

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Artificial Intelligence

www.elsevier.com/locate/artint

Predicting voting outcomes in the presence of communities, echo chambers and multiple parties



Jacques Bara^a, Omer Lev^b, Paolo Turrini^{a,*}

^a University of Warwick, Coventry CV47AL, United Kingdom

^b Ben-Gurion University, Beer-Sheva 8410501, Israel

ARTICLE INFO

Article history:

Received 9 August 2021

Received in revised form 4 August 2022

Accepted 8 August 2022

Available online 12 August 2022

Keywords:

Social networks
Opinion dynamics
Voting
Communities
Echo chambers
Influence gap

ABSTRACT

When individuals interact in a social network their opinions can change, at times quite significantly, as a result of social influence.

In elections, for example, while they might initially support one candidate, what their friends say may lead them to support another. But how do opinions settle in a social network, as a result of social influence?

A recently proposed graph-theoretic metric, the *influence gap*, has shown to be a reliable predictor of the effect of social influence in two-party elections, albeit only tested on regular and scale-free graphs. Here, we investigate whether the influence gap is able to predict the outcome of multi-party elections on networks exhibiting community structure, i.e., made of highly interconnected components, and therefore more resembling of real-world interaction. To encode communities we build on the classical model of caveman graphs, which we extend to a richer graph family that displays different levels of *homophily*, i.e., how many connections and opinions are intertwined.

Our contribution is three-fold. First, we study the predictive power of the influence gap in the presence of communities. We show that when there is no clear initial majority the influence gap is not a good predictor of the election outcome. When we instead allow for varying majorities, although the influence gap improves as a predictor, counting the initial partisan majority does consistently better, across all levels of homophily. Second, we study the combined effect of the more predictive metrics, as function of the homophily levels. Using regression models, we demonstrate that the influence gap combined with the initial votes count does increase the overall predictive power for some levels of homophily. Third, we study elections with more than two parties. Specifically, we extend the definition of the influence gap to any number of parties, considering various generalisations, and show that the initial votes count has an even higher predictive power when compared to influence gap than it did in the two-party case.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

* Corresponding author.

E-mail address: p.turrini@warwick.ac.uk (P. Turrini).

1. Introduction

The role of social networks – both person-to-person and online – in shaping collective decisions is widely recognised in today’s society. Recent political developments, from fake news to data leaks, have also brought to light their powerful role in potentially altering election results, and, therefore, their vulnerability to various forms of manipulation [1–4].

As a consequence of the growing importance of social networks – particularly online social networks – in our public discourse, the research on collective decision-making is now looking at the spread of information as a key factor in determining voting outcomes [5–7], with important ramifications for network engineering [8]. The multi-agent systems community, in particular, explored the algorithmic side of social dynamics in existing voting models (e.g., Tsang and Larson [9]), studying novel possibilities to manipulate collective outcomes [10–13] and making the study of social choice on social networks an active research enterprise [14].

When studying the effect of social influence on collective decisions, establishing predictive metrics is paramount. A recent *Nature* paper by Stewart et al. [15] has shown how a simple graph-theoretic metric, the *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, through computational simulations of a voter model, backed by social network experiments with human subjects, Stewart et al. [15] found strong correlations between the outcome of the voters’ decisions and their proposed metric. These results suggest that an increased presence in a voter’s 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 is the strategic positioning on a social network 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 [16–18]. Moreover, it allows for a further understanding of the effects of manipulation, for example through the strategic placements of bots or zealots to alter the network dynamics.

The results in Stewart et al. [15] are, however, based on a number of limiting assumptions, notably the analysis is carried out on two-party elections on regular graphs of degree 3 and large scale-free graphs.

The main goal of this paper is to examine whether the metric’s usage can be expanded beyond this fairly narrow family of graphs to more realistic-looking graphs and into the multi-party case. Our focus is on graphs which are characterised by the presence of a community structure [19,20], allowing for phenomena such as echo chambers and *homophily* [21,9]. These are well-established patterns of real-world social networks, which can be observed, for example, in how Americans are sorting themselves into partisan communities [22], and should, in our view, be accounted for by any reasonable model of how social influence affects collective decisions.

Our contribution We analyse multi-party elections on graphs with community structure, exploring the predictive power of the influence gap and comparing it against various other metrics. We first model communities using what we call *homophilic relaxed-caveman* graphs, which build on the classical clique-like community model of caveman graphs [23]. The homophilic relaxed-caveman graphs introduce variance and reality-resembling interactions by determining connections as a noisy function of the degree of homophily. The key difference from Stewart et al. [15] is that we look at networks with community structure, while they only consider regular and scale-free graphs. We compute how homophily and rewiring interact in these graphs to change the influence gap (Fig. 3) and we then examine the influence gap as a predictor of the final voting outcome, using Stewart et al. [15]’s empirically backed opinion diffusion model and parameters.

We see that with equal initial representation but varying influence gap, the latter no longer correlates with the final voting outcome as it did in Stewart et al. [15] (Fig. 5). Once we look at settings which include varying initial partisan majorities, the gap generally correlates well with the final voting outcome (Fig. 6). However, simply counting the initial majorities is an even better predictor (Fig. 6) and remains consistently so for different levels of homophily and rewiring (Fig. 7). Using regression models, we then determine the interrelation between influence gap and initial majority to strengthen the predictions (Table 2).

Finally, we move to the study of the influence gap in elections with more than two parties. We look at multiple extensions of the originally proposed definitions, analysing their behaviour on simple graphs such as rings, cycles and stars. Using the extended setup we provide a theoretical analysis on clique-like communities and we observe empirically how the initial votes count has an even stronger predictive power when compared to influence gap than it did in the two-party case (Fig. 8).

Other related literature Discussion of how opinions and ideas spread in society flourished as a research field since Rogers’s seminal work [24], who introduced many of the concepts still underlying the field. Research then expanded to cases where agents have limited information [25,26], including on graph structures [27,28]. There was a particular focus on “information cascades” or “herd mentality”, where choices are made sequentially, both when there is a ground truth [29–31] and where there is none [32,33]. 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 [34] and theoretical [35,36,10] (see overviews

[37,38]), including attempts to find influential nodes in the social graph [39]. Computational models of opinion diffusion have looked at the fixed-point properties of the graph dynamics, in connection with consensus formation [17] and its complexity [18]. An important stream of research has looked at how to control opinion diffusion by external intervention, for example through bribery [13], false-name attacks [40] or information control [11], and we see our results as introducing effective heuristics for outcome prediction in those frameworks.

Paper structure In Section 2 we introduce our setup and the basic graph-theoretic terminology, in particular the model of caveman-graphs and some basic observations on the influence gap. In Section 3 we present our homophilic extension, together with the algorithms to control the homophily level and the rewiring probability. Section 4 focuses on two-party elections, introducing the opinion diffusion dynamics and measuring the predictive power of the influence gap and other key metrics. These results are further discussed in Subsection 5.2, where we combine metrics and establish the effects of homophily and rewiring. In Section 6 we delve into the multi-party case. We first consider potential extensions of the influence gap, studying them on simple graphs and then using computer-aided simulations and larger graphs in Section 7 to observe their predictive power in a setting with three parties. We conclude in Section 8 presenting various follow-up research directions.

2. Influence gap and communities

Consider an election with parties \mathcal{P} , where for each party $P \in \mathcal{P}$ there are $N_P > 0$ supporters (voters). The N total voters are placed on an undirected graph $\mathcal{G} = (V, E)$, where they are represented as nodes in the node set $V = \{1, 2, \dots, N\}$ while social connections are given by edges in the edge set $E \subseteq V \times V$.

Let $p : V \rightarrow \mathcal{P}$ be the party assignment, representing individuals' initial opinions, such that, for example, $p(n)$ is the party of voter n . We can also apply this to a subset of nodes $U \subseteq V$, and denote by $p(U)$ the parties that all nodes in the subset U vote for, i.e., $p(U) = \{p(n) \mid n \in U\}$. We denote the neighbourhood of $n \in V$, i.e., the set of social connections of voter n , by $\mathcal{N}_n = \{m \mid (n, m) \in E\}$. A poll of a node n includes n itself and its neighbourhood, $\mathcal{N}_n^n = n \cup \mathcal{N}_n$. For a node $n \in V$, denote by Δ_n^P be the fraction of n 's poll that votes for party P and let $\Delta_n = (\Delta_n^P \mid P \in \mathcal{P})$ be the vector of such polls; for brevity we shall often use Δ_n instead of $\Delta_n^{p(n)}$ when no ambiguity arises.

2.1. Influence gap

Like much of the research on opinion dynamics (including, in particular, Stewart et al. [15]), we shall, at least at first, focus on the two-party (or two-opinion) setting and we shall henceforth refer to the parties using two colours, red and blue. In other words, we fix a partisan structure $\mathcal{P} = \{R, B\}$ where the initial number of voters for the red (blue) party is N_R (N_B), noticing that with only two parties it suffices to consider only the fraction of a node's poll that votes for their own party, i.e., Δ_n . The multiparty generalisation ($|\mathcal{P}| \geq 2$) has non-trivial ramifications for a number of concepts, and will be investigated in detail from Section 6 onwards.

In the two-party case, *influence assortment* [15] is, intuitively, the relative advantage of a party against their rival, 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 P denoted by A_P . The *influence gap* (IG), which we denote by G_P , is the average advantage in influence assortment of party $P \in \{R, B\}$ (resp., $G_{P'}$ denotes its dual for party $P' \neq P$). Below are the formal definitions (as per Stewart et al. [15]); note the use of the Kronecker delta $\delta(i, j)$, which is 1 if $i = j$ and 0 otherwise.

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

$$A_P = \frac{1}{N_P} \sum_{n \in V} a_n \delta(p(n), P) \quad (2)$$

$$G_P = A_P - A_{P'} \quad (3)$$

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 their own poll, and thus how much they can be influenced to vote for 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, A_P . The influence gap, G_P can therefore be understood as the difference in how "strategically placed" a party is – how much do its supporters interact with other parties (and thus, are able to be influenced by them).

Throughout the paper we focus on the case of a strong party assignment (SPA), following Alon et al. [10], where each party is assigned a fixed fraction of nodes – initially this will be a half but in later sections we consider non-equal representation. 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.

2.2. Caveman graphs

A *caveman graph* $\mathcal{G} = (V, E)$ is a set of l isolated cliques each of size k [23]. These graphs encode a very basic form of community-structure without showing interesting variety or empirical relevance [41],¹ but we shall use them for the most basic theoretical insights and will build on them to develop more realistic-looking structures.

2.2.1. Influence gap in caveman graphs

In caveman graphs the influence gap shows some interesting features in connection with the partisan majorities in each clique. To see why this is the case, consider a graph of l cliques (labelled $1, \dots, l$), each of size k . Let us take, without loss of generality, the perspective of the red party first. The fraction of a poll of a red node n in clique $c \in \{1, \dots, l\}$ – for the total number of red nodes in clique c , x_c – is exactly $\Delta_n = \frac{x_c}{k}$ red nodes, as, by definition, they are connected to all of the nodes inside their clique. This holds for all red nodes inside c , of course. Thus, the influence assortment of any red node in c is the same, so that $a_n = a_m$ for all red 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 $\exists M, M' \in \{1, \dots, l\}$ 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 *marginal* cliques. 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, A_R .

$$\begin{aligned} A_R &= \frac{1}{N_R} \sum_{n \in V} a_n \delta(p(n), R) \\ &= \frac{1}{N_R} \left(\sum_{c=1}^{M'} \left(x_c \frac{x_c}{k} \right) + \sum_{c=M'+1}^l x_c \left(\frac{x_c}{k} - 1 \right) \right) \\ &= \frac{1}{N_R} \left(\sum_{c=1}^l \frac{x_c^2}{k} - (N_R - X_{M'}) \right) \\ A_R &= \frac{1}{N_R} \left(\sum_{c=1}^l \frac{x_c^2}{k} + X_{M'} \right) - 1 \end{aligned} \tag{4}$$

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, A_B , in terms of red counts x_n is thus as follows.

$$A_B = \frac{1}{N_B} \left(\sum_{c=1}^l \frac{x_c^2}{k} + X_M + N - 2N_R - Mk \right) \tag{5}$$

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

$$G_R = \frac{X_{M'}}{N_R} - \frac{X_M}{N_B} + \frac{N_B - N_R + Mk}{N_B} - 1 + \frac{N_B - N_R}{N_R N_B} \sum_{c=1}^l \frac{x_c^2}{k} \tag{6}$$

¹ A connected version is formed by rewiring a single edge per clique to a node in an adjacent clique along a central cycle [23], also known as a caveman graph. For the purposes of this paper a caveman graph will be taken to mean the unconnected set of isolated cliques.




Fig. 1. A caveman graph with equal representation and IG of $G_R = -1/3$. Notice that red has one strict majority while blue has two, i.e., $M' = M = 1$. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)




Fig. 2. A caveman graph with $G_R = G_B = 0$ and different majorities, since $X_M = X_{M'} = x_1 = 4$. Notice how the IG does not “see” a win for the blue party.

Influence gap with equal representation We now show, using the expressions above, what happens in caveman graphs where parties have equal representation, i.e., $N_B = N_R = n/2$. When this is the case a number of terms in Equation (6) vanish, including the quadratic term, leaving only a linear equation in purely community-level quantities. In particular, note that the term $X_{M'} - X_M = (M' - M) \cdot k/2$ since all the cliques between M and M' are marginals. After some simplification this gives us an equation for G_R in terms of the number of cliques that are (strictly) dominated by red – as opposed to individual node quantities.

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

A corollary of this is when $G_R = G_B = 0$, when by the above equation caveman graphs with equal representation must also have as many strict red majorities as strict blue ones.

Influence gap and initial majority We now turn our attention on how predictive the influence gap is in these structures. We start with the restricted class of caveman graphs with equal representation, which we analysed above.

To do so, we first consider a simple opinion diffusion model, e.g., at each time step voters stick to their own opinion unless most of their friends think differently, in which case they flip with probability Δ_n . In Fig. 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. We note that in this case IG seems to be a good predictor of the overall outcome.

However, when equal representation is relaxed, IG can also fail to predict robust configurations with clear winners while majority does, as illustrated in Fig. 2.

2.2.2. Towards more realistic communities

We want to establish how good of a predictor the IG is in graphs with community structure. While we observed interesting edge cases in the simple caveman graphs, it is now important to extend these to more capture more realistic ones.

We do so by generalising them in two ways:

- Allowing some edges to be rewired with a set probability, i.e., relaxed-caveman graphs [42].
- Allowing such probability to depend on homophily, i.e., how likely like-minded voters are to be connected to one another.

A *relaxed-caveman* graph is a modified version of the basic caveman graph, whereby edges are rewired with some given probability [42]. Concretely, given a probability p_0 and iterating over all edges E of a set of isolated cliques, an edge $(u, v) \in E$ is rewired as (u, n) , for some randomly selected $n \in V$; 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 [42–44] and generation [45,46] methods.

It is important to note that, at low rewire probabilities, the results above on caveman graphs can be extended to relaxed-caveman graphs (and their homophilic variant, see below) by considering the effects of a small number of rewired edges as perturbations $\mathcal{O}(\min(N_R, N_B)^{-1})$, 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 $\mathcal{O}(N_R^{-1})$ or $\mathcal{O}(N_B^{-1})$.

The next section will present the model of homophilic relaxed-caveman graphs, where rewiring is linked with homophily, and use it to compare the predictive power of a number of metrics, including the influence gap.

3. Homophilic relaxed-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, unlike real-world networks, where the two are highly intertwined [22,21]. 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 [9]. This allows us to generate synthetic graphs with communities where the graph structure is dependent on the party assignment, following the observed behaviour [22] that people tend to cluster with people who share their views.

Algorithm: Homophilic Relaxed-Caveman Graph, \mathcal{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** \mathcal{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 $p : V \rightarrow \mathcal{P}$. The probability to rewire p_0 can be thought of as the likelihood of changing a pre-existing friendship to a new one, 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 their party as they are 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. Community detection methods [42,43], such as modularity-based methods [47], may be used to reestablish communities in this case.

We now show how the influence gap is distributed in homophilic relaxed-caveman graphs with equal partisan split ($N_R = N_B$). 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. We plotted the results in Fig. 3.

We find that for values of $h < 0.85$, the absolute value of the influence gap increases monotonically with the rewire probability p_0 . That is, taking a cross-section for some $h < 0.85$ of Fig. 3 is monotonic in p_0 . At higher homophily we find the curves are peaked at around $p_0 \approx 0.7$.

We see a more complex phenomenon when taking cross-sections in p_0 , that is, curves change from nearly flat at low p_0 to being peaked at higher p_0 . 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 the 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 their own preference is enough to push the poll into a (weak) majority $\Delta_n \geq 0.5$ then the maximal value of the influence gap is reached, when most/all nodes face this situation.

4. Predicting two-party elections: the model

In this section we examine whether the findings of Stewart et al. [15] still hold for two-party elections when communities are present. Our conclusion, in summary, is that influence gap does not predict influence dynamics when initial votes are equally split, which is a central assumption in their analysis. When this assumption is relaxed, the correlation is




Fig. 3. A surface plot of the mean absolute value of the influence gap, across different homophily factors h and a range of rewiring probabilities p_0 . Each point is measured from 10^4 party assignments each generating a single graph.

restored, but a more predictive and computationally simpler metric exists, namely counting the initial majority of either party. This holds across all levels of homophily and rewiring, with some interesting specificities.

Note that here, we measure predictive power of a metric by the correlation over multiple simulations between the metric value evaluated at the beginning of the election and the final outcome. In this way a metric is ‘good’ at prediction only in so far as there is statistical correlation between effectively an initially measured heuristic and the actual result of the election, and not necessarily if there is a causal link.

Next, we provide the details of our experiment and results, which uses the same behavioural parameters of Stewart et al. [15], starting from the voter model.

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, in line with Stewart et al. [15]. A voter’s knowledge is restricted by their social network, only knowing the voting intentions of their neighbours as well as their own, serving as a form of poll, to which they wish to conform. The game lasts for a fixed amount of time, during which players can change their voting intentions synchronously.² The winning party is the one to hold a majority above a threshold $V > 0.5$ when the updating process is done, otherwise we consider it a deadlock.

The agents follow the stochastic behavioural model developed by Stewart et al. [15], informed by a social experiment with human subjects who were given pay-offs depending on the success of the party they voted for at the end of the game (in other words the actual they vote for not the one they initially were assigned to). At any given time, a voter, according to the behavioural model, would vote for their 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*.

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 Stewart et al. [15] when replicating their findings with simulations.

4.2. Benchmark metrics

Since an exact analysis of the opinion dynamics faces significant complexity barriers, Stewart et al. [15] proposed using the Influence Gap as a prediction tool of the dynamic’s outcome. To the IG proposal we add three more metrics, which we will compare against:

² Unlike the iterative voting model of Tsang and Larson [9], in which every voter updates separately to the others, this is more akin to the synchronous updates of Alon et al. [10].

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. [15] 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




Fig. 4. Example time series of a voting game, using Stewart et al. [15] behavioural model on a homophilic relaxed-caveman graph with rewiring 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.

Majority The initial majority of the red party.

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 classical political science metric [48], developed to measure gerrymandering in two-party elections, in which we examine how many votes were “wasted”, i.e., could have been eliminated without changing the outcome.

4.3. Experimental setup

We have $N = 20$ voters, each assigned a party and then are placed in a graph G , generated by a number of different assignments of h and p_0 . Following the setting of Stewart et al. [15] (which was also based on human-subject data as well as simulations), each simulation game runs for a total of 240 seconds in which the first 83 seconds are considered to be the early phase after which the game transitions to the late phase. Every 3.3 seconds³ a voter v , can update their voting intention, with probability p_{ij}^v , which are sampled from the empirical parameter distributions found from the human social experiment [15]. In total, for a single simulation N samples are taken from 6 distributions each (for each stage – early/late – and each voting outcome – Red victory/Blue victory/Deadlock). After 240 seconds have elapsed the vote share across the entire graph is measured.

An example of a time series produced by a simulation is shown in Fig. 4. The different sections of the plot, partitioned by dashed and dotted lines, represent different strategies. For example, in the early phase ($t < t_*$) both parties are deadlocked and thus agents vote for their assigned party with probability $p_{deadlock,early}^{agent}$. The convergence of a time series is not guaranteed because, as in the original social experiment, the game finishes after 240 seconds.

³ Note that we report time in seconds to be physically interpretable in relation to the original experiments. In our implementation we bound time by time-steps, not wall clock or CPU time, where each time-step corresponds to 3.3 s: the simulation runs for 72 time-steps and the early phase consists of 25 time-steps.




Fig. 5. Simulations of the behavioural model of Stewart et al. [15] acting on relaxed- and homophilic relaxed-caveman graph, starting with equal representation. In both cases the influence gap correlates very weakly with voter skew, having a Pearson’s $\rho < 0.3$, and passing a significance test with $p < 10^{-7}$. **Left:** relaxed-caveman graphs with rewiring parameter $p = 0.3$. **Right:** homophilic relaxed-caveman with rewiring probability $p_0 = 1$ and homophily factor $h = 0.3$.




Fig. 6. For a hRC with $p_0 = 0.4$ and $h = 0.6$, we plot the final vote skew against the IG and the majority of the red party.

We simulate over 10^4 elections for each set of hRC parameter values (p_0, h) while varying the initial number of red nodes, $N_R \geq 10$. Parameters were chosen to cover a reasonable range, $p_0 \in \{0, 0.4, 1\}$ and $h \in \{0, 0.1, \dots, 1\}$, so that in total there were over 33,000 simulated elections. At the start of each election we measure the influence gap, the deterministic voter skew and the efficiency gap (see Section 4.2); the final outcome of the election is also recorded as a vote skew towards the red party $N_R/N - 1/2$. For example, Fig. 6 shows the elections occurring on hRCs with $(p_0, h) = (0.4, 0.6)$ and illustrates how two metrics, the IG and majority, both correlate strongly with the final outcome (note that the figure shows a case where parties did not have equal power initially).

5. Predicting two-party elections: results

Where Stewart et al. [15] 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 there. In both the relaxed-caveman and the homophilic relaxed-caveman with given parameter sets, the Pearson correlation coefficient is small – $\rho < 0.3$ – as seen in Fig. 5.

When we allow for unequal initial setup, results show more complexity. In Fig. 6 we plot as an example how the initial influence gap and majority compare to the final election outcome, for a hRC graph with intermediate rewiring $p_0 = 0.4$ and intermediate homophily $h = 0.6$. We see in both cases very strong correlations with a high Pearson correlation coefficient (PCC) $\rho > 0.9$. Note that by symmetry near identical distributions would be found, if we had simulated $N_R < 10$.

In other words, for this set of parameter values, both IG and majority are strong predictors of the final outcome – in so far as their initial values correlate strongly with the result. However we can begin to see how the influence gap does not predict the outcome as accurately as the initial majority.

5.1. The effects of rewiring and homophily

We now present our findings taking rewiring and homophily into account, a snapshot of which is shown in Fig. 7.

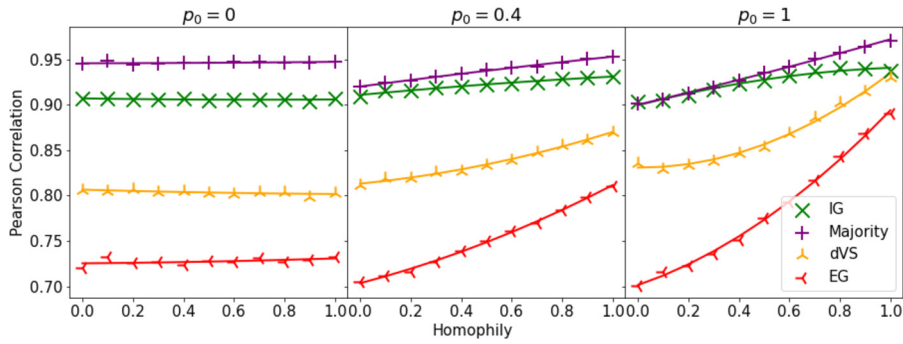


Fig. 7. As model parameters of the homophilic relaxed-caveman are varied independently, the Pearson correlation coefficient (PCC) between the final voter skew and the majority (purple), the influence gap (IG, green), the deterministic voter skew (dVS, orange) and the efficiency gap (EG, red) are plotted. Quadratic lines of best fit for multiple rewiring probabilities are plotted, on the left $p_0 = 0$, in the middle $p_0 = 0.4$ and on the right $p_0 = 1$. The influence gap only outperforms the majority once, at $(p_0, h) = (1, 0)$.

Before delving into discussing them, it is important to note that neither homophily nor rewiring correlate, on their own, with the final outcome of the election, with Pearson correlation coefficients of $\rho = -0.0256$ and $\rho = 0.0106$ respectively. This does not imply that neither have an effect at all – indeed, Fig. 7 shows otherwise – simply that they do not impact the election single-handedly.

High homophily When $h > 0.5$, hRC graphs are rife with echo chambers in which an agent has friends mostly of the same party and opinion as themselves. Their polls are therefore fairly homogeneous and they see no compelling reason to change their vote, meaning the graph structure plays little to no role. Extending this intuition to all agents in the graph, very little diffusion of opinions occurs and thus the final outcome of the election will not be too dissimilar to the starting one. As such, counting the number of votes at the start will closely resemble – and thus predict – the number of votes at the end, and majority here can better predict the outcome than IG. Moreover, this logic can help explain the general trend that the correlations of most metrics increase drastically at high homophily.

More formally, at high h the average poll will show a majority towards their 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 $i = \text{win}$, so that their strategy is $p_{\text{win},j}^v$. 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.

Low homophily Conversely, for $h < 0.5$, an agent’s poll, at least initially, likely shows a super-minority $\Delta_v < 1 - V$ or equivalently $i = \text{lose}$ such that over 40% of agents will change their votes. In other words, most polls are quite diverse and an agent will see mostly nodes of the opposite party and “doubt” their opinion more often.

In this scenario the dynamics become less stable, and more complex diffusion occurs. Therefore, most metrics perform either as well as or worse at $p_0 = 1$ than $p_0 = 0$. Only the deterministic voter skew sees an increase in predictive power as p_0 increases, at all values of homophily. We suspect this is because a single deterministic step causes the network to evolve towards a more stable configuration and, as such, closer to the final outcome.

Another effect of this is that the predictive powers of both majority and influence gap monotonically both decrease as homophily decreases; in fact for $h < 0.5$ the difference in PCC between IG and majority shrinks considerably.

As we move up the rewiring level, the majority PCC is linear while the IG PCC is a flat quadratic. Both the efficiency gap and deterministic voter skew, however, show a move from linear to quadratic as communities become more diluted. Moreover, both metrics are consistently outperformed by influence gap, though still have reasonably high correlations; once again majority bests all others.

Overall, only once does the influence gap predict outcomes better than the initial majority, at $p_0 = 1$ and $h = 0$, in other words, when nodes are extremely diverse and seek out alternative opinions to their own. The difference, however, is tiny at $|\Delta\rho| = 0.00144$ – and likely due to the stochastic nature of the elections. We conclude that simply counting the number of party votes is a better predictor than the metric of Stewart et al. [15].

5.2. Regression models

In order to further explore the role of majority and influence gap (x_M and x_G respectively, following typical regression notation) as predictive or explanatory variables we use linear regression (Equation (8)) to build several models of the vote 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 and tested on the remaining 30%.

$$y = \beta_M x_M + \beta_G x_G + \beta_0 \tag{8}$$

Table 2

Coefficients of regression, β , and of determination, R^2 , for regression models of the dynamic voting game outcome on homophilic relaxed-caveman graphs for two-party elections. Note that adjusted- R^2 are the same as the R^2 values up to 3 significant figures while p -values for all non-zero coefficients are 0.000.

Metric	β_M	β_G	β_0	R^2
Majority	0.0507		0.0363	0.881
IG		0.226	0.00474	0.837
Both	0.0333	0.0869	0.0144	0.902

The regression confirms our observations, that majority is a better predictive tool than influence gap. This does not, however, render IG useless. In particular the joint model outperforms both individual models that either use majority or influence gap exclusively as shown in Table 2, despite some colinearity between features. In other words, IG is an informative metric that can build upon the simple predictions of majority, but alone is not as effective.

6. Extending the influence gap to multiple parties

In many political systems more than two parties are competitive and relevant. While the US has two main parties, Canada has 5 parties with more than 5% in parliament; the UK has 3 such parties, and the situation is similar in most Continental European democracies and around the world. In this section we will discuss the challenges in expanding the definitions for IG to the multi-party setting, by showing a few possible definitions for the influence assortment and gap. We then apply these definitions to small graphs to highlight the subtle differences and to motivate the choice of one definition over the other.

6.1. Influence assortment

We examine two possible extensions to multi-party settings. Both assume, naturally, that plurality is the reasonable extension to majority (which always exists in two-party settings, but not guaranteed to exist with more parties). The first assumes the “force” of the plurality winner is all there is, and thus even when an agent does not support it, the negative weight it carries is only as much as the plurality winner could muster.

$$a_n = \begin{cases} \Delta_n^{p(n)} & \text{if } \Delta_n^{p(n)} = \max(\Delta_n), \\ -\max(\Delta_n) & \text{otherwise.} \end{cases} \tag{9}$$

An alternative definition takes an all-against-one stance: rather than taking the most dominant party, we sum over all other parties. That is, when an agent does not support the plurality winner, their negative weight is the sum of all the parties one did not support.

$$a_n = \begin{cases} \Delta_n^{p(n)} & \text{if } \Delta_n^{p(n)} = \max(\Delta_n), \\ -(1 - \Delta_n^{p(n)}) & \text{otherwise.} \end{cases} \tag{10}$$

Both definitions reduce to the original two-party definition of influence assortment and have reasonable intuition behind them.

6.2. Influence gap

There are two main ways to define the influence gap. First, to define it as the difference between influence assortments of party P and the (next-) most influential party. Second, as the difference between party P and the (next-) plurality party (i.e., plurality runner-up).

$$G_P = A_P - \max_{Q \neq P} A_Q \tag{11}$$

$$G_P = A_P - A_Q \quad \text{where } Q = \arg \max_{P' \neq P} N_{P'} \tag{12}$$

These two gaps largely correlate with one another, since influence and partisan split are well correlated – though not without exceptions. Specifically, the latter definition can produce non-unique values for the same network.

To more carefully explore our suggested extensions, we will now examine several settings of multi-party situations, and for that we shall explore the case of three parties: $\mathcal{P} = \{\circ, \diamond, \square\}$.

Table 3
Comparison of definitions, applied on $N = 3$ graphs.

Graph	a_n	G_P	G_\circ	G_\diamond	G_\square
Cycle	(*)	(*)	0	0	0
$\circ - \diamond - \square$	(*)	(11)	0	$-1/6$	0
		(12)	$0, 1/6$	$-1/6$	$0, 1/6$

Table 4
Comparison of definitions, applied on lines of size $N = 4$.

Line	a_n	G_\circ	G_\diamond	G_\square
$\circ - \circ - \diamond - \square$	(*)	$1/3$	$-1/2$	$-1/3$
$\circ - \diamond - \square - \circ$	(*)	$1/6$	$-1/6$	$-1/6$
$\diamond - \circ - \square - \circ$	(*)	$-1/12$	$1/12$	$-13/12$
$\diamond - \circ - \circ - \square$	(*)	$1/6$	$-1/6$	$-1/6$

Table 5
Comparison of definitions, applied on stars of size $N = 4$, with different centres.

Centre	a_n	G_\circ	G_\diamond	G_\square
\diamond	(9)	0	-1	0
	(10)	0	$-5/4$	0
\circ	(*)	$1/4$	$-1/4$	$-1/4$

6.3. $N = 3$ examples

Consider three voters who each vote for a different party. There are only two connected networks to consider: the cycle and the line. Due to the symmetry of the cycle, all influence assortments are equal and thus all influence gaps are 0. For the line where \diamond is the middle node, there is symmetry between the \circ and the \square (Table 3).

For the line graph, notice that the definition of assortment leads to \diamond being comparatively heavily influenced (Equation (10)), or only slightly influenced (Equation (9)). Moreover, the definition of the influence gap can lead to multiple results for the same graph – this is clearly problematic. In general the latter definition of IG, Equation (12), produces non-unique results if there are joint winners – this can occur frequently and as such we hence only consider the first definition, Equation (11).

6.3.1. $N = 4$ examples

Consider a network of 4 voters, two of which vote for the \circ party, one votes \diamond and another votes \square . Already with this set-up there are at least a dozen configurations to create a connected network; here we consider a few of these to illustrate the simple examples that complicate choosing a definition, starting with the 4 non-trivial line graphs (since swapping \square and \diamond produces a qualitatively identical result) (Table 4).

We see that for such line graphs, both assortment definitions coincide. In general, for any sized line graph (or regular graph of degree 2), definitions (9) and (10) always coincide. For any number of parties in \mathcal{P} , a node either has one neighbour or two. When n has exactly one neighbour – in other words, n is at the end of the line – $\Delta_n^{p(n)} = 1/2$ or 1, so that $\Delta_n^{p(n)} = \max(\Delta_n)$ hence both definitions coincide. In the bulk of the line, with two neighbours, there are three cases:

- $p(n) = p(n - 1)$ or $p(n + 1)$; $p(n)$ has plurality in the poll of n , therefore (9) = (10) trivially.
- $p(n - 1) = p(n + 1) \neq p(n)$; $p(n - 1)$ holds plurality but is the only alternative party to $p(n)$, hence $\max(\Delta_n) = 1 - \Delta_n^{p(n)}$ ($= 2/3$).
- $p(n - 1) = p(n) = p(n + 1)$; all three parties locally hold joint plurality, therefore (9) = (10) trivially.

Using similar analysis, ring graphs (with degree 2) also cause the two definitions to coincide (Table 5).

On the other hand as soon as a node has degree 3 or higher with at least 3 parties we start to see non-unique solutions. For example, examine star graphs (i.e., graphs in which one node is connected to all others, and no other edges exist) when there are two \circ 's, one \diamond and one \square . When the \diamond is at the centre, \circ holds plurality but is not the only alternative party. Hence Definition (10) will produce a more negative assortment and thus a more extreme influence gap.

More complex edge configurations exist, but we look at two more to further illustrate the problem: a clique and a near-clique – i.e., a clique where a single edge ($\circ - \diamond$) has been removed (Table 6).

Table 6
Comparison of definitions, applied on near-cliques of size $N = 4$, with different centres.

Graph	a_n	G_o	G_\circ	G_\square
Near-Clique	(9)	$5/6$	$-5/6$	$-5/3$
	(10)	$5/6$	$-5/6$	$-23/12$
Clique	(9)	1	-1	-1
	(10)	$3/4$	$-5/4$	$-5/4$

Henceforth we consider only Definition (9) for influence assortment and Definition (11) for influence gap, due to their faithfulness to the original concepts and, in the case of IG, that they are well-defined, i.e., do not produce multiple results on the same graph.

7. Influence gap on large graphs with multiple parties

In this section, we use the previously identified extensions of influence gap and assortment, to handle larger social networks with multiple parties. We first provide an analysis of their behaviour on cliques, for any number of parties, and then we show experimental evidence of their predictive power on hRC graphs.

7.1. Cliques and plurality centres

Consider a clique of size N with a set of parties $\mathcal{P} = \{0, \dots, |\mathcal{P}| - 1\}$ – in other words, every voter has perfect information on all other voters – and let N_P be the number of voters for party P . In this situation *influence* perfectly correlates with voter split, i.e., the fraction of the population. There are two main cases to consider: 1) a single plurality winner, 2) multiple joint winners (a “deadlock” in the parlance of Stewart et al. [15]).

In case 1, let us denote by $P = 0$ the singular winner and note that $\max(\Delta_n) = \frac{N_0}{N}, \forall n \in V$ since it is a clique. Moreover, for all voters of a single party P , the influence assortment is the same, due to the symmetry of the system, and hence the influence assortment at the party level is equal to the assortment at the node level. From Definition (9) we can conclude that $A_0 = \frac{N_0}{N}$ and $A_P = -\frac{N_0}{N}, \forall P \in \mathcal{P} \setminus 0$. As such, the influence gap for the plurality winner is $+2\frac{N_0}{N}$, and for the losers $-2\frac{N_0}{N}$.

In case 2, let W be the number of jointly winning parties $\mathcal{W} = \{0, \dots, W - 1\}$, with N_0 number of voters. In this case for $P < W, A_P = \frac{N_0}{N}$ with the remainder as $A_P = -\frac{N_0}{N}$ for $P \geq W$; the influence gap follows simply as $+2\frac{N_0}{N}$ for the winners and $-2\frac{N_0}{N}$ for the losers.

Combining the two cases, we find the influence gap for a clique of size N with W (joint) winners ($P < W$), who have won with N_0 number of voters each. The magnitude of the gap is the same for all parties, but whose sign is determined by being a winning party or not.

$$G_P = \begin{cases} +\frac{2N_0}{N} & \text{for } P < W, \\ -\frac{2N_0}{N} & \text{otherwise} \end{cases} \tag{13}$$

From a clique we can analytically bound the influence gap for a broader class of networks that exhibit some core-peripheral structure by removing edges systematically from the clique. Consider removing an edge between two nodes u and v , neither of which vote for a plurality party, P_0 . In doing so, in the polls of u and v , P_0 (and all other plurality winners) still holds plurality and even has a higher poll fraction $\Delta_u^{P_0} = \Delta_v^{P_0} = \frac{N_0}{(N-1)}$. As such a_u and a_v decrease, so that overall G_{P_0} increases while $G_{p(u)}$ and $G_{p(v)}$ decreases.

Repeating this procedure one can keep removing edges until plurality winners form a centre or core (of size WN_0) – in other words, all voters of $P < W$ are connected to all other nodes – while those in the periphery each have WN_0 neighbours. Thus all nodes have polls that show plurality towards $P < W$ with fraction $\Delta_n^{P < W} = \frac{1}{W + \frac{1}{N_0}} \approx \frac{1}{W}$ such that the influence gap is simple.

$$G_P = \begin{cases} +\frac{2N_0}{WN_0+1} & \text{for } P < W, \\ -\frac{2N_0}{WN_0+1} & \text{else.} \end{cases} \tag{14}$$

Any network between this plurality-core and the clique has an influence gap bounded by Equations (13) and (14): $\frac{2N_0}{WN_0+1} \leq |G_P| \leq \frac{2N_0}{N}$, with the same dependency on winning for the sign.

7.2. Predicting three-party elections

Equipped with the formal multiparty definition for *influence assortment* (Equation (9)) and *influence gap* (Equation (11)), we can replicate the methodology of Section 4 in order to simulate elections with 3 parties – red, blue and green. The voter




Fig. 8. For a three-party election occurring on hRC graphs, the Pearson correlation coefficient (PCC) between the final vote skew and the initial number of party voters (pluses) and influence gap (crosses) are plotted for each party colour (red, blue or green). As before for different values of the rewiring – $p_0 = 0$ on the left, $p_0 = 0.4$ in the middle and $p_0 = 1$ on the right – the homophily is varied. In black are lines of best fit for each metric, across all parties. For visual clarity we leave out the deterministic voter skew and efficiency gap, the latter being canonically used for the American two-party system hence unsuitable for the multi-party case.

model easily extends to 3 or more parties; with probability p_{ij}^v a voter v in state i in phase j of the election sticks to her current vote, and with probability $1 - p_{ij}^v$ she votes for the party with highest poll fraction (excluding her own). Similarly the algorithm to generate the homophilic relaxed-caveman graph works for any arbitrarily sized set of parties \mathcal{P} .

For the benchmark metrics, we focus only on party votes – counting the number of voters for each party. For simplicity, we omit the analysis of the simplified dVS dynamic and the two-party intended EG.

Looking at Fig. 8, there is a very clear and consistent rift between the influence gap and the party votes – IG is always the weaker predictor. Moreover, while in the two-party case the difference between PCCs can be as small as $|\Delta\rho| \geq 0.00141$, the smallest difference for the three-party case is $|\Delta\rho| \geq 0.0191$. In other words, in the three-party case the influence gap never gets arbitrarily close in its predictive accuracy to the simple vote count; there seems to be a bound on how close the IG will get.

Another curious effect in moving to the three-party case is the behaviour with high rewiring probability. In this case, the PCC for influence gap becomes noticeably worse for low homophily, causing a widening divide between it and the party votes, in direct contrast to the two-party case where the difference shrinks.

To conclude, note that for more than three parties and random initial assignment and rewiring, the share of any winning party in any neighbourhood is smaller in expectation. Therefore, the magnitude of the assortment values a_n (for any node $n \in V$) is smaller, and so is the influence gap, leading it to be an even noisier predictor.

7.3. Multiparty regression

We repeat the regression testing, as in Section 5.2, with the same methodology: 70% of the data is used to train with the remainder 30% used to test and validate. In the multiparty case note that no longer do we have a single x_G , x_M nor a single y but generally there are $|\mathcal{P}|$ for each variable. In other words we now have x_G^P and x_M^P for $P \in \mathcal{P}$. This introduces high inter-party collinearity with high variance inflation factors ($VIF > 11$) particularly for the voter count, as voter numbers are constrained by the total number of voters, i.e., $\sum_{P \in \mathcal{P}} x_M^P = N$. As such, for a target variable y^P that is the final voter count for party P , we only use the initial voter count for party P .

$$y^P = \beta_M^P x_M^P + \sum_{P' \in \mathcal{P}} \beta_G^{P'} x_G^{P'} + \beta_0 \tag{15}$$

The multi-party regression qualitatively reproduces the results for the two-party case: initial voter count is a better predictor than IG. As before, however, the inclusion of IG as an additional feature improves this accuracy (Table 7).

8. Conclusion

We studied the predictive power of a recently proposed metric, the influence gap, to compute the results of many-party elections on networks with community structure. To do so, we proposed a novel model, the homophilic relaxed-caveman, as a means to generate synthetic graphs with communities that may exhibit echo chambers. In order to extend our analysis to the multi-party case we then proposed multiple definitions for the influence assortment and gap. Their similarities and subtle differences were illustrated first in small simple graphs and then characterised for large graphs – such as cliques and plurality centres.

We showed that the presence of communities suppresses the power of the influence gap as a predictor of the final outcome. A much simpler metric – the initial majority – is a far better predictor, albeit one which can be improved somewhat

Table 7

Coefficients of regression, β , and of determination, R^2 , for regression models of the dynamic voting game outcome on homophilic relaxed-caveman graphs for three-party elections. Note that adjusted- R^2 are the same as the R^2 values up to 3 significant figures while p -values for all non-zero coefficients are 0.000.

Metric	Party	β_M^P	β_G^{Red}	β_G^{Blue}	β_G^{Green}	β_0	R^2
Voter count	Red	1.18				-1.23	0.938
	Blue	1.18				-1.22	0.937
	Green	1.18				-1.23	0.937
IG	Red		3.43	-1.71	-1.71	6.67	0.900
	Blue		-1.72	3.42	-1.72	6.66	0.901
	Green		-1.71	-1.71	3.44	6.67	0.901
Both	Red	0.798	1.21	-0.614	-0.612	1.34	0.951
	Blue	0.799	-0.610	1.21	-0.610	1.34	0.951
	Green	0.800	-0.595	-0.592	1.22	1.34	0.951

by combining it with the influence gap for an even better predictor. Surprisingly, the initial advantage is an even stronger metric in the multi-party cases.

Having measured the efficacy of several metrics, this poses the question of whether a metric exists that can predict the voting dynamics most accurately, while still being easy to compute.

A further important 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.

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. [15] did, to influence members behaviour. Just like the forceful agents in Acemoglu et al. [5], 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 (in other words edges can dynamically change over time), 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?

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work was 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] L. Jasný, J. Waggle, D.R. Fisher, An empirical examination of echo chambers in us climate policy networks, *Nat. Clim. Change* 5 (2015) 782–786.
 [2] K. Garimella, G. De Francisci Morales, A. Gionis, M. Mathioudakis, Political discourse on social media: echo chambers, gatekeepers, and the price of bipartisanship, in: *Proceedings of the 2018 World Wide Web Conference, WWW '18*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2018, pp. 913–922.
 [3] C. Shao, G.L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, F. Menczer, The spread of low-credibility content by social bots, *Nat. Commun.* 9 (2018) 4787.
 [4] A. Borodin, O. Lev, N. Shah, T. Strangway, Big city vs. the great outdoors: voter distribution and how it affects gerrymandering, in: J. Lang (Ed.), *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, 2018, pp. 98–104.
 [5] D. Acemoglu, A. Ozdaglar, A. ParandehGheibi, Spread of (mis)information in social networks, *Games Econ. Behav.* 70 (2010) 194–227.
 [6] B. Wilder, Y. Vorobeychik, Controlling elections through social influence, in: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, Richland, SC, 2018, pp. 265–273.
 [7] F. Corò, E. Cruciani, G. D'Angelo, S. Ponziani, 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, 2019, pp. 201–207.
 [8] M. Castiglioni, D. Ferraioli, N. Gatti, Election control in social networks via edge addition or removal, in: *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (2020) 1878–1885.
 [9] A. Tsang, K. Larson, 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), AAMAS '16*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2016, pp. 368–375.

- [10] N. Alon, M. Feldman, O. Lev, M. Tennenholtz, How robust is the wisdom of the crowds?, in: Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI), IJCAI'15, AAAI Press, 2015, pp. 2055–2061.
- [11] D. Baumeister, A.-K. Selker, A. Wilczynski, 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, 2020, pp. 132–140.
- [12] U. Grandi, P. Turrini, A network-based rating system and its resistance to bribery, in: S. Kambhampati (Ed.), Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9–15 July 2016, IJCAI/AAAI Press, 2016, pp. 301–307.
- [13] R. Bredereck, E. Elkind, Manipulating opinion diffusion in social networks, in: C. Sierra (Ed.), Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19–25, 2017, ijcai.org, 2017, pp. 894–900.
- [14] U. Grandi, Social choice and social networks, in: U. Endriss (Ed.), Trends in Computational Social Choice, AI Access, 2017.
- [15] A.J. Stewart, M. Mosleh, M. Diakonova, A.A. Arechar, D.G. Rand, J.B. Plotkin, Information gerrymandering and undemocratic decisions, *Nature* 573 (2019) 117–121.
- [16] Z. Christoff, D. Grossi, Stability in binary opinion diffusion, in: LORI, in: Lecture Notes in Computer Science, vol. 10455, Springer, 2017, pp. 166–180.
- [17] V. Auletta, D. Ferraioli, G. Greco, On the complexity of reasoning about opinion diffusion under majority dynamics, *Artif. Intell.* 284 (2020) 103288.
- [18] D. Chistikov, G. Lisowski, M. Paterson, P. Turrini, Convergence of opinion diffusion is PSPACE-complete, in: Proceedings of the AAAI Conference on Artificial Intelligence 34 (2020) 7103–7110.
- [19] M. Girvan, M.E.J. Newman, Community structure in social and biological networks, *Proc. Natl. Acad. Sci.* 99 (2002) 7821–7826.
- [20] G.W. Flake, S. Lawrence, C.L. Giles, F.M. Coetzee, Self-organization and identification of web communities, *Computer* 35 (2002) 66–70.
- [21] E. Bakshy, S. Messing, L.A. Adamic, Exposure to ideologically diverse news and opinion on Facebook, *Science* 348 (2015) 1130–1132.
- [22] B. Bishop, *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*, Mariner Books, 2009.
- [23] D.J. Watts, Networks, dynamics, and the small-world phenomenon, *Am. J. Sociol.* 105 (1999) 493–527.
- [24] E.M. Rogers, *Diffusion of Innovations*, 5th ed., Simon and Schuster, 2003.
- [25] J. Farrell, G. Saloner, Standardization, compatibility, and innovation, *Rand J. Econ.* 16 (1985) 70–83.
- [26] M.L. Katz, C. Shapiro, Network externalities, competition, and compatibility, *Am. Econ. Rev.* 75 (1985) 424–440.
- [27] L.E. Blume, The statistical mechanics of strategic interaction, *Games Econ. Behav.* 5 (1993) 387–424.
- [28] D. López-Pintado, D.J. Watts, Social influence, binary decisions and collective dynamics, *Ration. Soc.* 20 (2008) 399–443.
- [29] S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades, *J. Polit. Econ.* 100 (1992) 992–1026.
- [30] E. Mossel, J. Neeman, O. Tamuz, Majority dynamics and aggregation of information in social networks, *Auton. Agents Multi-Agent Syst.* 28 (2014) 408–429.
- [31] M. Feldman, N. Immorlica, B. Lucier, S.M. Weinberg, 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), 2014.
- [32] N. Alon, M. Babaioff, R. Karidi, R. Lavi, M. Tennenholtz, Sequential voting with externalities: herding in social networks, in: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), Valencia, Spain, 2012, p. 36.
- [33] D.J. Watts, A simple model of global cascades on random networks, *Proc. Natl. Acad. Sci. USA* 99 (2002) 5766–5771.
- [34] J.J. Brown, P.H. Reingen, Social ties and word-of-mouth referral behavior, *J. Consum. Res.* 14 (1987) 350–362.
- [35] M.S. Granovetter, The strength of weak ties, *Am. J. Sociol.* 78 (1973) 136–1380.
- [36] J. Goldenberg, B. Libai, E. Muller, Talk of the network: a complex systems look at the underlying process of word-of-mouth, *Mark. Lett.* 12 (2001) 211–223.
- [37] V. Mahajan, E. Muller, F.M. Bass, New product diffusion models in marketing: a review and directions for research, *J. Mark.* 54 (1990) 1–26.
- [38] H.P. Young, Innovation diffusion in heterogeneous populations: contagion, social influence, and social learning, *Am. Econ. Rev.* 99 (2009) 1899–1924.
- [39] D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: International Conference on Knowledge Discovery and Data Mining (KDD), Washington D.C., 2003, pp. 137–146.
- [40] M. Brill, V. Conitzer, R. Freeman, N. Shah, False-name-proof recommendations in social networks, in: C.M. Jonker, S. Marsella, J. Thangarajah, K. Tuyls (Eds.), Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9–13, 2016, ACM, 2016, pp. 332–340.
- [41] W. Xiao, J. Ren, F. Qi, Z. Song, M.-X. Zhu, H.-F. Yang, H.-Y. Jin, B. Wang, T. Zhou, Empirical study on clique-degree distribution of networks, *Phys. Rev. E, Stat. Nonlinear Soft Matter Phys.* 76 (2007) 037102.
- [42] S. Fortunato, Community detection in graphs, *Phys. Rep.* 486 (2010) 75–174.
- [43] M.T. Schaub, J.-C. Delvenne, M. Rosvall, R. Lambiotte, The many facets of community detection in complex networks, *Appl. Netw. Sci.* 2 (2017) 4.
- [44] R. Aldecoa, I. Marín, Surprise maximization reveals the community structure of complex networks, *Sci. Rep.* 3 (2013) 1060.
- [45] P.W. Holland, K.B. Laskey, S. Leinhardt, Stochastic blockmodels: first steps, *Soc. Netw.* 5 (1983) 109–137.
- [46] C. Lee, D.J. Wilkinson, A review of stochastic block models and extensions for graph clustering, *Appl. Netw. Sci.* 4 (2019).
- [47] M.E.J. Newman, Fast algorithm for detecting community structure in networks, *Phys. Rev. E* 69 (2004) 066133.
- [48] N. Stephanopoulos, E. McGhee, Partisan gerrymandering and the efficiency gap, *Univ. Chic. Law Rev.* 82 (2015) 831–900.