performance of a system [2, 8]. Both were used in the original analysis of this model for plurality voting and seem quite appropriate when the system has converged to a stable state.

In our experiments, however, voting rarely converged, limiting the usefulness of PoS and PoH. Thus, we propose two additional variants of these measures better suited for non-converged systems. In place of the Price of Honesty we study the mean social welfare (the mean strategic SW of the winner from each round of an election) divided by the truthful social welfare (Mean:Truthful). Price of Strategy is replaced by the mean SW divided by the optimal SW (Mean:Optimal), pictured in Figure 3 and Figure 4 respectively. These measures provide a more accurate representation of the system, though average values of the PoS and PoH are included in Table 3 to illustrate that they follow the same qualitative trends as our new metrics.

Similarities can be seen between PoH and the Mean:Truthful ratio, and the PoS and Mean:Optimal ratio suggesting that the comparisons are valid. The generally lower values of PoH compared to Mean:Truthful suggest that the final strategic result is not as good as the mean strategic result, or that over time strategies tend to become less beneficial. In general, the opposite trend seems indicated by the comparison of PoS and Mean:Optimal SW which suggests that the final strategic result is closer to optimal than the mean strategic result. This seems somewhat contradictory and warrants closer inspection.

While we see little difference when it comes to whether

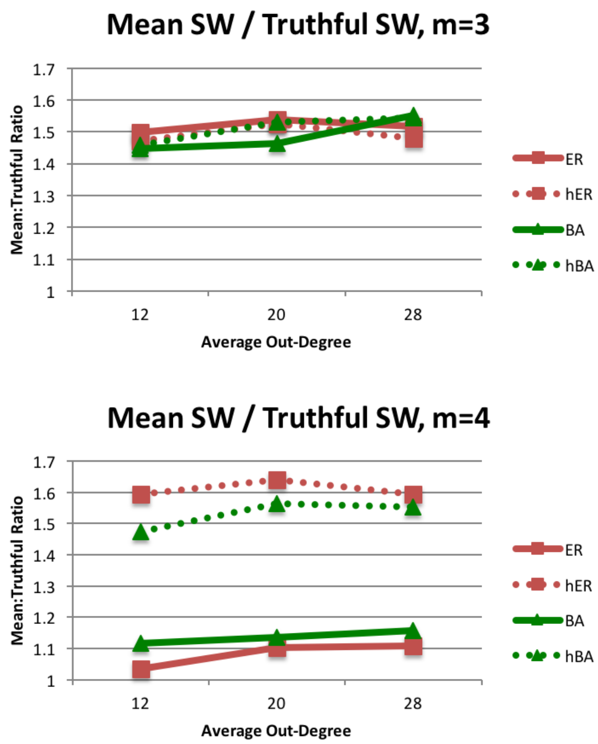

Figure 3: Mean Social Welfare over Truthful Social Welfare for each network type.

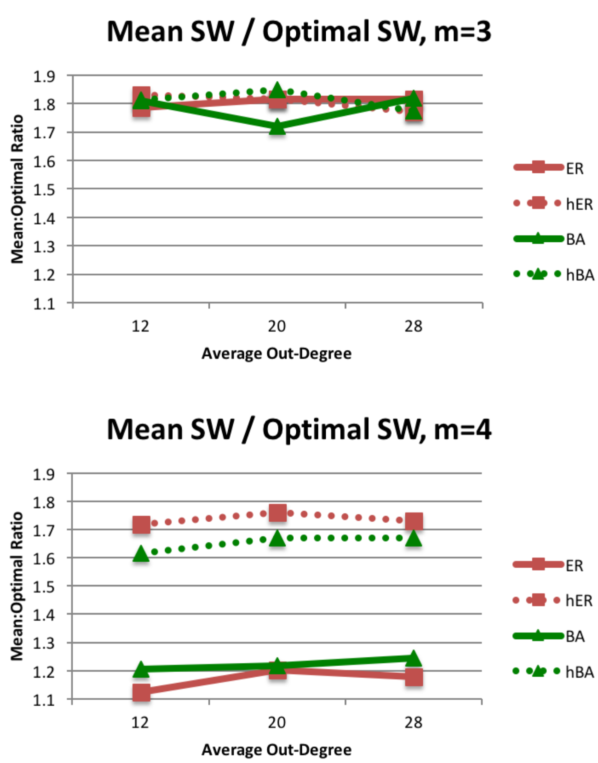

Figure 4: Mean Social Welfare over Optimal Social Welfare for each network type.

the underlying social network was generated using ER or BA, we do note that homophily is important. The difference between graphs with and without homophily can be observed most readily when $m = 4$. We can see from the Mean:Optimal ratio that without homophily voters benefit much more from strategization than those with homophily. We observe a similar effect from the Mean:Truthful ratio: homophily leads to a lower social welfare.

6.4 Simulations with 5 candidates

While we conducted smaller experiments for the case where $m = 5$ we still report our preliminary findings as they raise some interesting questions. Our results can be seem in Figure 5.

First, we note that the Mean:Truthful ratio is now consistently below 1, indicating that the actual (strategic) outcome is always better than the honest outcome and suggesting that as the space for strategization grows it becomes more beneficial. Evidence for this is also found by observing that the Mean:Optimal ratio is closer to 1 than in previous experiments.

We also noticed a considerable difference in the proportion of strategic voters. With 4 candidates, a lower degree led to more strategization while with 3 and 5 candidates degree seemed to have little effect on strategization levels. However, with 5 candidates there is significantly more strategization occuring compared to the 3-candidate case; consistently 35-40% of candidates strategize. Interestingly, with 4 candidates, there is a case where strategization is at approximately 50%, much higher than seen here. This difference may be related to the differing population sizes but is mildly surprising as the opportunities for strategization are much larger with a larger ballot. We also noted that there were instances when $m = 5$ where a voter would not vote for any of their top k-candidates, including one instance where as many as 13 voters in a single round did not approve any of their top k candidates. This observation needs to be investigated further, as it opens up a number of questions with respect to the strategy space of voters.

6.5 Comparison with Plurality

The results found in this paper have both similarities and differences to those found under plurality voting [16]. Average PoS and PoH[1] seem to be quite similar when $m = 4$ (plurality data is not available for $m = 3$) for ER and BA graphs with a slightly higher PoS for hER and hBA graphs. In general, homophily tends to reduce the benefit of strategization, however k-approval seems to be affected more strongly than plurality.

Curiously, the fraction of agents voting strategically is quite different in k-approval. In both $m = 3$ and $m = 4$, the fraction strategizing was higher than in plurality, however (excluding BA and hBA for $m = 3$) the graphs follow a different curve. In plurality, strategization goes up with degree and here the trend is the opposite.

As the plurality simulations consistently converged within several rounds the number of updates is quite a bit lower, averaging 40 to 80 ballot updates per election. By contrast, k-approval averaged over 2000 updates per election. This massive increase is explained by the fact that in plural-

[1] To simplify comparison between the results, we write PoS and PoH rather than Mean:Optimal and Mean:Truthful for this subsection only.

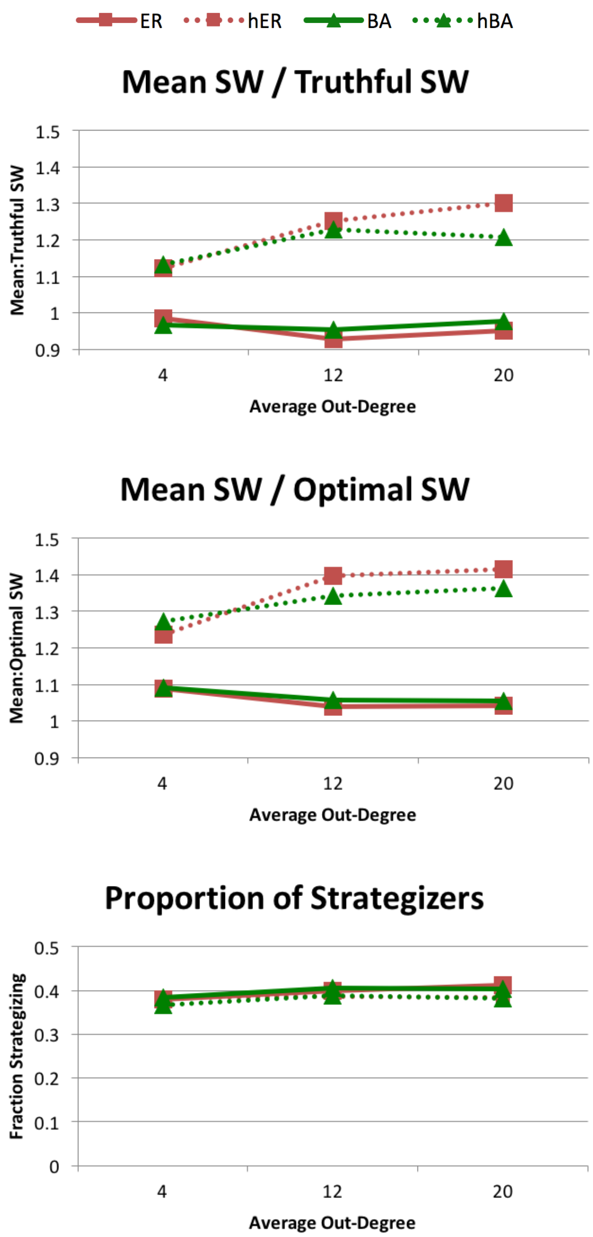

Figure 5: Several metrics showing data for 5 candidates over average voter out-degree of 4, 12, and 20.

ity elections converged within 6 rounds while with approval voting the elections did not converge and were terminated after 20 rounds. In plurality the number of agents strategizing quickly drops to zero while under approval the number of strategizers typically remains constant or cycles between large and small numbers in alternating rounds.

7. CONCLUSIONS

In this paper we have studied the effects of homophily and a variety of social network structures on strategic voting behaviour in a k-approval iterated voting system. Our model uses a population of voters connected in a social network, each with a preference on some issue and several candidates, also with preferences. In each round, the voters observe the support for each candidate amongst their neighbours and determine the expected utility from a vote for each candidate, selecting the k candidates with the top utilities.

We notice several interesting patterns emerge from our data. As with previous work under plurality, strategization is slightly more prevalent and much more useful for networks without homophily, indicating that social welfare is increased when a voter's associates hold a wide variety of opinions. We also noted that as the space for strategization increased (when the number of candidates increased) strategization led to a larger utility gain, as might be expected. Curiously, the fraction of voters strategizing tends to either go down or remain approximately constant as the average degree of the network changes. With plurality voting it was observed that there seems to be a ceiling on the amount of strategization occurring in the network, it is possible that our experiments were simply at this ceiling much of the time.

The results of this study indicate that there is significant room for future work in this area. A more in-depth study of approval voting is warranted, as well many additional voting systems are available for study. In future work, it may be desirable to modify the model in such a way that would make simulations more likely to converge and avoid repetitive behaviour.

The model used in our experiments allows for two different ways of constructing a ballot. Currently, a voter views the ballots of all its neighbours, sums the approvals and calculates the prospects of each individual candidate and selects the k candidates with the highest prospects. This method was chosen due to the intuitive aspect of simply approving of the candidates that give the highest expected utility. The disadvantage in this method is that in considering only specific candidates information about which sets of k candidates are most approved by the voter's neighbours is lost. The alternative method involves calculating the prospects for each individual ballot, rather than each candidate. This would retain some useful information however it would vastly increase the computational complexity of the simulations as, in effect, there would be one candidate for each possible ballot.

A metric more relevant to approval voting might be the notion that a voter can always safely approve of their favourite candidate and also approve of other candidates. This property is always satisfied for $m = 3$ and 4, and only very occasionally not met when $m = 5$. We hypothesize that if voters were allowed a variable number of approvals, this property would always be met.

Presently, the preference structure of the voters in our model is limiting in several ways. The requirement that voters have single-peaked preferences means there will be a tendency to elect the candidate with the median opinion, and in fact when $k > \frac{m}{2}$ that candidate will always have the most honest approvals. Removing single-peaked preferences could be difficult while keeping intact the preference structures we have given voters. A simple modification to the model could give voters and candidates multi-dimensional preferences to reflect the fact that each agent may have a distinct opinion on several issues. This allows for slightly more variance in preferences over candidates while leaving the possibility for a simple utility function. Unfortunately, this would likely not remove all bias towards electing the median candidate but it may reduce the likelihood of such an event.

Finally, extending this work to yet more election methods could yield interesting comparisons between the methods. Different voting methods could serve two purposes: First, running experiments with alternative methods would, of course, teach about the behaviour of voters under those methods and may yield surprises as with the lack of convergence with approval voters. Second, different voting methods might serve to highlight aspects of this model that could be further refined or may not generalize well, and may give clues as to how to construct a more accurate model.

8. ACKNOWLEDGEMENTS

We would like to thank Alan Tsang for allowing us to build upon code written for his paper studying this model with the plurality rule and for his answers to numerous questions during the writing this paper.

REFERENCES

[1] A.-L. Barabási. Scale-free networks: a decade and beyond. *science*, 325(5939):412–413, 2009.

[2] S. Brânzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *AAAI*, volume 13, pages 138–144, 2013.

[3] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *European Workshop on Logics in Artificial Intelligence*, pages 18–30. Springer, 2004.

[4] 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.

[5] Y. Halberstam and B. Knight. Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter. Technical report, National Bureau of Economic Research, 2014.

[6] M. O. Jackson et al. *Social and economic networks*, volume 3. Princeton University Press, Princeton, 2008.

[7] 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.

[8] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 313–330. ACM, 2014.

[9] 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.

[10] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *American Political Science Review*, 87(01):102–114, 1993.

[11] W. Poundstone. *Gaming the vote: Why elections aren't fair (and what we can do about it)*. Macmillan, 2008.

[12] R. Reyhani and M. C. Wilson. Best-reply dynamics for scoring rules. In *20th European Conference on Artificial Intelligence*. IOS Press, 2012.

[13] 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 (AAMAS 2015)*, pages 705–713, 2015.

[14] A. Slinko and S. White. Is it ever safe to vote strategically? *Social Choice and Welfare*, 43(2):403–427, 2014.

[15] W. D. Smith. Range voting. *The paper can be downloaded from the author's homepage at http://www. math. temple. edu/˜ wds/homepage/works. html*, 2000.

[16] A. Tsang and K. Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 368–375. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[17] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *nature*, 393(6684):440–442, 1998.