Predicting Voting Outcomes in Presence of Communities

Jacques Bara
University of Warwick
jack.bara@warwick.ac.uk

Omer Lev
Ben-Gurion University
omerlev@bgu.ac.il

Paolo Turrini
University of Warwick
p.turrini@warwick.ac.uk

ABSTRACT

Individuals in a social network may form their views as a result of the influence exerted by their connections. In elections, for example, while they might initially support one candidate, social influence may lead them to support another. Here, we investigate whether a recently proposed metric, *influence gap*, designed to measure the effect of social influence in voting on social networks, is able to predict the outcome of a vote on networks exhibiting community structure, i.e., made of highly interconnected components, and therefore more resembling of real-world interaction.

To encode communities, we extend the classical model of caveman graphs to a richer graph family that displays levels of *homophily*, i.e., where connections and opinions are highly intertwined. We show that, across these graphs, there are important cases when the influence gap correlation is a weak predictor due to communities, and a simpler metric, counting the initial partisan majority, provides a more accurate prediction overall. Using regression models, we further demonstrate that the influence gap combined with the more successful metrics does increase their predictive power for some levels of homophily.

KEYWORDS

Social Networks, Opinion Dynamics, Voting, Communities

ACM Reference Format:
Jacques Bara, Omer Lev, and Paolo Turrini. 2021. Predicting Voting Outcomes in Presence of Communities. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

In today's societies the role of social networks in shaping collective decisions is widely recognised. Recent political developments, involving fake news and private data usage, have also brought to the fore their vulnerability to manipulation attempts, both accidental – those stemming from the underlying network structure [24, 31] – and intentional – those pursued by malign actors [11, 41].

As a consequence of the growing importance of social networks in our public discourse, the research on collective decision-making has witnessed a shift towards the study of social dynamics [18], looking at the spread of information [1, 19, 47] or edge updates [15] as a key factor in determining voting outcomes . The multi-agent systems community, in particular, explored a computational treatment of social dynamics in existing voting models (e.g., Tsang and Larson [44]), studying novel possibilities to manipulate collective outcomes [4, 7, 12, 28], making social choice on social networks an active research topic [27].

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

Following this social take on voting research, a recent *Nature* paper by Stewart et al. [43] has shown how a static graph-theoretic metric, *influence gap*, is highly predictive of how influence dynamics will impact the result of an election, turning minority views (with well-placed supporters) into strong majorities.

More specifically, Stewart et al. [43], through computational simulations of a voter model, backed by social network experiments with human subjects, found strong correlations between the outcome of the voters' decisions and their proposed metric. These results suggest that an increased presence in a voters' social neighbourhood (what they call *influence assortment*) is a good predictor of a party's chances to win elections. In other words, when voters update their preferences looking at their connections, it's the strategic positioning of a party's electorate that matters, rather than the initial majority (what they call *information gerrymandering*). Undoubtedly, a metric that allows us to forego the equilibrium computation of a highly complex system is an important practical tool, significantly simplifying the analysis of the opinion diffusion dynamics, a notoriously complex problem [5, 16]. Moreover, it allows for a further understanding of the effects of manipulation, for example through the strategical placements of bots or zealots to alter the network dynamics.

The results in Stewart et al. [43] are, however, based on a number of limiting assumptions, notably that the analysis is carried out on regular graphs of small degree – each node has three incoming and three outgoing edges – and large scale-free graphs. The main goal of this paper is to examine whether their metric's usage can be expanded beyond this fairly narrow family of graphs to more realistic-looking graphs. In particular, we focus on graphs which are characterized by the presence of a community structure [22, 25], allowing for phenomena such as echo chambers and *homophily* [6, 44]. These are well-established patterns of real-world social networks, e.g., in the evidence on how Americans are sorting themselves into partisan communities [9], and should, in our view, be accounted for by any reasonable model of how collective decisions are affected by social influence.

Our Contribution. Here we study a family of graphs – *homophilic relaxed-caveman* graphs – which, building upon the classical clique-like community model of caveman graphs [45], introduces further variance and reality-resembling interactions by determining connections as a noisy function of the degree of homophily. The key difference thus to Stewart et al. [43], is that we look at networks with community structure, while they only consider regular and scale-free graphs. We compute the influence gap for these graphs and show how homophily and rewiring interact to produce a rather unexpected pattern, which will feature in our main findings. We then examine the influence gap against the final voting outcome, using an empirical opinion diffusion model and parameters [43].

We see that with equal representation but varying influence gap, the latter no longer correlates with the final voting outcome as it did in Stewart et al. [43]. Once we extend the scope of the statistics, though, to include a multitude of starting partisan majorities, the gap generally correlates strongly with the final voting outcome. In doing so, however, we find that trivially counting the initial majorities generally provides an even better predictor, which is in contrast with Stewart et al.'s proposal. Using regression models, we then determine the complex interrelation between influence gap and initial majority, which we can combine to refine the predictions as a function of homophily and rewiring probability.

Other Related Literature. Discussion of how opinions and ideas spread in society flourished as a research field since Rogers's seminal work [39], which introduced many of the concepts still underlying the field. Since then, research expanded to cases where agents have limited information [20, 32], including on graph structures [10, 35]. There was a particular focus on "information cascades" or "herd mentality", when choices are made sequentially, both when there is a ground truth [8, 21, 37] and where there is none [3, 46]. We use this basic assumption that people wish to conform to their surroundings in this paper, as well.

A closely related avenue of research concerns opinion diffusion models, where agents are recipients of social influence and opinions spread in a network. Research on this has been both empirical [14] and theoretical [4, 26, 29] (see overviews in Mahajan et al. [36] and Young [49]), including attempts to find influential nodes in the social graph [33]. Computational models of opinion diffusion have looked at the fixed-point properties of the graph dynamics, in connection with consensus formation [5] and its complexity [16]. An important stream of research has looked at how to control opinion diffusion by external intervention, for example through bribery [12] or false-name attacks [13], which is naturally connected to the view we are taking here.

Paper Structure. In Section 2 we introduce our setup and the basic graph-theoretic terminology, in particular the model of (relaxed) caveman-graphs. We also present some analytical considerations on the computation of the influence gap in such graphs. In Section 3 we present our proposed homophilic extension, together with the algorithms to control the homophily level and the rewiring probability. Section 4 introduces the opinion diffusion dynamics, measuring the predictive power of influence gap and other metrics in our models. These results are further discussed in Section 5, where we compare metrics and establish the effects of homophily and rewiring. We conclude in Section 6 presenting various follow-up research directions.

2 INFLUENCE GAP AND COMMUNITIES

2.1 Influence Gap

Our social graph is defined by a structure $G = (V, E)$ where $V = \{1, 2, ..., N\}$ is the set of nodes, our agents (or voters), and E is a set of undirected, unweighted edges that signify a social connection/acquaintance. Like much of the research on opinion dynamics (including, in particular, Stewart et al. [43]), we shall focus on the two-party (or two-opinion) setting and we shall henceforth refer to the parties as colors, red and blue. The case with more than two parties remains an important follow-up research direction.

Over our social graph we therefore overlay a partisan structure $\mathcal{P} = \{\mathcal{R}, \mathcal{B}\}$, and each node is associated with a particular party by a function, the *party assignment*, $p : V \to \mathcal{P}$. The number of voters for the red (blue) party is R (B), and each node $n \in V$ sees the votes of its neighbours, $\mathcal{N}(n) = \{m \mid (n, m) \in E\}$, and itself, henceforth its *poll*, $\mathcal{N}(n) \cup n$. In particular, each voter sees Δ_n, the fraction of its poll voting for their party $p(n)$, again, including themselves. *Influence assortment* [43] is, intuitively, the relative advantage of a party against its rivals, and acts on two different levels: on the level of a single node n, denoted by a_n; and on the level of a party, say red, denoted by $\mathcal{A}_R$.

The following table defines these notions formally, including the *influence gap* IG_R, as the advantage in assortment of party $\mathcal{R}$ (resp., IG_B denotes its dual). Note the use of the Kroneker delta $\delta_{i,j}$, which is 1 if $i = j$ and 0 otherwise.

$$a_n = \begin{cases} \Delta_n & \Delta_n \geq \frac{1}{2} \\ -(1 - \Delta_n) & \Delta_n < \frac{1}{2} \end{cases} \tag{1}$$

$$\mathcal{A}_R = \frac{1}{R} \sum_{n \in V} a_n \delta_{p(n),\mathcal{R}} \tag{2}$$

$$IG_R = \mathcal{A}_R - \mathcal{A}_B \tag{3}$$

Influence assortment on the level of nodes, a_n, can be thought of as the extent to which an agent's party is present in its own poll, and thus how much they can be influenced to vote to a different party. Its value (regardless of sign) highlights how homogeneous a node's neighbourhood is, while its sign indicates if a node belongs to the majority party in the local neighbourhood. The mean of node assortments over nodes of a single party is then the party assortment, $\mathcal{A}_R$ (resp., $\mathcal{A}_B$). The influence gap, IG_R (resp., IG_B), can therefore be understood as the difference in how "strategically placed" a party is – how much do its supporters interact with other parties (and therefore, able to be influenced by them). Henceforth, we omit the party index from IG whenever clear from the context.

Throughout the paper we focus on the case of a strong party assignment (SPA), following Alon et al. [4], where each party is assigned a fixed fraction of nodes – typically this will be a half but in later sections we consider non-equal representation. [1]

2.2 Caveman Graphs

A caveman graph $G = (V, E)$ is a set of l isolated cliques each of size k [45].[2] These graphs are collections of cliques that encode a very basic form of community-structure formation without showing interesting variety nor appearing much in empirical networks [48]. A *relaxed-caveman* graph is a modified version of this, whereby the edges are rewired with some given probability. Concretely, given

[1] Weak party assignment (WPA), in contrast, assigns to every node a party P with some probability, e.g., they are red with probability $\frac{3}{4}$. While we leave the treatment of WPA for future research, we note that our random graph generation models with non-equal representation are de facto working with a constrained form of WPA.

[2] A connected version of caveman graphs is formed by rewiring a single edge per clique it to a node in an adjacent clique along a central cycle [45], also commonly referred to as a caveman graph. For the purposes of this paper a caveman graph will be taken to mean the unconnected set of isolated cliques.

a probability p and iterating over all edges E of a set of isolated cliques, an edge $(u, v) \in E$ is rewired as (u, n), for $u, v, n \in V$ [23]; if (u, n) already exists, nothing happens, such that all new edges are between nodes of different cliques. This extension provides a fairly diverse and intuitively clear set of communities, without the need to rigorously define the concept of community itself or to delve into the plethora of community detection [2, 23, 40] and generation [30, 34] methods.

Cavemen graphs and their relaxed variant display interesting analytical properties when it comes to computation of the influence gap. We now show that in the two party case, where each party has equal number of voters, this has a unique solution and how this paves the way to an intuitive explanation of the IG's predictive power under a reasonable model of influence dynamics. The IG computation procedure can be extended beyond the two party case although, in general, will not have a unique solution.

First, consider a graph of l cliques, labelled $c \in \{1, \cdots, l\}$, each of size k. A red node n in clique c, with total number of red nodes x_c, sees exactly $\Delta_n = {}^{x_c}/_{k}$ red nodes in its poll, as, by definition, it sees all of the nodes inside its clique. This clearly holds for all other red nodes inside c, as well. Thus, the influence assortment of any red node in c is equal, so that $a_n = a_m$ for all nodes $n, m \in c$; the sum of influence assortments over all red nodes in c is therefore $x_c a_n$.

Once all x_c are known, we can group and loosely order cliques together into three types: those in which red holds a strict majority, those in which there's an exact tie and those in which blue holds a strict majority. To this end $\exists M, M' \in \{1, \cdots, l\}$ with $M' \leq M$ such that for $n \in V$:

$$x_c \begin{cases} > 1/2 & \text{for } c \leq M' \\ = 1/2 & \text{for } M' < c \leq M \\ < 1/2 & \text{for } M < c \end{cases}$$

In other words, M' is the number of cliques in which the red party hold a strict majority, M is the number of cliques where they hold a weak majority in and $M - M'$ is the number of cliques with equal representation. Finally, denoting the sum of x_c's from $c = 1$ up to $c = d$ (for $1 \leq d \leq l$) as $X_d \equiv \sum_{c=1}^{d} x_c$, knowing the cliques with red majority M allows us to calculate the red assortment on the level of the party, $\mathcal{A}_R$.

$$\begin{aligned} \mathcal{A}_R &= \frac{1}{R} \sum_{n \in V} a_n \delta_{p(n),\mathcal{R}} \\ &= \frac{1}{R} \left(\sum_{c=1}^{M} \left(x_c \frac{x_c}{k} \right) + \sum_{c=M+1}^{l} x_c \left(\frac{x_c}{k} - 1 \right) \right) \\ &= \frac{1}{R} \left(\sum_{c=1}^{l} \frac{x_c^2}{k} - (R - X_M) \right) \\ \mathcal{A}_R &= \frac{1}{R} \left(\sum_{c=1}^{l} \frac{x_c^2}{k} + X_M \right) - 1 \end{aligned} \tag{4}$$

We can find the equivalent for the blue party by making two observations. First, the influence assortment of any blue node $n \in c$ is $b_n = -a_m$ for a red node $m \in c$. Second, the number of blue nodes in clique c is $y_c = k - x_c$ hence the number of blue nodes up to clique d is $Y_d \equiv \sum_{c=1}^{d} y_c = dk - X_d$. We note that since there are M' cliques that contain a strict red majority, then equivalently in these M' cliques blue is strictly a minority. Following similarly from the red party, the influence assortment of the blue party, $\mathcal{A}_B$, in terms of red counts x_n is thus as follows.

$$\mathcal{A}_B = \frac{1}{B} \left(\sum_{c=1}^{l} \frac{x_c^2}{k} + X_{M'} + N - 2R - M'k \right) \tag{5}$$

Finally, this gives us an expression for the influence gap IG in favour of the red party for a general set of l isolated cliques, under any party assignment, noting that $N = B + R = lk$.

$$IG = \frac{X_M}{R} - \frac{X_{M'}}{B} + \frac{B - R + M'k}{B} - 1 + \frac{B - R}{RB} \sum_{c=1}^{l} \frac{x_c^2}{k} \tag{6}$$

In the case of a strong party assignment with equal representation, i.e., $B = R = {}^{N}/_{2}$, we can constrain M and M' by noting that there must be as many strict red majorities as strict blue minorities $M = l - M'$. Furthermore, the term $X_M - X_{M'} = (M - M') \cdot {}^{k}/_{2}$ since all the cliques in $M - M'$ have equal representation.

$$\begin{aligned} IG &= \frac{2}{N}(X_M - X_{M'}) + \frac{2M'k}{lk} - 1 \\ &= \frac{2}{lk}(M - M')\frac{k}{2} + \frac{2M'}{l} - 1 \\ &= \frac{M + M'}{l} - 1 \\ IG &= 0 \end{aligned} \tag{7}$$

Let us now notice, in light of this observation, how in graphs with rigid communities, IG can fail to be predictive with equal representation. Consider to this extent a simple update model where voters never change their mind unless most of their friends are different, in which case they stick to their party with probability Δ_n. In Figure 1, for example, we expect such an opinion diffusion model to converge to an outcome where the red party conquers the clique on the left but loses the other two. This, note, would also hold under a number of reasonable update rules alike. For these basic models, we start seeing how community-based structures "locally" converge to robust equilibria.

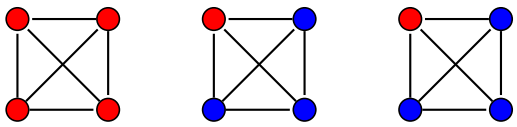

Figure 1: A caveman graph with split majority and IG of 0.

Strong party assignments on isolated cliques thus present a case for which influence gap can, even at the intuitive level, fail to predict the outcome of a particular voting game, counter to Stewart et al. [43]. As we pointed out earlier, such communities are clearly somewhat artificial and unrealistic. Instead, relaxations of these graphs present a more realistic model of social networks while having non-zero influence gap. For low rewire probabilities we note that the results above can be extended further to relaxed-caveman graphs by considering the effects of a small number of rewired edges as perturbations $O(1/\min(R, B))$, since at worst for a single rewire the assortment of a node gets changed by ± 1 and thus

its contribution to the influence gap changes by $O(1/R)$ or $O(1/B)$. The observation also carries over to our homophilic extension, which we present next.

3 HOMOPHILIC CAVEMAN GRAPHS

Relaxed-caveman graphs rewire the edges of the original caveman graph *without looking at party assignment*. This means that, effectively, the resulting graph, though exhibiting a rich community structure, abstracts away from the relation between connections and opinions, which is typical of real-world networks, where the two are highly intertwined [6, 9]. To address this issue, we propose a modification to the relaxed-caveman graph as the *homophilic relaxed-caveman* (hRC) graph model in a similar fashion to the homophilic Erdős-Rényi and Barabási-Albert graphs used in Tsang and Larson [44]. This allows us to generate synthetic graphs with communities where the graph structure is dependent on the party assignment, following the observed behavior [9] that people tend to cluster with people who share their views.

Algorithm: Homophilic Relaxed-Caveman Graph, G

(1) **Initialise** G as a set of l cliques each of size k
(2) **for** $(u, v) \in E$:
- Choose at random an $n \in V, n \neq v$
- **if** $p(u) = p(n)$:
  - $\tilde{p} = p_0 h$
- **else**:
  - $\tilde{p} = p_0(1 - h)$
- Rewire (u, v) as (u, n) with probability $\tilde{p}$

(3) **return** G

The algorithm above describes how to generate the hRC graph, starting from a set of disconnected cliques $G = (V, E)$. It becomes hRC graph $\mathcal{G}_p(l, k, p_0, h)$, with l communities, each of size k, with *rewiring probability* p_0 and *homophily factor* h, given a party assignment $p : V \to \mathcal{P}$. The probability to rewire p_0 can be thought of as the likelihood of changing a pre-existing friendship to a new friend, while the homophily is the probability of agent u's new friend n voting for the same party.

We highlight two important subclasses of the hRC model:

- For $h = 0.5$ we recover the relaxed-caveman graph with rewire probability $p = p_0/2$ since a node is equally likely to be rewired to its party as it is to the opposing party, but with probability $p_0 \cdot 0.5$.
- For high values of p_0 nodes from different cliques intermingle sufficiently enough that the community structure begins to fade. This becomes problematic due to a wide variety in the definition and detection of community [23, 40], such as via modularity-based methods [38].

We now show how the influence gap is distributed in homophilic relaxed-caveman graphs with equal partisan split ($R = B$). To do so we generated hRC graphs across the entire range of homophily and rewire probability. Due to the randomized nature of the model, for any given set of parameters h and p_0, we produced 10,000 graphs and found the mean of their influence gaps, towards whichever party had the higher influence assortment. Note how this, therefore, represents the ability for either party to open an advantage over its opponent, not for a specific one. If, instead, we measured the gap towards a specific party, say red, then we would expect the mean of the gap to be 0 due to symmetry.

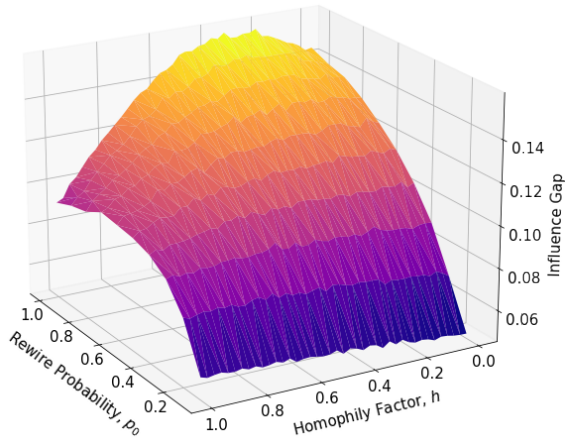

Figure 2: A surface plot of the mean absolute value of influence gap, across different homophily factors h and a range of rewire probabilities p_0. Each point is measured from 10^4 party assignments each generating a single graph.

We find that across all values of h, the absolute value of influence gap increases monotonically with the rewire probability p_0. That is, any cross-section in h of Figure 2 is monotonic in p_0. This is consistent with our results on non-relaxed-caveman graphs (Equation 7), in which the influence gap was 0 (the effects of homophily are irrelevant for such low p_0 values, since these graphs are statistically indistinguishable from the equivalent relaxed-caveman).

The relationship to homophily, however, is more complex, particularly when the community structure is less apparent, as with high p_0 value. Specifically, the influence gap is asymmetrically unimodal, with a peak at around $p_0 = 1$ and $h = 0.3$. Notice that this means parties are marginally better off when the likelihood of forming a community between the party members is lower than with the opposing party. This asymmetry, we believe, is in large part due to the definition of influence assortment for a node. When the neighbourhood $\mathcal{N}(n)$ of an agent n is just slightly in support of the opposite party but its own preference is enough to push the poll into a (weak) majority $\Delta_n \geq 0.5$ then the maximal value of influence gap is reached, when most/all nodes face this situation. The asymmetric behaviour of influence gap in hRC graphs is reflected in our overall metric comparison (Figure 7), where the predictive power of IG is overtaken by majority exactly from this peak on.

4 PREDICTING VOTING OUTCOMES

In this section we examine whether the findings of Stewart et al. [43] still hold when communities are present. Our conclusion, in summary, is that influence gap does not capture the result of the influence dynamics in case of a given partisan split, incidentally also a central assumption in their analysis, with a rather poor correlation. When the assumption of a given representation is relaxed, the correlation is restored, but a more predictive and computationally simpler metric exists, namely the initial majority of either party.

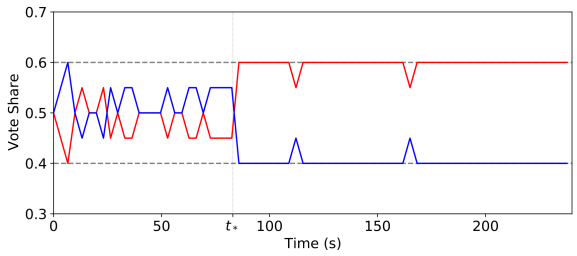

Figure 3: Example time series of a voting game, using Stewart et al. [43] behavioural model on a homophilic relaxed-caveman graph with rewire parameter $p_0 = 0.3$ and homophily $h = 0.3$. The two parties initially have equal vote share but eventually the red party reaches a super-majority of $V = 0.6$ (the dashed lines are V and $1 - V$). The vertical dotted line at t_* represents the transition between early and late phases of the game.

4.1 Voter Model

For N voters, at least half are assigned to the red party and the remainder to the blue, and all are placed in an influence network. A voter's knowledge is restricted to a subset of the entire graph: only knowing the voting intentions of its neighbours as well as its own, serving as a form of poll, to which it wishes to conform. The game lasts for a fixed amount of time, during which players can change their voting intentions synchronously[3]. The winning party is the one to hold a super-majority above a threshold $V > 0.5$ when the updating process is done, otherwise we consider it a deadlock.

The agents follow a stochastic behavioural model developed by Stewart et al. [43], informed by a social experiment with human subjects who were given pay-offs depending on the success of their assigned party. At any given time, a voter, according to the behavioural model, would vote for their assigned party with a probability that depends on a) what their surroundings predict will happen (win, lose or deadlock) and b) the stage of the game (early or late). In other words, for each individual there exists a family of six parameters p_{ij}, where $i \in \{\text{win, lose, deadlock}\}$ is the poll's prediction and $j \in \{\text{early, late}\}$ is the stage of the game, that are precisely these probabilities, henceforth *strategies*.

Table 1: Mean agent strategies $\bar{p}_{ij}$, such that an agent with a neighbourhood poll predicting state i during phase j will stick to their party, on average, with probability $\bar{p}_{ij}$. These values were inferred from social experiments (see the Supplementary Material of Stewart et al. [43] for an extensive discussion of how these numbers are obtained and why these are used independently of the graph structure).

$\bar{p}_{ij}$	*Early*	*Late*
Win	0.975	0.979
Deadlock	0.964	0.911
Lose	0.598	0.574

[3]Note the difference from the iterative voting model, as used by Tsang and Larson [44], in which every voter updates separately to the others. Whereas, ours is more akin to the update model of Alon et al. [4].

Since voters are not a homogeneous bunch and will have different strategies, each parameter is sampled from the empirical distribution of the social experiment. Thus while p_{ij} is a random variable, which has an empirical distribution with mean given by Table 1, for each voter v a set of 6 parameters p_{ij}^v are realisations of the random variables. We point out that the behavioural parameters can be used independently of the initial structure, as voters' are unaware of the graph they sit in, an assumption also made by [43] when replicating their findings with simulations.

4.2 Benchmark Metrics

As the exact analysis of the opinion dynamics faces important complexity barriers – Stewart et al. [43] proposed using the IG as a prediction tool. To IG we add three more metrics, which we will compare against:

Majority Which party has the majority in the beginning state.

Deterministic voter skew (dVS) A deterministic simplification of the update dynamics; at each time step every agent synchronously conforms to the strict majority party in their poll, keeping the current choice in case of a tie, and after σ steps the voter skew is measured. In principle, one could evolve the system for as many steps σ as in the stochastic process, but we use $\sigma = 1$, as errors due to the simplification may be propagated and worsened with more steps.

Efficiency gap (EG) A political science metric [42], developed to measure gerrymandering, in which we examine how many votes were "wasted", i.e., could have been eliminated without changing the outcome.

4.3 Experiments

Each simulation of the voter game runs for 240 seconds that starts with the early phase of the game and transitions to the late phase after 83 seconds (both times chosen to be consistent with the social experiments of Stewart et al. [43]). The $N = 20$ voters are assigned a party – 10 are assigned red while the other 10 are given the blue party – and then are placed in a graph G, generated by a number of different assignments of h and p_0.

During the game, every 3.3 seconds a voter, v, can update their intention, with probability p_{ij}^v, which are sampled from the empirical parameter-distributions found from the human social experiment [43]. In total, for a single simulation N samples are taken from 6 distributions each. After 240 seconds have elapsed the vote share across the entire graph is measured.

For clarity we outline in Table 2 the three synthetic datasets we produced, each broader and more expansive than the last. The first studies the effects of communities being present in relaxed and homophilic relaxed-caveman graphs with equal representation. The second broadens the notion of correlation to include unequel representations, focusing only on hRC graphs for a particular set of parameter values. The final and most comprehensive delves more deeply into the effects of community strength (as controlled by the rewire p_0) and echo chambers (homophily h) on the predictions of the metrics noted above.

An example of a time series produced by a simulation is shown in Figure 3. The different sections of the plot, partitioned by dashed and dotted lines, represent different strategies. For example, in the

Table 2: Summary table of the synthetic datasets (labelled 1, 2 and 3, from top to bottom) generated of homophilic relaxed caveman graphs. Metrics were evaluated at the start of the dynamics, while voter skews were measured post-election. The size of a dataset is measured in the number of elections simulated.

Majority	*Parameters*	*Metrics*	*Size*
0	$h = 0.3, p_0 = 1$	IG	10^5
0-9	$h = 0.3$, $p_0 \in \{0, 0.1, 0.5, 0.9, 1\}$	Majority, IG	$5\text{x}10^6$
0-9	$h \in \{0, 0.1, \cdots, 1\}$, $p_0 \in \{0, 0.2, \cdots, 1\}$	Majority, IG, dVS, EG	$1.2\text{x}10^7$

early phase ($t < t_*$) both parties are deadlocked and thus agents vote for their assigned party with probability $p^{\text{agent}}_{\text{deadlock,early}}$. The convergence of a time series is not a guaranteed because, as in the original social experiment, the game finishes after 240 seconds.

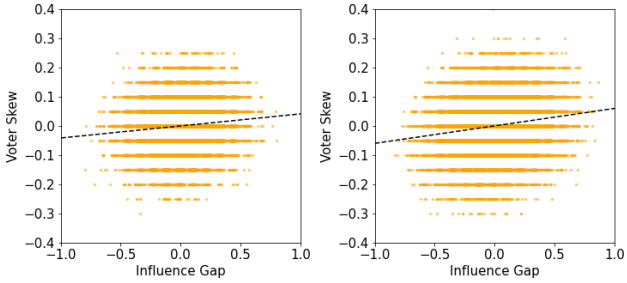

Figure 4: Simulations of the behavioural model of Stewart et al. [43] acting on relaxed- and homophilic relaxed-caveman graph from dataset 1. In both cases the influence gap correlates very weakly with voter skew, having a Pearson's $\rho < 0.2$, and passing a significance test with $p < 10^{-5}$. Left: for the relaxed-caveman graph with rewire parameter $p = 0.3$ gives rise to a Pearson correlation of $\rho = 0.104$. Right: for the homophilic relaxed-caveman with rewire probability $p_0 = 1$ and homophily factor $h = 0.3$ the correlation is slightly more pronounced with $\rho = 0.160$.

4.4 Results

Where Stewart et al. [43] found strong correlation between IG and election outcome in scale-free (Barabási-Albert) graphs for a given initial voter skew, we find the contrary in RC and hRC graphs. Starting with equal representation – same number of red nodes as blue – we find that the presence of communities suppresses the correlation noted in the original paper Stewart et al. [43]. In both the relaxed-caveman and the homophilic relaxed-caveman with given parameter sets, the Pearson correlation coefficient is small – $\rho < 0.2$ – as seen in Figure 4.

Communities in the hRC graphs reduce the predictive quality of the influence gap, not only due to the dynamic effects, but because of its constraints on the influence gap itself. As shown by Equation 7, for equi-sized cliques with equal representation the influence gap is 0. Minor modifications to the graph via edge rewires induce a bounded effect on the influence gap where on average the value is not very different from 0.

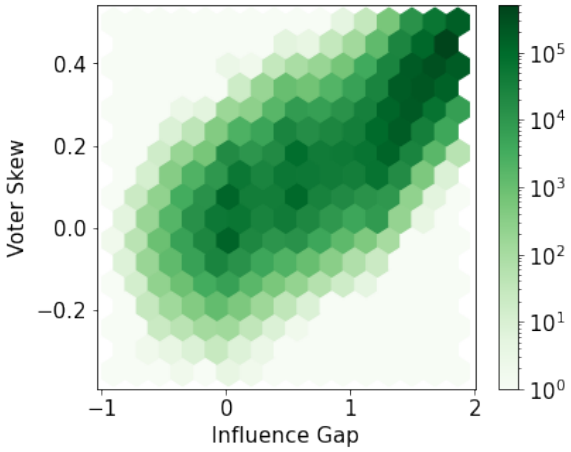

Figure 5: Histogram of influence gap and final voter skew for homophilic relaxed-caveman graphs generated by dataset 2. Colours denote the frequency of each bin. The relationship is statistically strong, with Pearson correlation of $\rho = 0.891$.

We observe a more interesting behaviour as we extend the scope of the statistics to include multiple initial voter skews, i.e., initial partisan majorities. As we vary the red party's non-negative majority more graphs with higher IG – in favour of red – are generated and correspondingly see elections that are increasingly biased to the red party. Looking at Figure 5, a visibly strong correlation with a Pearson correlation of $\rho = 0.891$ arises, restoring the high correlation Stewart et al. [43] originally observed. Notice how the left image in Figure 4 can be seen as an embedded subset of Figure 5, as a small patch of points centered around 0. Further, increasing the majority shifts the outcomes (the green patches in the graph) to higher influence gap and voter skew and reduces its size. The initial majority, therefore, presents itself as a useful control variable (of influence gap) and as a predictive metric.

In fact, as shown by Figure 6, for a particular set of parameter values, majority predicts the election outcome more accurately than influence gap, with Pearson $\rho = 0.904$. The trivial act of counting the number of nodes a party has provided a stronger correlation than the more complex metric proposed in Stewart et al. [43]. Although troublesome, this does not spell the end for influence gap. Using several regression models (see Section 5.1) knowing both majority and IG improves the predictions from either individually. In other words having counted the nodes, knowing the IG can still improve predictions for the final outcome.

Varying homophily and rewiring. We now investigate the effects of network generation on the dynamics and, more importantly for this work, the effectiveness of metrics. The homophilic relaxed-caveman model allows us to change both the strength of the communities, through the rewire probability p_0, and the echo chambers

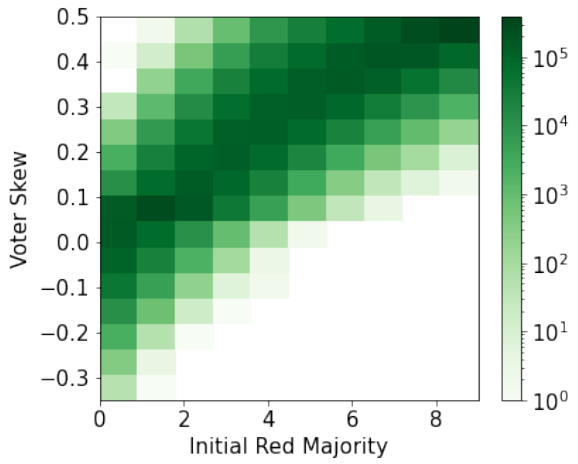

Figure 6: Histogram of initial majority and final voter skew for homophilic relaxed-caveman graphs generated by dataset 2. Colours denote the frequency of each bin. Having a Pearson correlation of $\rho = 0.904$, majority predicts the election outcome better than IG.

that form, through the homophily h. For low rewiring the graphs generated are almost a set of equisized cliques, while increasing p_0 tends to dilute communities together. Agents become increasingly entrenched in their partisan social bubble for high homophily while at the low end they actively seek out opposing views[4].

Homophily plays little role at low rewire, as shown by Figure 7; at $p_0 = 0$ all hRC graphs are reduced to the unconnected caveman graph and as such homophily has no impact. As p_0 increases, the homophily becomes more important and in particular a region $h < 0.5$ emerges in which IG predicts an election more accurately than the initial majority (for an intuitive explanation see Section 5.2) suggesting that graph structure is a more important factor when friendship groups are more diverse. The predictive powers of both majority and efficiency gap monotonically and linearly increase with homophily; the gradient of their relationship also increases as the hard community structure is diluted. The other two, however, show non-monotonic behaviour in opposite ways: the correlation of IG is peaked at around $h = 0.7$ while there is a minimum for the dVS at $h = 0.3$.

Curiously the predictions of IG are not affected greatly by rewiring more edges, at least comparatively to the other three metrics. This suggests influence gap has some degree of robustness, in the sense that it performs as well in homophilic/heterophilic social networks, i.e. in a multitude of cases. Take the low homophily case, for example, where both majority and EG perform worse as the communities become more diluted (higher p_0) while the dVS actually improves considerably. At the other end, for high h, all metrics generally improve considerably, the efficiency gap in particular.

[4]To an extent low homophily is the exception and not the rule in empirical social networks. Certainly there are many people who remain open-minded to different opinions, however this is often not the case in the digital world [17].

5 COMBINING METRICS

5.1 Regression Models

In order to explore the role of majority and influence gap (x_M and x_I respectively, following typical regression notation) as predictive or explanatory variables in dataset 2, we use linear regression (Equation 8) to build several models of the voter skew, y. Two models are single-featured using only the initial majority or the influence gap and the third is a joint model built using a multiple regression of both features. All three are trained on the same 70% of the data – multiple majorities for single homophily hRC graphs – and tested on the remaining 30%.

$$y = \beta_M x_M + \beta_I x_I + \beta_0 \tag{8}$$

Table 3: Coefficients of regression, β, and of determination, R^2, for regression models of the dynamic voting game outcome on homophilic relaxed-caveman graphs with $h = 0.3$.

Metric	β_M	β_I	β_0	R^2
Majority	0.0497		0.0445	0.817
IG		0.229	-0.00328	0.793
Majority, IG	0.0300	0.101	0.0132	**0.844**

The regression confirms our observations (see Section 4.4), that majority is a better predictive tool than influence gap. This does not, however, render it useless. In particular the joint model outperforms both individual models that either use majority or influence gap exclusively as shown in Table 3, despite some colinearity (with a variance inflation factor of 19.4) between features.

5.2 Topology and Correlation

At high homophily, $h > 0.5$, hRC graphs are rife with echo chambers in which an agent has friends mostly of the same party and opinion as themselves. Its polls are therefore incredibly homogeneous and it sees no compelling reason to change its vote, meaning the graph structure plays little to no role. Extending this intuition to all agents in the graph, very little diffusion of opinions occur and thus the final outcome of the election will be incredibly reminiscent of and similar to the start. In other words counting the number of votes at the start will closely resemble – and thus predict – the number of votes at the end, and as such majority here can better predict the outcome than IG.

Moreover, this logic can help explain the general trend that the correlation of most metrics increase drastically at high homophily. Since the exogenous party assignment will resemble the final outcome the election itself becomes immensely straight-forward and thus prediction becomes increasingly trivial. As such any reasonable metric should perform better purely on the basis that there has not been great change in the system. Initial majority is likely the best metric in this regime because it is precisely equivalent to measuring the initial voter skew.

On the other hand, at low homophily, $h < 0.5$, most polls are very diverse: an agent will see mostly nodes of the opposite party. In this case enough of her friends disagree with her own opinion, that the

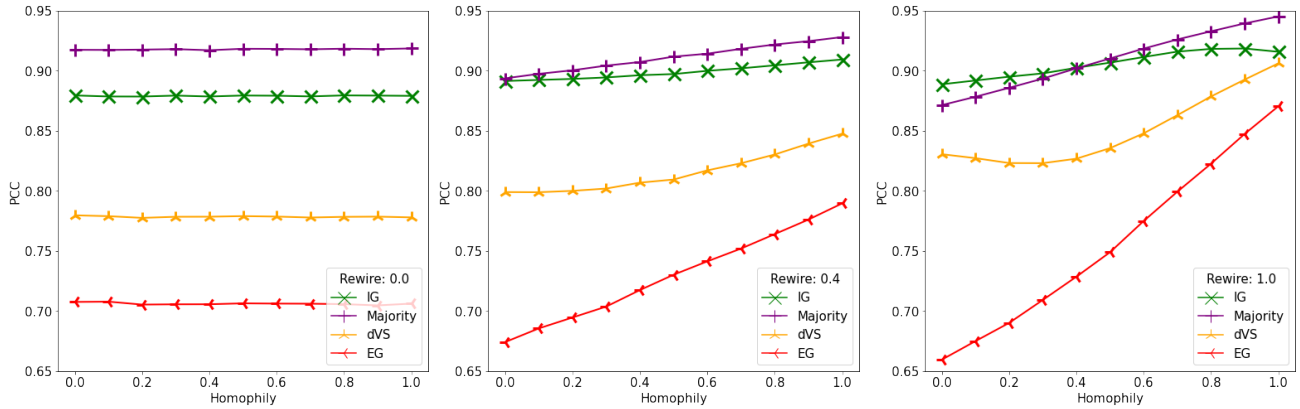

Figure 7: As model parameters of the homophilic relaxed-caveman are varied independently, in dataset 3, the Pearson correlation coefficient (PCC) between the final voter skew and the majority (purple), the influence gap (IG, green), the deterministic voter skew (dVS, orange) and the efficiency gap (EG, red) are plotted. Curves for multiple rewire probabilities are plotted, on the left $p_0 = 0$, in the middle $p_0 = 0.4$ and on the right $p_0 = 1$. Influence gap begins to outperform majority around $p_0 = 0.4$.

agent has a higher probability of changing her views, doubting her choice more frequently. In this scenario the dynamics become less 'stable', and more complex diffusion occurs. The graph structure, therefore, is able to play an important role in how these changes occur and in favour of which party. Hence, initial majority gives poorer predictions.

More formally, at high h the average poll will show a majority towards its agent's party $\Delta_v > 0.5$ or even a super-majority $\Delta_v > V$. During phase j of the voting game, most voters v see a prediction of i = win, so that their strategy is $p^v_{\text{win},j}$. Since the empirical distribution for $p_{\text{win},j}$ is heavily biased towards $p_{\text{win},j} \geq 0.9$ most voters will stick to their initial opinion. Conversely for $h < 0.5$ an agent's poll, at least initially, likely shows a super-minority $\Delta_v < 1 - V$ or equivalently i = lose such that over 40% of agents will change their votes.

6 DISCUSSION AND CONCLUSION

We proposed a novel graph model, the homophilic relaxed-caveman, as a means to generate synthetic graphs with communities that may exhibit echo chambers. Starting with the recently proposed metric of influence gap as predictor of voter skew in dynamic opinion formation models, we provide algorithms to compute it in our graphs.

Moreover, we show that the presence of communities suppresses the power of the influence gap, unless a broader understanding of correlation, one that considers multiple starting majorities, is utilised. When doing so, the correlation with the final outcome is strengthened and the value of the metric's predictions is restored. However, this is not without caveats: for some parameter-regimes of the homophilic relaxed-caveman model, a trivial aggregation of initial voter intentions is an even more informative metric. This, as we further analysed, is due to the presence of echo chambers that are conducive to more predictable dynamics.

Having measured the efficacy of several metrics, this poses the question of whether a metric exists that can predict the voting dynamics most accurately. We envisage such metric to be explicitly dependent on party assignment and should be predictive while computationally easy to compute.

Another potential direction concerns the difference in voters' behaviour. In our hRC models, the homophily level and the probability of rewiring are the same for both parties. However, we may want to distinguish between electorates that have different levels of approaching others. For example more open-minded voters, who do not mind accepting connections that do not share their view, while others are more close-minded. Preliminary results suggest that a party can leverage a higher homophily factor in this way to structure the network to its favour.

Additionally, it is important to explore the models with more than two parties. With two parties, influence assortment and gap are very well-defined measures. In contrast, straightforward multi-party extensions are not unique and the choice of one depends on the context. For example, does defining a_i through the "net influence of the party with the plurality of positive influence" [43] make sense? Moreover, voters preferences may come into play: do they have a ranked list of parties or a single favoured party?

The dynamic model has only shown that the outcome of an election on networks with community can be biased by the structure. However, we may want to model manipulation explicitly, by allowing parties to insert artificial bots or zealots, as Stewart et al. [43] did, to influence members behaviour. Just like the forceful agents in Acemoglu et al. [1], these can be modelled as nodes that have a party affiliation but never update their view.

Finally, agents may be allowed to actively seek new friendships, whether to express their views more widely or to receive more opinions - the distinction becoming important when considering directed networks. Can a party utilise these dynamic connections to its advantage?

ACKNOWLEDGMENTS

This research was partly supported by ISF grant 1965/20, GIF grant I-2527-407.6/2019 and EPSRC grant EP/S022244/1 for the MathSys CDT.

REFERENCES

[1] Daron Acemoglu, Asuman Ozdaglar, and Ali ParandehGheibi. 2010. Spread of (mis)information in social networks. *Games and Economic Behaviour* 70, 2 (2010), 194 – 227.

[2] Rodrigo Aldecoa and Ignacio Marín. 2013. Surprise maximization reveals the community structure of complex networks. *Scientific Reports* 3, 1 (14 Jan 2013), 1060.

[3] Noga Alon, Moshe Babaioff, Ron Karidi, Ron Lavi, and Moshe Tennenholtz. 2012. Sequential voting with externalities: herding in social networks. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. Valencia, Spain, 36.

[4] Noga Alon, Michal Feldman, Omer Lev, and Moshe Tennenholtz. 2015. How Robust is the Wisdom of the Crowds?. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)* (Buenos Aires, Argentina) *(IJCAI'15)*. AAAI Press, 2055–2061.

[5] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. 2018. Reasoning about Consensus when Opinions Diffuse through Majority Dynamics. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 49–55.

[6] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.

[7] Dorothea Baumeister, Ann-Kathrin Selker, and Anaëlle Wilczynski. 2020. Manipulation of Opinion Polls to Influence Iterative Elections. In *Proceedings of the Nineteenth Conference on Autonomous Agents and MultiAgent Systems (AAMAS-20)*. International Foundation for Autonomous Agents and Multiagent Systems, Auckland, New Zealand, 132–140.

[8] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100, 5 (October 1992), 992–1026.

[9] Bill Bishop. 2009. *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart.* Mariner Books.

[10] Lawrence E. Blume. 1993. The Statistical Mechanics of Strategic Interaction. *Games and Economic Behavior* 5, 3 (June 1993), 387–424.

[11] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. 2018. Big City vs. the Great Outdoors: Voter Distribution and How It Affects Gerrymandering. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, Jérôme Lang (Ed.). Stockholm, Sweden, 98–104.

[12] Robert Bredereck and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 894–900.

[13] Markus Brill, Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2016. False-Name-Proof Recommendations in Social Networks. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls (Eds.). ACM, 332–340.

[14] Jacqueline Johnson Brown and Peter H. Reingen. 1987. Social Ties and Word-of-Mouth Referral Behavior. *Journal of Consumer Research* 14, 3 (December 1987), 350–362.

[15] Matteo Castiglioni, Diodato Ferraioli, and Nicola Gatti. 2020. Election Control in Social Networks via Edge Addition or Removal. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 02 (Apr. 2020), 1878–1885.

[16] Dmitry Chistikov, Grzegorz Lisowski, Mike Paterson, and Paolo Turrini. 2020. Convergence of Opinion Diffusion is PSPACE-Complete. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 05 (Apr. 2020), 7103–7110.

[17] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2020. Echo Chambers on Social Media: A comparative analysis. arXiv:2004.09603 [physics.soc-ph]

[18] James S. Coleman. 1990. *Foundations of Social Theory.* Harvard University Press.

[19] Federico Corò, Emilio Cruciani, Gianlorenzo D'Angelo, and Stefano Ponziani. 2019. Exploiting Social Influence to Control Elections Based on Scoring Rules. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 201–207.

[20] Joseph Farrell and Garth Saloner. 1985. Standartization, Compatability, and Innovation. *Rand Journal of Economics* 16, 1 (1985), 70–83.

[21] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. Reaching Consensus via non-Bayesian Asynchronous Learning in Social Networks. In *Proceedings of the 17th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*.

[22] G. W. Flake, S. Lawrence, C. L. Giles, and F. M. Coetzee. 2002. Self-organization and identification of Web communities. *Computer* 35, 3 (2002), 66–70.

[23] Santo Fortunato. 2010. Community detection in graphs. *Physics Reports* 486, 3-5 (Feb 2010), 75–174.

[24] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 913–922.

[25] M. Girvan and M. E. J. Newman. 2002. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99, 12 (2002), 7821–7826.

[26] Jacob Goldenberg, Barak Libai, and Eitan Muller. 2001. Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Marketing Letters* 12, 3 (2001), 211–223.

[27] Umberto Grandi. 2017. Social Choice and Social Networks. In *Trends in Computational Social Choice)*, Ulle Endriss (Ed.). AI Access.

[28] Umberto Grandi and Paolo Turrini. 2016. A Network-Based Rating System and Its Resistance to Bribery. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, Subbarao Kambhampati (Ed.). IJCAI/AAAI Press, 301–307.

[29] Mark S. Granovetter. 1973. The Strength of Weak Ties. *Amer. J. Sociology* 78, 6 (May 1973), 136–1380.

[30] Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social Networks* 5, 2 (1983), 109 – 137.

[31] Lorien Jasny, Joseph Waggle, and Dana R. Fisher. 2015. An empirical examination of echo chambers in US climate policy networks. *Nature Climate Change* 5, 8 (2015), 782–786.

[32] Michael L. Katz and Carl Shapiro. 1985. Network Externalities, Competition, and Compatibility. *The American Economic Review* 75, 3 (June 1985), 424–440.

[33] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *International conference on Knowledge discovery and data mining (KDD)*. Washington D.C., 137–146.

[34] Clement Lee and Darren J. Wilkinson. 2019. A review of stochastic block models and extensions for graph clustering. *Applied Network Science* 4, 1 (Dec 2019).

[35] Dunia López-Pintado and Duncan J. Watts. 2008. Social Influence, Binary Decisions and Collective Dynamics. *Rationality and Society* 20, 4 (November 2008), 399–443.

[36] Vijay Mahajan, Eitan Muller, and Frank M. Bass. 1990. New Product Diffusion Models in Marketing: A Review and Directions for Research. *The Journal of Marketing* 54, 1 (January 1990), 1–26.

[37] Elchanan Mossel, Joe Neeman, and Omer Tamuz. 2014. Majority Dynamics and Aggregation of Information in Social Networks. *Autonomous Agents and Multi-Agent Systems* 28, 3 (May 2014), 408–429.

[38] M. E. J. Newman. 2004. Fast algorithm for detecting community structure in networks. *Phys. Rev. E* 69 (Jun 2004), 066133. Issue 6.

[39] Everett M. Rogers. 2003. *Diffusion of innovations* (5th ed.). Simon and Schuster.

[40] Michael T. Schaub, Jean-Charles Delvenne, Martin Rosvall, and Renaud Lambiotte. 2017. The many facets of community detection in complex networks. *Applied Network Science* 2, 1 (15 Feb 2017), 4.

[41] Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. The spread of low-credibility content by social bots. *Nature Communications* 9, 1 (2018), 4787.

[42] Nicholas Stephanopoulos and Eric McGhee. 2015. Partisan Gerrymandering and the Efficiency Gap. *University of Chicago Law Review* 82 (03 2015), 831–900.

[43] Alexander J. Stewart, Mohsen Mosleh, Marina Diakonova, Antonio A. Arechar, David G. Rand, and Joshua B. Plotkin. 2019. Information gerrymandering and undemocratic decisions. *Nature* 573, 7772 (2019), 117–121.

[44] Alan Tsang and Kate Larson. 2016. The Echo Chamber: Strategic Voting and Homophily in Social Networks. In *Proceedings of the 15th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (Singapore, Singapore) *(AAMAS '16)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 368–375.

[45] Duncan J. Watts. 1999. Networks, Dynamics, and the Small-World Phenomenon. *Amer. J. Sociology* 105, 2 (1999), 493–527.

[46] Duncan J. Watts. 2002. A Simple Model of Global Cascades on Random Networks. *Proceedings of the National Academy of Sciences of the United States of America* 99, 9 (April 2002), 5766–5771.

[47] Bryan Wilder and Yevgeniy Vorobeychik. 2018. Controlling Elections through Social Influence. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) *(AAMAS '18)*. Richland, SC, 265–273.

[48] Weike Xiao, Jie Ren, Feng Qi, Zhiwei Song, Meng-Xiao Zhu, Hong-Feng Yang, Hui-Yu Jin, Bing Wang, and Tao Zhou. 2007. Empirical study on clique-degree distribution of networks. *Physical review. E, Statistical, nonlinear, and soft matter physics* 76 (10 2007), 037102.

[49] H. Peyton Young. 2009. Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review* 99, 5 (2009), 1899–1924.