

Towards control of opinion diversity by introducing zealots into a polarised social group

Antoine Vendeville

Department of Computer Science
Centre for Doctoral Training in Cybersecurity
University College London, United Kingdom

Benjamin Guedj

Department of Computer Science
Centre for Artificial Intelligence
University College London, United Kingdom
Inria, France

Shi Zhou

Department of Computer Science
Centre for Artificial Intelligence
University College London, United Kingdom

This is the corrected version of a paper that was published as A. Vendeville, B. Guedj and S. Zhou, Towards control of opinion diversity by introducing zealots into a polarised social group, *Complex Networks & Their Applications X*, pp. 341–352, 2022 [1]. Since then, a small mistake was found in equation (11). We fixed the mistake and updated Figures (3) and (4). Our overall findings are unchanged.

We explore a method to influence or even control the diversity of opinions

within a polarised social group. We leverage the voter model in which users hold binary opinions and repeatedly update their beliefs based on others they connect with. Stubborn agents who never change their minds (“zealots”) are also disseminated through the network, which is modelled by a connected graph. Building on earlier results, we provide a closed-form expression for the average opinion of the group at equilibrium. This leads us to a strategy to inject zealots into a polarised network in order to shift the average opinion towards any target value. We account for the possible presence of a *backfire effect*, which may lead the group to react negatively and reinforce its level of polarisation in response. Our results are supported by numerical experiments on synthetic data.

1 Introduction

We are interested in controlling opinions on connected networks with arbitrary degree distribution. In recent years, recommendation algorithms on social platforms have greatly enhanced confirmation bias by showing users content that is the most susceptible to match their interests — the so-called “filter bubble” effect [2]. As a consequence more and more isolated, tightly clustered online communities of similar-minded individuals have arisen in various domains such as politics [3, 4, 5], healthcare [6, 7] or science [8]. Because of the so-called *backfire effect*, presenting these users with opposing information might have the adverse effect of reinforcing their prior beliefs [9, 10]. In this paper we provide a simple method to shift diversity of opinions within towards a chosen target level, with and without the presence of a backfire effect.

To this end we rely on the well-known voter model, in which each user holds one of two possible opinions (e.g. liberal or conservative, pro or anti-abortion) and updates it randomly under the distribution of others’ beliefs. Independently introduced by Clifford and Sudbury [11] and Holley [12] in the context of particles interaction, this model has since been used to describe in a simple and intuitive manner social dynamics where people are divided between two parties and form their opinion by observing that of others around them. We assume some of the users are stubborn and never change opinion. We call them *zealots* as in [13, 14]. They can represent lobbyists, politicians or activists for example. Long time dynamics and limiting behaviour of such processes have been subject to several studies [14, 15, 16].

To achieve our goal we extend a previous result from the literature. Namely, authors of [14] found an expression for the average number of opinion-1 users at equilibrium in the $n \rightarrow \infty$ limit for a fully-connected network. Here we prove that their result also holds on expectation for any connected graph, assuming the placement of zealots is done uniformly at random. This allows us to find the optimal number of zealots to inject in a polarised community in order to reach any target average opinion at equilibrium, with and without the presence of a backfire effect. This effect we model simply by assuming that the injection of any number of zealots entails the *radicalisation* of some non-zealous users, turning them into zealots with the opposite opinion.

Our findings are illustrated through numerical simulations. Four different user graph topologies are considered: a fully-connected group where everyone is influenced by everyone else, an Erdős-Rényi random graph, a Barabási-Albert scale-free graph and a Watts-Strogatz small-world network. We find empirical averages to be close approximations of theoretical values. All code used for the simulations is available online.¹

2 Related Literature

Perhaps the earliest milestone in the study of opinion dynamics are the works from French [17] and Degroot [18] who studied how a society of individuals may or may not come to an agreement on some given topic. Assuming the society is connected and people repeatedly update their belief by taking weighted averages of those of their neighbours, they showed that consensus is reached. That is, everyone eventually agrees. Various other models have been developed since, to tackle the question of under which circumstances and how fast a population is able to reach consensus. Amongst others, [19] introduces immutable innate preferences, [20] studies the effect of homophily, [21] assumes individuals are perfectly rational and [22] accounts for the influence of external events.

The voter model was introduced independently by Clifford and Sudbury [11] and Holley [12] in the context of particles interaction. They proved that consensus is reached on the infinite \mathbb{Z}^d lattice. Several works have since looked at different network topologies, wondering whether consensus is reached, on which opinion and at what speed. Complete graphs [23, 24, 25, 26], Erdős-Rényi random graphs [25, 26], scale-free random graphs [27, 25], and other various structures [25, 26] have been addressed. Variants where nodes deterministically update to the most common opinion amongst their neighbours have also been studied [28, 29].

An interesting case to consider is the one where zealots – i.e. stubborn agents who always keep the same opinion, are present in the graph. Such agents may for example represent lobbyists, politicians or activists, i.e. entities looking to lead rather than follow and who will not easily change side. One of those placed within the network can singlehandedly change the outcome of the process [13, 25]. If several of them are present on both sides, consensus is usually not reachable and instead opinions converge to a steady-state in which they fluctuate indefinitely [14, 16].

Recently, Mukhopadhyay [15] considered zealots with different degrees of zealotry and proved that time to reach consensus grows linearly with their number. They also showed that if one opinion is initially preferred — i.e. agents holding that opinion have a lesser probability of changing their mind — consensus is reached on the preferred opinion with a probability that converges to 1 as the network size increases. Klamser *et al.* [30] studied the impact of zealots on a dynamically evolving graph, and showed that the two main factors shaping their influence are their degrees and the dynamical rewiring probabilities.

¹<https://github.com/antoinevendeville/howopinionscrystallise>

With the increasing importance of social networks in the political debate and information diffusion, there has been a recent surge in research aiming at controlling opinions, often with the goal to reduce polarisation. In the context of the Voter and Friedkin-Johnsen models respectively, Yildiz *et al.* [16] and Goyal *et al.* [31] provide algorithms for selecting an optimal sets of stubborn nodes in order to push opinions in a chosen direction. Our work differs from the former in that we propose a more general strategy, that works on expectation for any connected graph. Yi and Patterson [32] formulate different constrained optimisation problems under the French-Degroot and the Friedkin-Johnsen models. They provide solutions in the form of optimal graph construction methods.

Still within the Friedkin-Johnsen paradigm, Chitra and Musco [33] prove that dynamically nudging edge weights in the user graph can reduce polarisation while preserving relevance of the content shown by the recommendation algorithm. Garimella *et al.* [34] propose a method to reduce polarisation through addition of edges in the network. The focus is put on which nodes to connect in order to get the best reduction in polarisation, while being sure that the edge is “accepted” — as extreme recommendations might not work because of the backfire effect. Finally, Cen and Shah [35] propose a data-driven procedure to moderate the gap between opinions influenced by a neutral or a personalised newsfeed. Importantly, they show that this can be done even without knowledge of the process through which opinions are derived from the newsfeed.

3 The Voter Model with Zealots

Consider a graph with n users labelled $1, \dots, n$ that can each hold opinion 0 or 1. Given an initial distribution of opinions, each user updates their opinion at the times of an independent Poisson process of parameter 1 by adopting the opinion of one of its neighbours chosen uniformly at random. Letting $x_i(t)$ denote the opinion of user i at time t , we say that consensus is reached if almost surely all users eventually agree, i.e. if

$$\forall i, j, \quad \mathbb{P}(x_i(t) = x_j(t)) \xrightarrow[t \rightarrow \infty]{} 1. \quad (1)$$

On any finite connected network, consensus is reached [36]. Intuitively, no matter the current number of opinion-0 and opinion-1 users, there exists a succession of individual opinion changes with strictly positive probability that results in everyone holding the same opinion.

Here we place ourselves in the case where the user network is a connected graph with arbitrary degree distribution. This means there is a path from any user to another, and degrees of different users are not correlated with one another. Edges may be directed or not. We are interested in the particular situation where some of the agents are stubborn and never change their opinions. We call such agents *zealots*. Whenever a clock associated with a zealot rings, their opinion is not updated. They form an inflexible core of partisans within a group who bear great power of persuasion over the whole population. Both the position of these agents and initial opinions are assumed to be independent from the network topology.

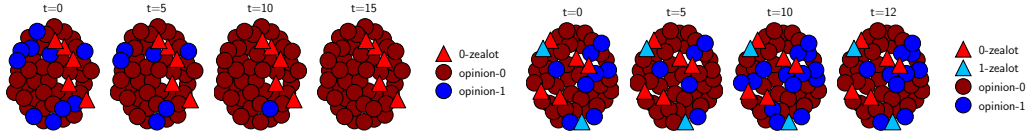


Figure 1: **(Left)** Example realisation of the model on a complete graph at different times with $n = 50, n_1 = 10, z_0 = 5, z_1 = 0$. Because there are no 1-zealots, everyone eventually adopts opinion 0. **(Right)** Same setting except $z_1 = 2$. Because there are zealots in both camps, the system reaches a state of equilibrium where no opinion prevails.

We call 1-zealot a stubborn agent with opinion 1 and denote by z_1 their number. Similarly, z_0 will denote the quantity of 0-zealots. The remaining $n - z_0 - z_1$ users are free to change opinions during the whole duration of the process. If $z_0 > 0, z_1 = 0$ or $z_0 = 0, z_1 > 0$ then via similar arguments as for eq. 1 consensus is reached on opinions 0 and 1 respectively. Here we are particularly interested in the case $z_0, z_1 > 0$, where there is always a strictly positive number of users with each opinion. This prevents consensus and instead the system reaches state of equilibrium in which it fluctuates indefinitely [14, 16]. We illustrate cases $\{z_0 > 0, z_1 = 0\}$ and $\{z_0, z_1 > 0\}$ in fig. 1.

Our results are valid as long as at least one of z_0 and z_1 is strictly positive and we formally require $z_0 + z_1 > 0$. Importantly, we assume that for any fixed tuple (z_0, z_1) , the position of zealots is drawn uniformly at random. Letting Z be a random vector encoding these positions we formally write $Z \sim \mathcal{U}$.

4 Expected Opinion Diversity at Equilibrium

Assuming $z_0 + z_1 > 0$, $N_1(t)$ converges to a state of equilibrium, which is characterised by a time-independent stationary distribution π . If $z_0 = 0$ (resp. $z_1 = 0$) then π is the constant distribution δ_n (resp. δ_0). This equilibrium does not depend on the initial opinions of non-zealots but on the topology of the graph and the position of zealots. Let N_1^* be a random variable distributed under π . Its average value informs us on the limiting opinion diversity of the group. Authors of [14] proved that this average was $nz_1/(z_0 + z_1)$ for complete graphs in the $n \rightarrow \infty$ limit. We now show that this results holds on expectation for any connected graph where the position of zealots is drawn uniformly at random.

Theorem 1. *For any connected user graph and any z_0, z_1 such that $z_0 + z_1 > 0$, we have*

$$\mathbb{E}_{Z \sim \mathcal{U}} \mathbb{E} [N_1^*] = n \frac{z_1}{z_0 + z_1}. \quad (2)$$

This theorem states that the proportion of opinion-1 users is expected to endlessly fluctuate around the ratio $z_1/(z_0 + z_1)$. For example, having twice as many zealots as the other camp will on average lead to count twice as many partisans. Thus the camp that boasts the biggest quantity of zealots is expected to be of bigger size in the long

run. The ratio $z_1/(z_0 + z_1)$ does not depend on n and can be interpreted as the average opinion diversity at equilibrium. Finally, if the graph is complete then all possible values of Z are equivalent and the left side of eq. 2 is simply $\mathbb{E}[N_1^*]$.

Theorem 1. From [16, Theorem 2.1], the vector of individual opinions $(x_1(t), \dots, x_n(t))$ converges in distribution to a random vector (x_1^*, \dots, x_n^*) . From [16, Proposition 3.2] we have that $\mathbb{E}[x_i^*]$ is equal to the probability that a random walk on the user graph initiated at node i is absorbed by the set of 1-zealots. Because here we consider a connected graph with z_0, z_1 zealots placed uniformly at random, it holds that $\mathbb{E}_{Z \sim \mathcal{U}} \mathbb{E}[x_i^*] = z_1/(z_0 + z_1)$. There are $n - z_0 - z_1$ non-zealots and z_1 1-zealots, thus we have

$$\mathbb{E}_{Z \sim \mathcal{U}} \mathbb{E}[N_1^*] = z_1 + (n - z_0 - z_1) \mathbb{E}_{Z \sim \mathcal{U}} \mathbb{E}[x_i^*] \quad (3)$$

and eq. (2) ensues. \square

4.1 Empirical Approximation

Let briefly discuss how to verify Theorem 1 through computer simulations. Assume we perform M simulations of our model on a given graph. Each time Z is drawn at random under \mathcal{U} at the beginning. Then the law of large numbers tell us that

$$\frac{1}{M} \sum_{m=1}^M \mathbb{E}[N_1^*] \xrightarrow{M \rightarrow \infty} n \frac{z_1}{z_0 + z_1}. \quad (4)$$

Now N_1^* is a random variable distributed under π . For a given realisation of $N_1(t)$, its expectation can be approximated empirically via

$$\frac{1}{T} \sum_{t \in \mathcal{T}} N_1(t) \xrightarrow{T \rightarrow \infty} \mathbb{E}[N_1^*] \quad (5)$$

where \mathcal{T} is a set of sufficiently large times and T denotes its size. Combining (4) and (5) we have:

$$\frac{1}{MT} \sum_{m=1}^M \sum_{t \in \mathcal{T}} N_1(t) \xrightarrow{M, T \rightarrow \infty} n \frac{z_1}{z_0 + z_1}. \quad (6)$$

Furthermore, we can also approximate the expected stationary distribution:

$$\frac{1}{MT} \sum_{m=1}^M \sum_{t \in \mathcal{T}} \mathbb{1}\{N_1(t) = k\} \xrightarrow{M, T \rightarrow \infty} \mathbb{E}_{Z \sim \mathcal{U}}[\pi_k]. \quad (7)$$

5 Control of Opinion Diversity

Theorem 1 is a useful tool for the control of opinions diversity at equilibrium. Let us represent a polarised online community by a connected graph of n individuals with a quantity $z_0 > 0$ of 0-zealots and no 1-zealot ($z_1 = 0$). For the sake of clarity we assume

the graph to be complete but the strategy holds on expectation when $Z \sim \mathcal{U}$ in the more general case of any connected graph. The completeness assumption is also fairly reasonable for certain social platforms, such as Facebook or Reddit where everyone in a group or subreddit sees the posts of everyone else.

Because $z_0 > 0$ and $z_1 = 0$, the community is homogeneous: each member will end up adopting opinion 0 no matter what. The presence of such polarised groups hinder democratic debate online, as they prevent the evolution of opinion and reinforce pre-existing beliefs. To mitigate this phenomenon, we suggest injecting 1-zealots into the group. Doing so means forcing $z_1 > 0$ and thus at equilibrium, the group will be more diverse and divided between the two camps.

From Theorem 1, the average opinion at equilibrium is given by $z_1/(z_0 + z_1) \in [0, 1]$ and equals 0 here. Let us choose a target value $0 < \lambda < 1$, representing the level of diversity we want the group to reach. Typically it should be around 1/2 if the goal is to transform the echo chamber into a diverse sphere of opinions. It is immediate that injecting a quantity

$$z_1^* = \frac{\lambda}{1 - \lambda} z_0 \quad (8)$$

of 1-zealots yields $z_1^*/(z_0 + z_1^*) = \lambda$. In practice we round this value to the previous or next integer. If $\lambda = 1/2$, then we should add exactly as many 1-zealots as there are 0-zealots.

Note that the equilibrium opinion is only determined by the quantity of zealots on each side and other users do not affect it. The expected opinion held by such a user converges to $z_1/(z_0 + z_1)$ and thus their number and initial opinions does not impact the equilibrium. Hence, one could “convert” users amongst the initial $n - z_0$ non-zealots into 1-zealots instead of adding external ones. That would lead to the same diversity, except that it comes with a limit on the quantity of 1-zealots because of the constraint $z_0 + z_1 \leq n$. Thus in that case, our strategy is only feasible if

$$z_0 + \frac{\lambda}{1 - \lambda} z_0 \leq n. \quad (9)$$

In a more general fashion, if we are limited in the number of zealots we can add to the network by z_1^{\max} , the optimal quantity of such agents is given by:

$$z_1^* = \min \left(z_1^{\max}, \frac{\lambda}{1 - \lambda} z_0 \right). \quad (10)$$

5.1 Backfire Effect

Numerous studies suggest that presenting certain users with opposing views might actually entrench them even deeper in their beliefs. This is known as the backfire effect. To account for it we study the scenario where in reaction to 1-zealots being created (either by addition or conversion), some of the non-zealous users will *radicalise* and become 0-zealots. Formally, we set that for each increment of z_1 , a quantity $\alpha < 1$ of non-zealous users become 0-zealots. The opinion at equilibrium is now given by

$$\frac{z_1}{z_0 + z_1 + \alpha z_1} \quad (11)$$

and diversity λ is exactly reached with

$$z_1^* = \lambda z_0 D^{-1} \quad (12)$$

where $D := 1 - \lambda - \lambda\alpha$. If $D > 0$ then $z_1^* > 0$ and we can inject this quantity of users into the network. If $D \leq 0$ however this becomes impossible as (12) is then either undefined or negative. In this case, we find that

$$z_1 \mapsto \left(\frac{z_1}{z_0 + z_1 + \alpha z_1} - \lambda \right)^2 \quad (13)$$

is strictly positive and decreasing towards $((1 + \alpha)^{-1} - \lambda)^2$ over $\mathbb{R}_{>0}$. Thus the larger z_1 the closer we get to the target diversity, up to a certain point. This means that the addition of 1-zealots entails the radicalisation of too many users into 0-zealots for it to ever counterbalance the spread of opinion 0. If we are limited in the number of zealots we can add to the network by z_1^{\max} , we have the natural optimal values for z_1 :

$$\begin{cases} z_1^* = \min(z_1^{\max}, \lambda z_0 D^{-1}) & \text{if } D > 0, \\ z_1^* = z_1^{\max} & \text{if } D \leq 0. \end{cases} \quad (14)$$

Note that in the case where we are converting existing users into zealots instead of injecting external ones, z_1^{\max} is given by the constraint

$$z_1^{\max} + \underbrace{z_0 + \alpha z_1^{\max}}_{\substack{\text{updated } z_0 \\ \text{after backfire}}} = n. \quad (15)$$

6 Numerical Experiments

We now validate Theorem 1 via computer simulations. Let us set $(n, z_0, z_1) = (100, 20, 40)$ and $\phi = nz_1(z_0 + z_1) \simeq 66.7$. We consider four different undirected graph models with varying parameters:

- Complete graph with an edge between each pair of users,
- Erdős-Rényi random graph with density $p = 0.1, 0.3, 0.5$,
- Watts-Strogatz small-world graph with initial connections to the 4 nearest neighbours and rewiring probability $\omega = 0.01, 0.05, 0.1$,
- Barabási-Albert scale-free graph with $m = 1, 3, 5$ initial nodes and m new connections at each step.

For each model and each parameter we generate one connected user graph then perform $M = 500$ simulations. At the beginning of each, position of zealots and initial opinions are drawn uniformly at random, then we let the model evolve for 200 time units. In Fig. 2 we plot values of $N_1(t)$ averaged over all simulations for each graph model and

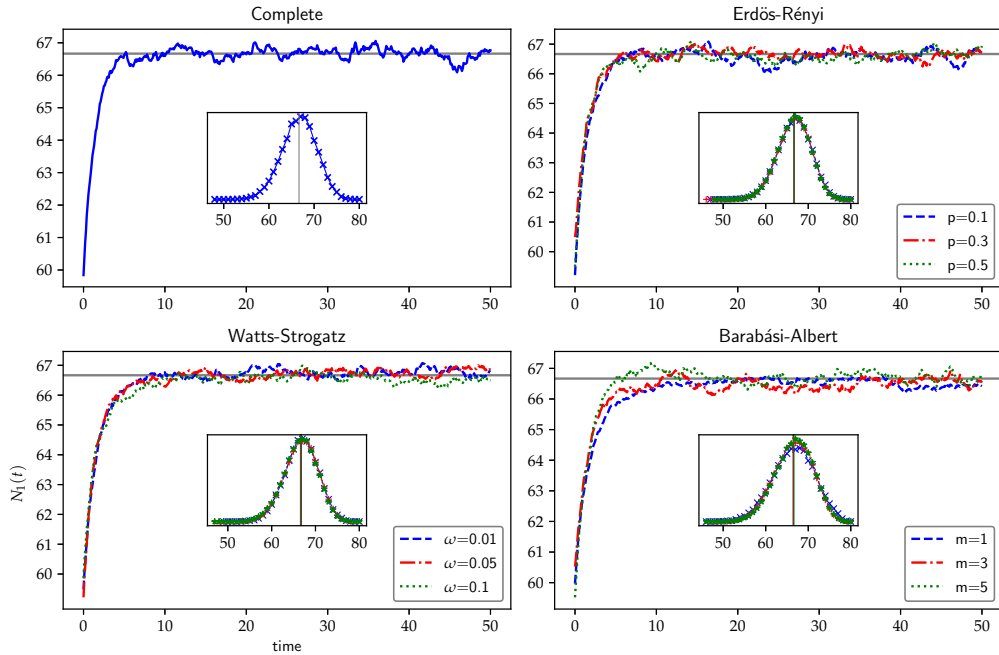


Figure 2: Evolution of $N_1(t)$ for various graph models and parameters. Values are averaged over 500 simulations. Horizontal grey lines indicate the limiting expectation given by Theorem 1. Insets show empirical distributions at equilibrium.

parameter—up to $t = 50$ and not 200 for the sake of clarity. We observe a good correspondence with the theory, as after about 20 time units all the empirical averages seem to fluctuate closely around ϕ .

Insets show empirical distributions of N_1^* obtained via eq. 7. For any given simulation, the empirical distribution is obtained by averaging all values of $N_1(t)$ starting at $t = 20$. This value was chosen by manual inspection, to make sure that the system had enough time to stabilise. Distributions for each choice of graph model and parameter are then averaged over all 500 simulations. We also show their means and its theoretical value ϕ . We observe very good correspondence between all the Erdős-Rényi and Watts-Strogatz graphs for all parameters, and a slightly worse one for the Barabási-Albert graphs. This might be due to the lower regularity in the graph topology, as when m increases, the distribution seems to get closer to the others. Finally, the means are all close to the theoretical value.

To quantify more precisely the discrepancy between theory and simulations, we look more closely at the difference between empirical $N_1(t)$ and ϕ . Again we only consider results from $t = 20$ onwards. For any graph model and parameter, the average relative error of the empirical means with respect to ϕ is calculated through:

$$\frac{1}{\phi} \frac{1}{M} \sum_{m=1}^M \left| \frac{1}{T} \sum_{t=20}^{100} N_1(t) - \phi \right| \quad (16)$$

Table 1: Relative errors between empirical averages of $N_1(t)$ at equilibrium and limiting expectation $nz_1/(z_0 + z_1)$, averaged over 500 simulations and with standard deviation.

Complete	$0.8\% \pm 0.6\%$	-	-
Erdős-Rényi	$p = 0.1$ $1.7\% \pm 1.2\%$	$p = 0.3$ $1.0\% \pm 0.7\%$	$p = 0.5$ $0.9\% \pm 0.6\%$
Watts-Strogatz	$\omega = 0.01$ $2.5\% \pm 1.9\%$	$\omega = 0.05$ $2.6\% \pm 1.9\%$	$\omega = 0.1$ $2.4\% \pm 1.9\%$
Barabási-Albert	$m = 1$ $4.7\% \pm 3.5\%$	$m = 3$ $3.3\% \pm 2.4\%$	$m = 5$ $2.8\% \pm 2.1\%$

where $M = 500$ is the number of simulations and T is the number of data points from $t = 20$ onwards. This value quantifies the average gap between empirical means of $N_1(t)$ at equilibrium and their theoretical value ϕ , expressed as a percentage relative to the value of ϕ . Results are presented in Table 1 alongside standard deviations.

As to be expected the complete network boasts the most precise results, with an average relative error of $0.8\% \pm 0.6\%$ corresponding to a difference of 0.5 ± 0.4 users. Then in order of decreasing accuracy we have successively Erdős-Rényi, Watts-Strogatz and finally Barabási-Albert graphs. Each time the precision increases with the density of the graph, as the network gets closer and closer to the complete case. The worst case is the Barabási-Albert network with parameter $m = 1$, yielding an average relative error of $4.7\% \pm 3.5\%$ which corresponds to a difference of 3.1 ± 2.3 users.

Finally, we turn to the problem of controlling the diversity of opinions under the presence of a backfire effect. We consider a complete graph of $n = 100$ users with $z_0 > 0, z_1 = 0$ and three different target diversities $\lambda = 0.1, 0.5, 0.9$. The intensity of the backfire effect takes various values $\alpha \in \{0.01, 0.1, 0.5\}$. The optimal quantity of 1-zealots is calculated according to (14), in the case where we are converting existing users and thus z_1^{\max} is given by eq. 15. In fig. 3 we plot z_1^* function of z_0 for each α . Additionally in fig. 4 we plot the error, that is the absolute difference between the diversity reached using z_1^* and the target λ .

As expected, for all λ considered, the lower α the higher quantities of 0-zealots it is possible to fight against. Peaks in the curves of fig. 3 correspond to tipping points after which $z_1^* = z_1^{\max}$, either because of the size constraint (15) or because D becomes non-positive. Thus diversity is reached exactly only before those peaks, as can be seen with nil errors in fig. 4. Notice that this peak does not necessarily appear—for example with $(\lambda, \alpha) = (0.9, 0.5)$, target diversity is never reached exactly. Finally, note that if we are looking to maximise the diversity of opinions—*i.e.* $\lambda = 1/2$, errors are rather small for a large range of z_0 values, meaning it is possible even in the presence of a backfire effect.

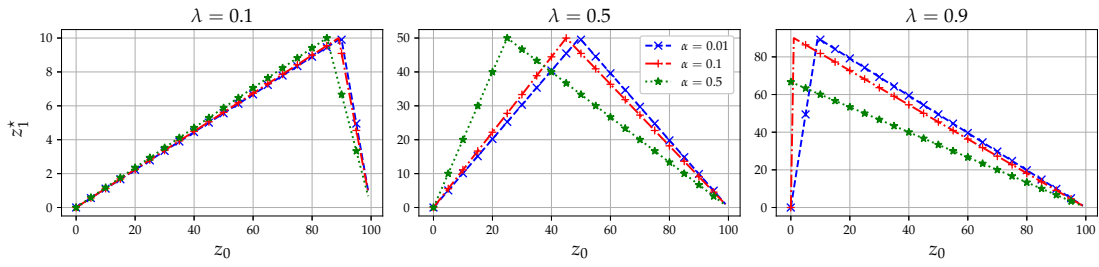


Figure 3: Controlling diversity of opinions with backfire effect. Optimal z_1^* function of z_0 for $n = 100$, $\lambda = 0.1, 0.5, 0.9$ and various intensities of the backfire effect α . Note that axes scales differ from one plot to another.

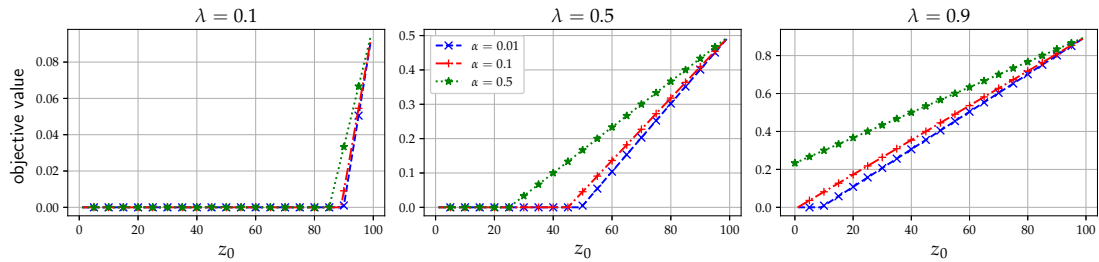


Figure 4: Absolute difference between diversity reached with $z_1^* = z_1^{\max}$ and target λ . Note that axes scales differ from one plot to another.

Conclusion and Future Work

In this paper we analysed the voter model with zealots on connected graphs with arbitrary degree distribution. Assuming that initial opinions and position of zealots are drawn uniformly at random, we extended existing results to provide closed-form expressions from (i) expected opinion distribution at equilibrium, and (ii) expected convergence time. We then used our findings to propose a simple method for the control of opinion diversity in a connected group of users that may or may not be subject to a backfire effect. Our analysis was supported through numerical simulations. Leads for further work include extensions such as considering more than two opinions or implementing variable degrees of zealotry.

References

- [1] Antoine Vendeville, Benjamin Guedj, and Shi Zhou. Towards control of opinion diversity by introducing zealots into a polarised social group. In *Complex Networks & Their Applications X*, pages 341–352, Cham, 2022. Springer International Publishing.

- [2] Eli Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. The Penguin Group, 2011.
- [3] Wesley Cota, Silvio C. Ferreira, Romualdo Pastor-Satorras, and Michele Starnini. Quantifying echo chamber effects in information spreading over political communication networks. *EPJ Data Sci*, 8(1):1–13, December 2019.
- [4] Michela Del Vicario, Fabiana Zollo, Guido Caldarelli, Antonio Scala, and Walter Quattrociocchi. Mapping social dynamics on Facebook: The Brexit debate. *Soc Netw*, 50:6–16, July 2017.
- [5] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Political discourse on social media: echo chambers, gatekeepers, and the price of bipartisanship. *Proc. of WWW2018*, pages 913–922, April 2018.
- [6] Daniel Allington, Bobby Duffy, Simon Wessely, Nayana Dhavan, and James Rubin. Health-protective behaviour, social media usage and conspiracy belief during the COVID-19 public health emergency. *Psychol Med*, page 1–7, 2020.
- [7] Harald Holone. The filter bubble and its effect on online personal health information. *Croat. Med. J.*, 2016.
- [8] Hywel T.P. Williams, James R. McMurray, Tim Kurz, and F. Hugo Lambert. Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Glob Environ Change*, 32:126 – 138, 2015.
- [9] Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. Exposure to opposing views on social media can increase political polarization. *PNAS*, 115(37):9216–9221, 2018.
- [10] Leonie Schaewitz and Nicole C. Krämer. Combating disinformation: Effects of timing and correction format on factual knowledge and personal beliefs. *Disinformation in Open Online Media*, pages 233–245, 2020.
- [11] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 12 1973.
- [12] Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann Probab*, 3(4):643–663, 08 1975.
- [13] Mauro Mobilia. Does a single zealot affect an infinite group of voters? *Phys. Rev. Lett.*, 91:028701, Jul 2003.
- [14] Mauro Mobilia, A. Petersen, and S. Redner. On the role of zealotry in the voter model. *J Stat Mech Theory Exp*, 2007:P08029–P08029, Aug 2007.
- [15] Arpan Mukhopadhyay. Voter and majority dynamics with biased and stubborn agents. *J Stat Phys*, 181:1239–1265, 2020.

- [16] M. E. Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Trans Econ Comput*, 1(4), Dec 2013.
- [17] J. R. French. A formal theory of social power. *Psychol. Rev.*, 63:181–94, 1956.
- [18] Morris H. DeGroot. Reaching a consensus. *J. Am. Stat. Assoc.*, 69(345):118–121, 1974.
- [19] Noah E. Friedkin and Eugene C. Johnsen. Social influence and opinions. *J. Math. Sociol.*, 15(3-4):193–206, January 1990.
- [20] Robert Axelrod. The dissemination of culture: A model with local convergence and global polarization. *J. Conflict. Resolut.*, 41(2):203–226, 1997.
- [21] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. *Games Econ. Behav.*, 46(1):1–22, jan 2004.
- [22] Ali Jadbabaie, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi. Non-bayesian social learning. *Games Econ. Behav.*, 76(1):210–225, 2012.
- [23] Yehuda Hassin and David Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Inf. Comput.*, 171(2):248–268, jan 2002.
- [24] E. Perron, D. Vasudevan, and M. Vojnovic. Using three states for binary consensus on complete graphs. *Proc. of IEEE INFOCOM 2009*, pages 2527–2535, 2009.
- [25] V. Sood, Antal Tibor, and S. Redner. Voter models on heterogeneous networks. *Phys. Rev. E*, 77:041121, Apr 2008.
- [26] M. E. Yildiz, R. Pagliari, Asuman Ozdaglar, and Anna Scaglione. Voting models in random networks. *Proc. of 2010 Information Theory and Applications Workshop (ITA)*, pages 1–7, Jan 2010.
- [27] John Fernley and Marcel Ortgiese. Voter models on subcritical inhomogeneous random graphs, 2019.
- [28] P. Chen and S. Redner. Majority rule dynamics in finite dimensions. *Phys. Rev. E*, 71:036101, Mar 2005.
- [29] Elchanan Mossel, Joe Neeman, and Omer Tamuz. Majority dynamics and aggregation of information in social networks. *Auton. Agent. Multi. Agent. Syst.*, 28(3):408–429, May 2014.
- [30] Pascal P. Klamser, Marc Wiedermann, Jonathan F. Donges, and Reik V. Donner. Zealotry effects on opinion dynamics in the adaptive voter model. *Phys. Rev. E*, 96, Nov 2017.

- [31] M. Goyal, D. Chatterjee, N. Karamchandani, and D. Manjunath. Maintaining ferment. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 5217–5222, 2019.
- [32] Yuhao Yi and Stacy Patterson. Disagreement and polarization in two-party social networks, 2019.
- [33] Uthsav Chitra and Christopher Musco. Analyzing the impact of filter bubbles on social network polarization. *Proc. of WSDM2020*, page 115–123, 2020.
- [34] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. Reducing controversy by connecting opposing views. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM '17*, pages 81–90, New York, NY, USA, February 2017. Association for Computing Machinery.
- [35] Sarah H. Cen and Devavrat Shah. Regulating algorithmic filtering on social media, 2020.
- [36] David J. Aldous and James Allen Fill. Reversible markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.