Social Influence of Competing Groups and Leaders in Opinion Dynamics

Catherine A. Glass[1] · David H. Glass[2]

Abstract

This paper explores the influence of two competing stubborn agent groups on the opinion dynamics of normal agents. Computer simulations are used to investigate the parameter space systematically in order to determine the impact of group size and extremeness on the dynamics and identify optimal strategies for maximizing numbers of followers and social influence. Results show that (a) there are many cases where a group that is neither too large nor too small and neither too extreme nor too central achieves the best outcome, (b) stubborn groups can have a moderating, rather than polarizing, effect on the society in a range of circumstances, and (c) small changes in parameters can lead to transitions from a state where one stubborn group attracts all the normal agents to a state where the other group does so. We also explore how these findings can be interpreted in terms of opinion leaders, truth, and campaigns.

✉ David H. Glass
dh.glass@ulster.ac.uk

Catherine A. Glass
catherine.glass@afbini.gov.uk

[1] Economics Research Branch, Agri-Food and Biosciences Institute, 18a Newforge Lane, Belfast BT9 5PX, UK

[2] School of Computing, Ulster University, Shore Road, Newtownabbey, County Antrim BT37 0QB, UK

🖄 Springer

1 Introduction

Opinions and beliefs are important as they impact our behaviour and choices, and are acquired to a great extent through social interaction. Opinion dynamics is a field of study that explores how a host of pertinent factors effect changes in opinions or beliefs. The first mathematical models of opinion dynamics can be traced back to French (1956), Harary (1959) and DeGroot (1974). French's theory of social power explored the impact of interpersonal interactions on opinion formation and Harary extended this work to further explore conditions leading to unanimity. DeGroot described a model in which the opinions of an agent (individual) are updated by taking an average of an agent's opinion and those in his group (now called the DeGroot model). Since then, opinion dynamics has attracted growing attention from researchers across a number of disciplines.

In economics, opinion dynamics is explored in the context of social learning, which has been the subject of much investigation in the literature (Banerjee 1992; Bikhchandani et al. 1992; Smith and Sørensen 2000; Acemoglu and Ozdaglar 2011; Jadbabaie et al. 2012; Mobius and Rosenblat 2014; Molavi et al. 2018). Social learning has been described as "the social aspect of belief and opinion formation" (Acemoglu and Ozdaglar 2011) and as such is highly relevant to many important economic phenomena including the adoption of new products and services, spread of new technologies or innovations, financial contagion on stock markets and word-of-mouth job search. Social learning models have been used to explore diffusion of innovations (Martins et al. 2009), product adoption (Ruf et al. 2017), the emergence of fads and fashions through herding of beliefs and information cascades (Banerjee 1992), word of mouth learning (Banerjee and Fudenberg 2004; Ellison and Fudenberg 1995), advertising and marketing (Schulze 2003; Sznajd-Weron and Weron 2003) and movie reviews (Barriere et al. 2017).

Opinion dynamics has also been investigated at length in the physics literature, particularly from the 1970s and 80s when sociophysics began to emerge (Weidlich 1972, 1971; Callen and Shapero 1974; Galam et al. 1982) as a new field of study. Techniques used in statistical physics were adapted to explore social phenomena such as cultural dissemination, evolution of language, population dynamics, social contagion, epidemic spreading, and opinion dynamics (Castellano et al. 2009). One such technique was the Ising model, a model used to explain critical phase transitions where very simple interactions can lead to qualitative changes on a macroscopic scale (Sznajd-Weron 2005). Models have explored group decision making in various contexts including firms and small committees (Galam 1997), voting behaviour (González et al. 2004; Bernardes et al. 2002), how to convince others (Stauffer 2003), making political predictions (Galam 2017), modelling the impact of social influence on the prices of options (Oster and Feigel 2015) and the modelling of industrial strikes in big companies (Galam et al. 1982). Across sociology, psychology, politics, public policy and business, there have been a number of relevant studies including: experiments to explore the effect of social influence on judgment shifts (Moussaïd et al. 2013), understanding the psychological factors affecting opinion formation with an application

to American politics (Duggins 2017), the initiation of smoking amongst adolescents (Sun and Mendez 2017) and consensus reaching in social network group decision making (Dong et al. 2018, 2017).

Models of opinion dynamics/social learning fall into two broad groups—Bayesian and non-Bayesian. Both Bayesian (Banerjee 1992; Bikhchandani et al. 1992; Smith and Sørensen 2000; Acemoglu and Ozdaglar 2011) and non-Bayesian (Jadbabaie et al. 2012; Molavi et al. 2018; Bala and Goyal 1998; DeMarzo et al. 2003) approaches are common in the economics literature while non-Bayesian approaches are more popular with sociophysicists and other researchers in opinion dynamics. Bayesian approaches use Bayes' rule to combine an agent's prior opinion with new evidence and information that becomes available to it. This becomes its posterior probability and provides a prior distribution for future updating cycles. A shortcoming of this approach, however, is the selection of an initial prior, a process requiring complex reasoning on the part of the agent. Non-Bayesian approaches to social learning involve so called "rules or thumb" (a classic example being the DeGroot model) in which agents modify their beliefs based on simple updating rules. For example, an agent could use a simple updating rule in which it combines its prior opinion with those of its friends and neighbours, giving greater weight to those closer to it.

Models of opinion dynamics / social learning can also be distinguished in terms of whether they use discrete opinions, continuous opinions or a combination of the two. Examples of models with a discrete opinion format include the voter model (Holley and Liggett 1975), the majority rule model (Galam 2002) and the Sznajd model (Sznajd-Weron and Sznajd 2000). In these models agents have a simple yes or no, 1 or 0, agree or disagree choice. All of these were inspired by the Ising model discussed above. Models with a continuous opinion format include the Deffuant-Weisbuch (DW) model in which each agent can take on an opinion value anywhere in the interval from 0 to 1. Two agents interact only if their opinions lie within a specified distance of each other and so this model is referred to as a bounded confidence model. A second model of this type is the Hegselmann and Krause (HK) model (Krause 2000; Hegselmann and Krause 2002, 2005) which is used in this paper. As with the DW model, opinions can take on any value between 0 and 1 and again agents are influenced only by those within their bound of confidence or neighbourhood. However, with the HK model, a given agent interacts with all agents within its neighbourhood at the same time.

An alternative model developed by Martins (2008a) and Martins (2008b) combines continuous opinions with discrete actions (CODA). Discrete actions are observed publicly but internal opinions are represented on a continuous scale and are updated using a Bayesian approach. Martins (2008a) applies CODA to both the voter model and the Sznajd model.

The standard Bayesian approach can be formulated in terms of identifying the equilibria of a dynamic game with incomplete information (see for example Acemoglu and Ozdaglar 2011). There are also game-theoretic connections with non-Bayesian approaches. For example, Bindel et al. (2015) showed that the repeated averaging process of a variant of the DeGroot model can be interpreted as the best-response dynamics in a complete information game that converges to the unique

Nash equilibrium. In terms of bounded confidence models, Di Mare and Latora (2007) have similarly presented a game theoretic formulation of the DW model, while Etesami and Başar (2015) have shown how an asynchronous version of the HK model (i.e. only one agent updates its opinion at a given time rather than all the agents doing so simultaneously) can be formulated in terms of a sequence of best response updates in a potential game. In both of these cases the steady state of the bounded confidence model is again equated with a Nash equilibrium.

Of interest in opinion dynamics is the study of individuals or groups exerting greater influence on agents' opinions than the normal agent. These include opinion leaders, stubborn agents/zealots and radical (extreme) agents. The opinion leader is an individual or group, who because of their expertise or prominence is able to stimulate change in public opinion, for example, mass media, political parties, community leaders, celebrities and bloggers. Studies have shown that agents may filter information received from mass media (including social media) through opinion leaders as they help normal agents to interpret and process it (Katz 1957; Roch 2005; Choi 2015). Stubborn agents or zealots are agents whose views are fixed. They can influence others more than normal agents because they will influence others, but are not influenced by others. In reality, of course, just how fixed opinions are will depend on the context. A local politician may bend somewhat to the will of their constituents while an extreme political group may remain fully inflexible. See Chen et al. (2016) for a study of the impact of varying degrees of stubbornness of opinion leaders on the group dynamics.

Radical agents are stubborn individuals or groups holding extreme views. There are a number of papers relevant to these groups in the opinion dynamics literature. The effect of zealotry was explored using the voter model (Mobilia et al. 2007; Yildiz et al. 2013; Acemoglu et al. 2013), the adaptive voter model (Klamser et al. 2017) and the naming game model (Verma et al. 2014; Waagen et al. 2015). Extremeness using the DW model has been explored in Deffuant et al. (2000), Deffuant (2006) and Gargiulo and Mazzoni (2008), and using the CODA model in Martins and Kuba (2010). Hegselmann and Krause (2015) used a variant of their HK model to represent radical groups and opinion leaders in a unified way and then investigated the influence of a single radical group / opinion leader. Chen et al. (2016) used the HK model to study four characteristics of successful opinion leaders by comparing competing opinion leaders. By keeping three characteristics fixed while varying the fourth, they reported that greater success in attracting followers is achieved in general by opinion leaders who are less stubborn (i.e. their opinions can change to some extent), have greater appeal and are less extreme. They also reported that higher reputation (the weight of influence on normal agents, see Sect. 2 below) helps the leader to attract more followers when the confidence is high, but can hinder the leader from attracting followers when the confidence is low.

In this paper, we adopt the common framework of Hegselmann and Krause (2015) and Chen et al. (2016) to examine the impact of competing stubborn agent groups on normal agents, though this can also be interpreted in terms of competing opinion leaders. We present a systematic study using computer simulations to investigate how varying the extremeness and size of stubborn agent groups, which corresponds to the reputation in Chen et al. (2016), affects the behaviour of the normal

agents. Unlike Chen et al. (2016), we do not vary stubbornness or appeal, but instead we explore the parameter space in much more detail and so are able to investigate the effects of varying both the extremeness and size of the stubborn agent group over the relevant space. This enables us to investigate how increasing or decreasing one of these features can compensate for changes in the other. It also allows us to explore how groups can optimize their approach to achieve key objectives such as maximizing follower numbers and maximizing societal influence. Overall, this leads to a more complex picture of the influence of these features. For example, our results show that there are many cases where a stubborn group of intermediate extremeness and size achieves the best outcome.

In the rest of the paper, Sect. 2 lays the theoretical foundations of the HK model in the presence of competing groups. Section 3 examines the impact of varying the number of agents in a stubborn group and the extremeness of the group. Section 4 explores in-depth the optimal strategies to play for success in the context of competing groups. Section 5 provides a wider discussion of the work, while Sect. 6 presents conclusions.

2 HK Model with Competing Groups

The standard HK model consists of n agents who update their opinions about a given topic at each time step by taking the average belief of their neighbours (Hegselmann and Krause 2002). Let agent i have belief $x_i(t)$ lying in the interval [0, 1] at time t. Agent j will be a neighbour of agent i at t if the belief of agent j lies within a certain bound of confidence of agent i, that is $|x_i(t) - x_j(t)| \leq \epsilon_i$, where ϵ_i is the confidence of agent i. Each agent's belief is then updated as follows:

$$x_i(t+1) = \frac{1}{|I(i,t)|} \sum_{j \in I(i,t)} x_j(t), \tag{1}$$

where $I(i,t) = \{j : |x_i(t) - x_j(t)| \leq \epsilon_i\}$, i.e. the set of agents who are in the neighbourhood of agent i.

Following recent work by Hegselmann and Krause (2015), we represent both groups and opinion leaders by means of the influence of a constant signal. To do this we include a group of agents, who all have the same opinion, R_1, in addition to the normal agents described above. These are stubborn agents who do not change their opinion at all, but they influence the *normal* agents in the usual way via Eq. (1) if they are within their neighbourhood. As in Hegselmann and Krause's approach, we consider the influence of this group of agents for different values of R_1 and for different numbers of agents within it, denoted #R_1. We also follow their proposal that this group of stubborn agents can be interpreted in different ways. First, they can be interpreted straightforwardly as a group of radical agents, especially in cases where their views are extreme, i.e. R_1 is close to 0 or 1. In general, however, we will refer to them as a group of stubborn agents since we will consider non-extreme groups. Second, they can be interpreted as a single opinion, or charismatic, leader in which case #R_1 represents the weighting given to the leader in comparison to normal agents

in the opinion updating process. This weighting is referred to as the *reputation* of the opinion leader by Douven and Riegler (2010) and Chen et al. (2016), who investigated an equivalent model where the opinion leader was given a weighting directly rather than being modelled by a number of separate agents.[1]

A third interpretation that is mentioned briefly by Hegselmann and Krause (2015) is that the group of stubborn agents represents the truth about some topic and the number of such agents determines the influence of the truth on normal agents in the updating process. Note that in this interpretation, it is only agents who are close enough to the truth (i.e. the truth lies within their confidence interval) that are influenced by it. This is in contrast to distance independent representations of truth in the HK model where all agents are influenced by the truth (Hegselmann and Krause 2006; Douven and Riegler 2010). Given this interpretation, our results indicate the number of agents who find the truth. This interpretation is particularly relevant in the context of social learning. Hegselmann and Krause (2015) also suggest other interpretations such as dogmatists, since the opinion of the group need not be extreme, and a campaign, where a message is communicated with a given intensity. For ease of exposition, we will frame the discussion mostly in terms of stubborn groups, while drawing attention to other interpretations in Sect. 5.

The current work differs from Hegselmann and Krause (2015) in two respects. First, they adopted a deterministic approach by fixing the initial opinions of the normal agents so that they are evenly distributed between 0 and 1 for all simulations. The initial opinion of agent i was set to $i/(n+1)$ for $i = 1, \ldots, n$. By contrast, we will adopt the strategy typically used in the opinion dynamics literature (including in Hegselmann and Krause 2002) of assigning initial opinions of normal agents randomly and then averaging over repeated simulations.

Second, instead of considering the influence of a single radical group on the normal agents, we follow Chen et al. (2016) by investigating the dynamics of the normal agents when there are two competing groups. By competing we mean that one group has a fixed opinion greater than 0.5 and the other less than 0.5. This raises a further set of questions about the relative influence of the groups depending on their respective fixed opinions and size of membership. For example, does a group that is less extreme in opinion generally acquire more followers? Or how does the size of a group affect the number of followers? The model is therefore a straightforward modification of Eq. (1) as follows:

---

[1] In fact, the model of Chen et al. 2016 can be seen as a generalization of the Hegselmann and Krause (2015) model since it does not require the opinion leader to have a fixed opinion and allows other agents to have greater confidence towards the opinion leader than they do towards other agents. For reasons of efficiency, we have obtained our results by implementing the model of Chen et al. 2016 since it replaces many stubborn agents with a weighting of one by a single agent with a greater weighting.

$$x_i(t+1) = \frac{1}{|I(i,t)|} \sum_{j \in I(i,t)} x_j(t), \text{ if } i \text{ is a normal agent}$$

$$= x_i(0) = R_1, \text{ if } i \text{ contributes to first group/opinion leader}$$

$$= x_i(0) = R_2, \text{ if } i \text{ contributes to second group/opinion leader}$$

(2)

where the set $I(i, t)$ is defined as before and includes all agents, whether normal or otherwise, who are within the confidence of agent i at time t.

Results are presented in the next section as various parameters are varied. In particular, we vary the confidence, ε, the number of agents in each stubborn group, $\#R_1$ and $\#R_2$, and the opinions of these groups, R_1 and R_2. Note that we assume homogeneous confidence levels such that $\varepsilon_i = \varepsilon$ for all agents, though heterogeneous confidence levels could also be explored (Chen et al. 2017; Pineda and Buendía 2015; Liang et al. 2013; Lorenz 2010).

The results are obtained by simulations where we iterate Eq. (2) until a steady state is reached, which we implement by ensuring that $|x_i(t+1) - x_i(t)| \leq 10^{-5}$ for all agents i. We focus on the number of agents who end up with the same opinion as a given stubborn group or, to be precise, we consider an agent i to have the same opinion as that of the stubborn group R_1 at the end of the simulation if $|x_i^* - R_1| < 10^{-3}$, where x_i^* is the opinion of agent i at the end of the simulation. Additionally, we also present some results for the root mean square deviation (RMSD) of the final opinions of normal agents from a radical group. For a stubborn group with opinion R_1, this is given by:

$$RMSD = \sqrt{\frac{\sum_{i=1}^{n}(x_i^* - R_1)^2}{n}},$$

(3)

where n is the number of normal agents, which is set to 200 in all cases. For each result, we ran the simulation 200 times using MATLAB and obtained average values for the number of followers and RMSD.

3 Results for Competing Groups

3.1 Varying the Number of Agents

First, we consider how the results depend on the number of stubborn agents, $\#R_1$ and $\#R_2$, in groups one and two, or equivalently, the reputation of two opinion leaders. In this case we set the opinion of agents in group one to be $R_1 = 1$ and the opinion for group two to be $R_2 = 0$. Figure 1 presents results for the number of normal agents who end up with the same opinion as group one, i.e. $R_1 = 1$. The results are presented as a function of $\#R_1$ and $\#R_2$ for four different values of the confidence, ε. Figure 1a shows that for $\varepsilon = 0.2$, group one attracts (1) very few followers close to the y-axis where $\#R_1$ is very low (less than 15), (2) many more followers (70–80) for slightly higher values of $\#R_1$ (about 40), and (3) intermediate numbers of followers (40–50) for higher values of $\#R_1$.
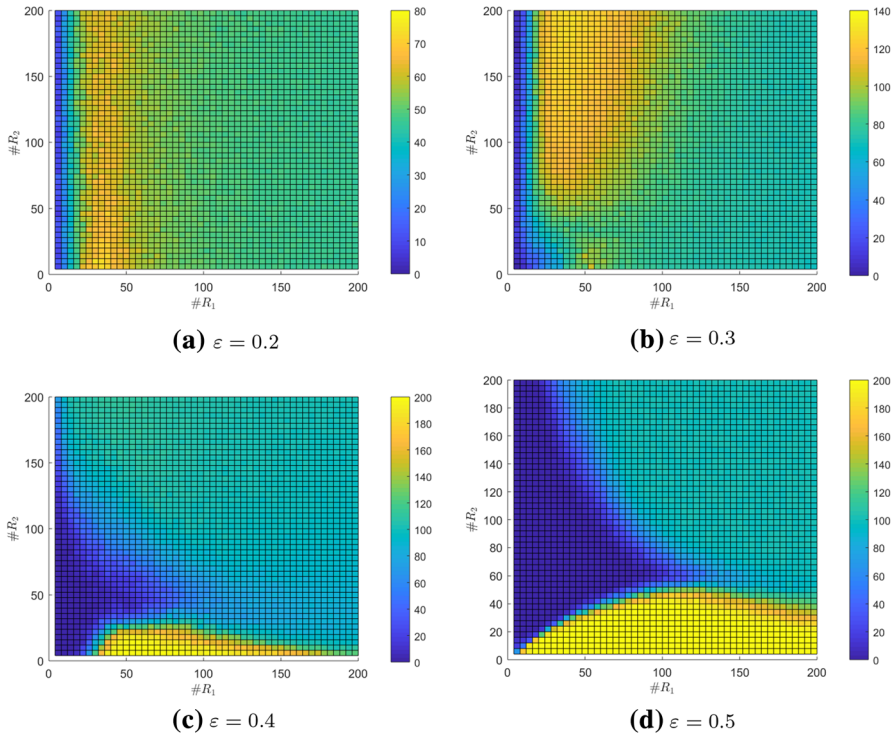
Fig. 1 The average number of normal agents who go to R_1 as a function of the number of agents $\#R_1$ in group one and the number of agents $\#R_2$ in group two. Note that the results are not all on the same scale

Intuitively, two features of these results might seem surprising. First, the number of agents in group two, $\#R_2$, has little effect on these results. However, recall that the confidence is 0.2 in this case and so there is a lot of separation between normal agents who are directly influenced by group one (those with opinions of 0.8 or higher) and those directly influenced by group two (those with opinions of 0.2 or lower). Second, one might have expected group one to attract more followers when its numbers are larger (high $\#R_1$), but instead it attracts more followers when its numbers are lower (but not too low). This can be explained by a well-known feature of the HK model. When the group size, $\#R_1$, is too big it exerts a large influence on normal agents within its neighbourhood, drawing them in too quickly for them to have much influence on other normal agents. However, when $\#R_1$ is lower it does not draw followers in so quickly and in the meantime they can influence other agents, so that in the end the group can acquire more followers.

We see in Fig. 1b that the results are quite different when $\varepsilon = 0.3$ compared to $\varepsilon = 0.2$. As in (a), group one still attracts very few followers when $\#R_1$ is very low and intermediate numbers of followers for higher values of $\#R_1$. It also attracts more followers for values of $\#R_1$ around 40–50, but this is now more dependent on the value of $\#R_2$, with the number of followers being much greater when $\#R_2$

is high. This can be explained by the same mechanism as discussed above. Since group two is large it has a lot of influence on its neighbours, and so it draws in its neighbours quickly meaning they have little influence on other normal agents. By contrast, group one is smaller and so influences its neighbouring normal agents more slowly meaning they can have a larger influence on other normal agents, resulting in more followers in the long run. This is illustrated by means of a typical individual run in figure S1 of the supplementary material.

We see that the results change again as the confidence increases to $\varepsilon = 0.4$ in Fig. 1c. Group one still attracts very few followers when $\#R_1$ is very low, but now it attracts most followers in a region close to the x-axis (low $\#R_2$). In fact, in this region it attracts all or almost all the normal agents. The low value of $\#R_2$ means that group two exerts little influence on its neighbours, while the high value of confidence means that group one can exert a lot of influence both directly and indirectly. Note, however, that if $\#R_1$ is too high, group one is no longer able to attract so many followers because it draws its neighbours in too quickly.

In Fig. 1d, where $\varepsilon = 0.5$, the trend from (c) continues and now we see three distinct areas. One of these is the region close to the x-axis (low $\#R_2$) where group one attracts all the normal agents, though this region has expanded in comparison to figure (c). A second region is at low values of $\#R_1$ where group one attracts no followers. This region has expanded considerably compared to the previous figures and is mostly due to the growing influence of group two, although some of it is due to neither group attracting any followers. (Note that the number of followers for group two is obtained just by swapping the axes.) In the third region, where neither $\#R_1$ nor $\#R_2$ are low, group one attracts half of the normal agents (with group two attracting the other half). It is worth noting that for $\varepsilon = 0.5$ (and also $\varepsilon = 0.4$), the normal agents would form a central consensus in the absence of stubborn groups. Hence, the stubborn groups lead to a polarization of the society in this case.

Figure 2 shows results that are closely related to Fig. 1, but instead of giving the number of followers of group one, it gives the difference between the number of followers of group one and the number of followers of group two. This enables us to identify when one group has an advantage over the other. For lower values of confidence (Fig. 2a and b), group one does better than group two when the size of group two is very small (less than 20), but this is only because group two attracts virtually no followers in this region. If group two is larger, group one can gain the advantage by having a relatively small group (around 40). So having a large group size is not the best strategy for gaining followers.

This pattern continues at $\varepsilon = 0.4$ in Fig. 2c, but now when the size of group two is large, group one is only able to gain a very small advantage over group two when the size of group one is around 40. At $\varepsilon = 0.5$ in Fig. 2d the only region where group one does better than group two is when group two is small and group one is larger. If group two is large, group one cannot acquire more followers than group two irrespective of the size of group one. The best outcome one can aim for in such a context is parity, which is achieved by just ensuring that its group size is above some threshold; no advantage is gained by increasing the group size further.

In Fig. 2d, it is clear that when $\#R_1$ and $\#R_2$ are small, a small change in $\#R_1$ or $\#R_2$ can result in a dramatic change in the number of followers acquired by the
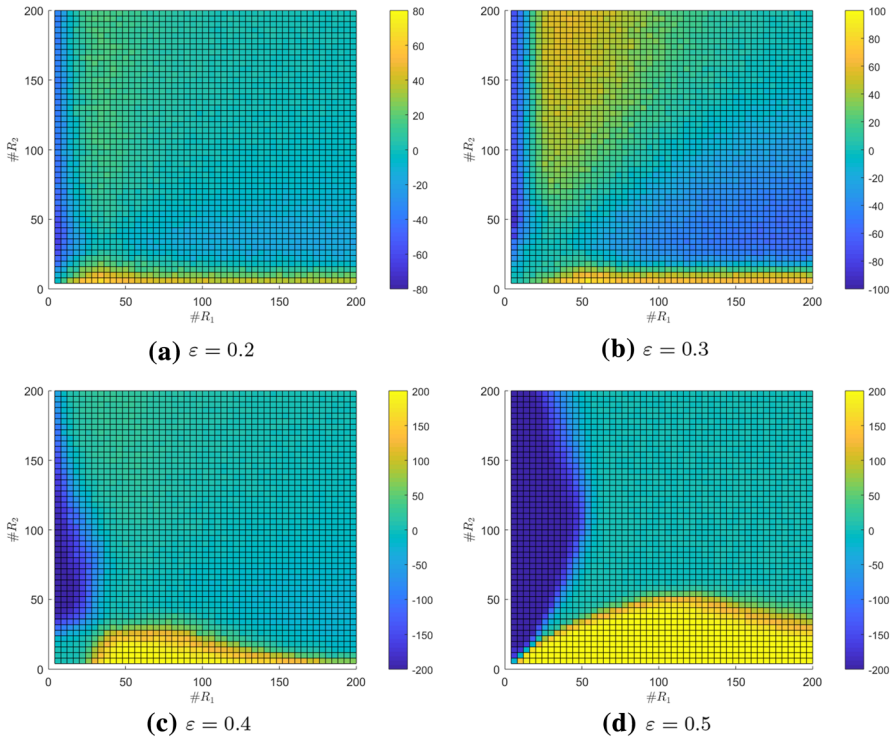
Fig. 2 The difference between the average number of normal agents who go to $R_1 = 1$ and the average number who go to $R_2 = 0$ as a function of the number of agents $\#R_1$ in group one and the number of agents $\#R_2$ in group two. Note that the results are not on the same scale
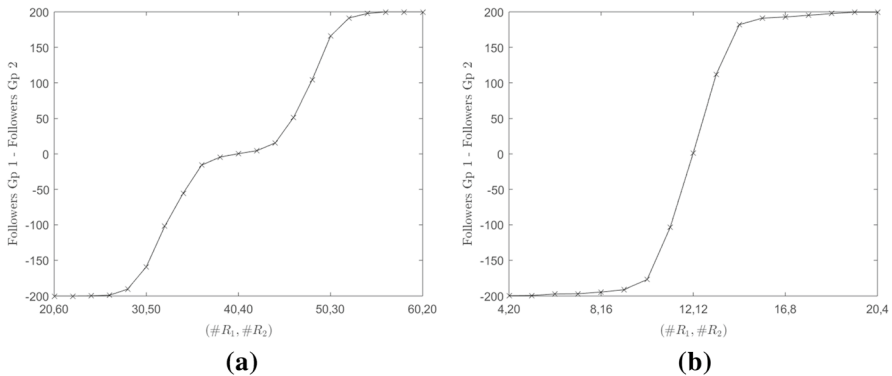

Fig. 3 The difference between the average number of normal agents who go to $R_1 = 1$ and the average number who go to $R_2 = 0$ along lines where $\#R_1$ and $\#R_2$ are both varied. In (**a**) $\#R_1$ and $\#R_2$ vary along the line from $\#R_1 = 20$, $\#R_2 = 60$ to $\#R_1 = 60$, $\#R_2 = 20$, while in (**b**) they vary along the line from $\#R_1 = 4$, $\#R_2 = 20$ to $\#R_1 = 20$, $\#R_2 = 4$

groups. This is explored in detail in Fig. 3. In Fig. 3a the difference in the number of followers between the groups is investigated as $\#R_1$ and $\#R_2$ are varied along a diagonal line from $\#R_1 = 20$, $\#R_2 = 60$ to $\#R_1 = 60$, $\#R_2 = 20$. These results have been averaged over 1000 runs to obtain more accurate values. At $\#R_1 = 20$, $\#R_2 = 60$ all the normal agents go to $R_2 = 0$. However, as $\#R_1$ increases and $\#R_2$ decreases, there is a transition to a state where neither group acquires any followers for values around $\#R_1 = \#R_2 = 40$, and then as they change further, another transition occurs to a state where all the normal agents go to $R_1 = 1$. Figure 3b shows corresponding results along a diagonal line from $\#R_1 = 4$, $\#R_2 = 20$ to $\#R_1 = 20$, $\#R_2 = 4$. In this case, the two transitions appear to have merged into one larger transition from a state where all the normal agents go to $R_2 = 0$ to another state where they all go to $R_1 = 1$. When $\#R_1 = \#R_2 = 12$ neither group acquires any followers, but there does not appear to be an intermediate state around this point; it just appears to be the midpoint of the larger transition. Figures 1 and 2 have already provided complementary perspectives on the results. Figure S2 in the supplementary material gives a third perspective, presenting the total number of normal agents who go to either group one or group two.

3.2 Varying the Extremeness of Groups

We now consider how the results depend on the extremeness of the stubborn agents in groups one and two. Recall that the opinion of groups one and two were set to $R_1 = 1$ and $R_2 = 0$ in the previous section. This scenario, where $1 - R_1 = R_2 = 0$, represents the most extreme case, but more generally we can represent the extremeness of groups one and two by $1 - R_1$ and R_2 respectively. We can vary the extremeness for each group from 0 to 0.5. Figure 4 shows the average number of normal agents who go to R_1 as the extremeness of R_1 and R_2 varies, while the number of stubborn agents in each group is kept fixed at $\#R_1 = \#R_2 = 40$. By symmetry all results for R_1 apply equally to R_2. The results are presented for four different values of confidence, ε.

Figure 4a shows the numbers of followers won over to R_1 at $\varepsilon = 0.2$. When R_1 is quite extreme, i.e. R_1 is greater than 0.85 (or equivalently in the figure when $1 - R_1$ is less than 0.15), the number of followers gained by R_1 is around 60–90 with the greatest gains being made by R_1 when R_2 is neither too extreme ($R_2 = 0$) nor too central ($R_2 = 0.5$) but when it has a value around 0.2. As R_1 moves further away from the extremes and towards the centre, the gains made by R_1 increase to around 90–120 normal agents as it expands its influence, though this depends on R_2 not moving too close to the centre. For example, at $R_1 = 0.7$ and $R_2 = 0.2$, R_1 gains around 110 followers but at $R_2 = 0.5$, it gains none. The highest gains of 120+ followers are made by R_1 as it approaches 0.5, as long as R_2 is less than around 0.15.

The triangular shaded area arises in the top right as the zones of influence of R_1 and R_2 overlap resulting in neither group gaining any followers. For example, at $\varepsilon = 0.2$, $R_1 = 0.6$ and $R_2 = 0.4$, R_1 has a direct influence on all normal agents with opinions from 0.4 to 0.8 and R_2's direct influence is from 0.2 to 0.6. In the overlapping range (0.4–0.6), all the normal agents are influenced by both R_1 and R_2
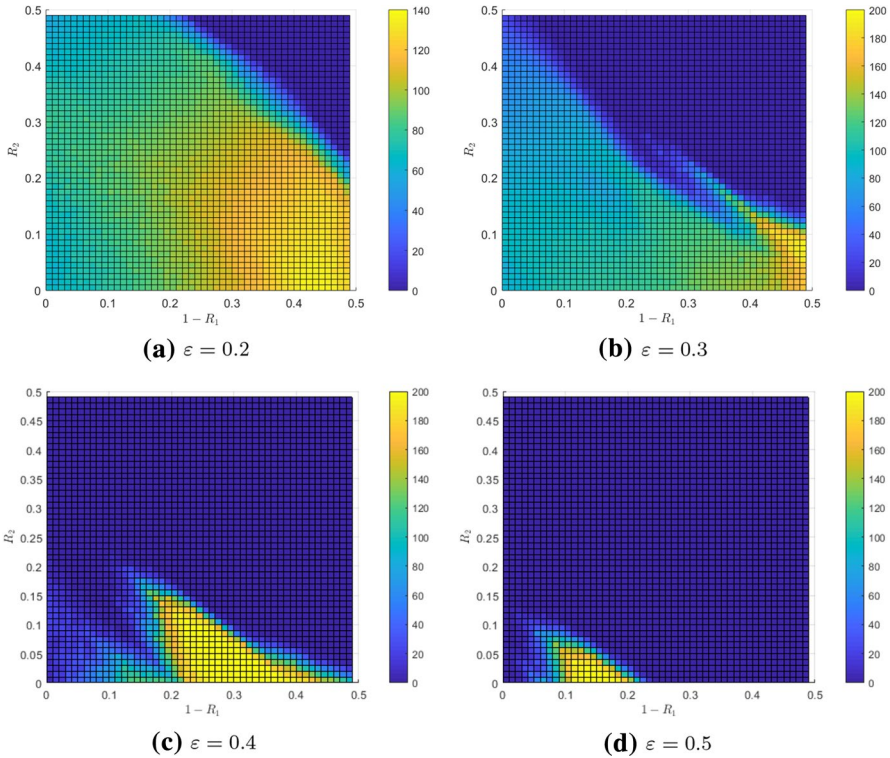
Fig. 4 The average number of normal agents who go to R_1 as a function of $1 - R_1$ and R_2. As $1 - R_1$ and R_2 increase the two groups become less extreme. Note that the results are not all on the same scale

simultaneously and the end result is that these agents converge to the midpoint of R_1 and R_2 resulting in neither group attracting any followers. Convergence to the midpoint assumes $\#R_1 = \#R_2$, but it is easy to show that in general they will converge to a point, x^*, that is a weighted combination of R_1 and R_2, with the weighting dependent on the respective sizes of the stubborn groups, $\#R_1$ and $\#R_2$, as follows

$$x^* = \frac{\#R_1 \times R_1 + \#R_2 \times R_2}{\#R_1 + \#R_2}. \tag{4}$$

From this line of reasoning, it follows that neither group will attract any followers when $R_1 - R_2 < \varepsilon$ (where we assume that $R_1 > R_2$) or equivalently when $1 - R_1 + R_2 > 1 - \varepsilon$. Hence, when $\varepsilon = 0.2$ we would expect no followers for either group above a line running from $1 - R_1 = 0.3$, $R_2 = 0.5$ to $1 - R_1 = 0.5$, $R_2 = 0.3$ in Fig. 4a, which roughly corresponds to the triangular area. In actual fact, the triangular area extends slightly below this line, indicating that even if the distance between the two stubborn groups is slightly greater than ε, the normal agents lying in between them are still able to form a central group. This behaviour is illustrated in figure S3 in the supplementary material, where the results of two individual runs are presented.
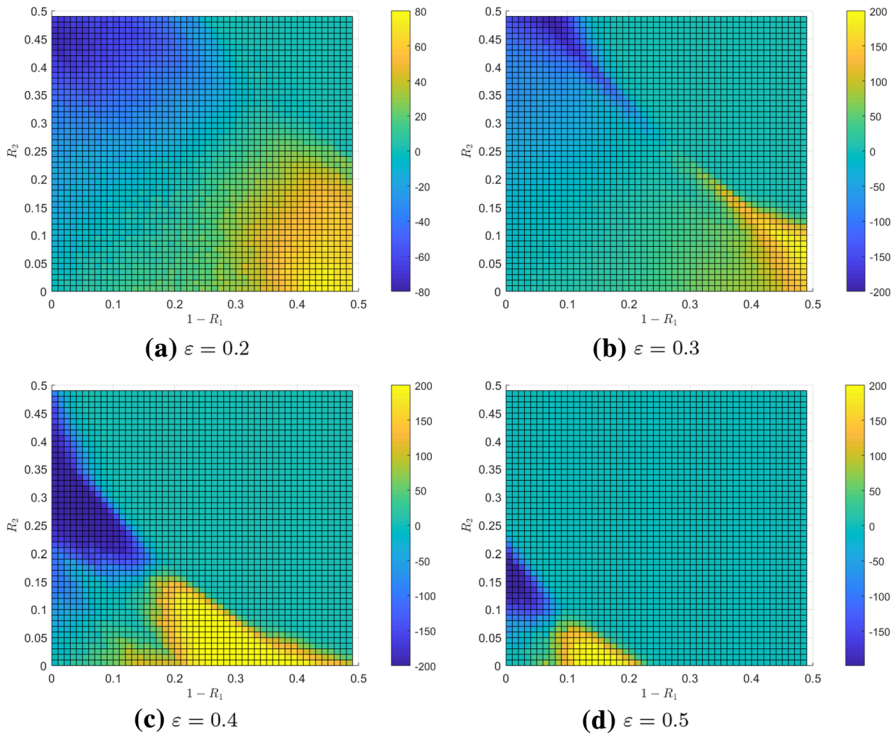
Fig. 5 The difference between the average number of normal agents who go to R_1 and the average number who go to R_2 as a function of $1 - R_1$ and R_2. Note that the results are not all on the same scale

Note that there is quite a sharp transition in Fig. 4a between the region where group one attracts a lot of followers and the region in the top right corner where it attracts none. This is very relevant to group one's strategy in terms of attracting followers. If R_2 is low (below about 0.15), the best strategy for group one is to become less extreme since as $1 - R_1$ increases group one attracts more followers. However, if R_2 is greater, this strategy is no longer valid in general, and if R_2 is greater than 0.25 it can go seriously wrong. Although, group one can attract more followers by becoming slightly less extreme, if it moves too close to the centre, it can quickly move from maximizing its number of followers to gaining none at all.

Figure 4b shows the numbers of followers won over to R_1 at $\varepsilon = 0.3$. At $R_2 = 0$, R_1 steadily increases its followers from 100 at $R_1 = 1$ to around 120 at $R_1 = 0.8$ and right up to around 180 followers at $R_1 = 0.5$. However, as R_2 rises above 0.1, R_1 is forced to move away from 0.5 to maximize follower numbers. Indeed, there is a much bigger area in which R_1 fails to attract any followers at $\varepsilon = 0.3$ compared to $\varepsilon = 0.2$ because there are more scenarios in which the zones of influence of R_1 and R_2 overlap. Following our earlier discussion, we would expect neither group to attract followers when $1 - R_1 + R_2 > 1 - \varepsilon = 0.7$. Hence we would expect no followers for either group above a line running from $1 - R_1 = 0.2$, $R_2 = 0.5$ to

$1 - R_1 = 0.5$, $R_2 = 0.2$ in Fig. 4b (see also Fig. 5b), which roughly corresponds to the triangular area, though as in Fig. 4a the actual area is slightly bigger than that. We also note that in the absence of stubborn groups, the standard HK model almost always gives consensus when $\varepsilon = 0.3$. Hence, there is a substantial area where the presence of the two groups prevents consensus from occurring.

Figure 4c and d show the numbers of followers won over to R_1 at $\varepsilon = 0.4$ and $\varepsilon = 0.5$ respectively. At these confidence levels most strategies of R_1 will fail to attract any followers, which can again be partly attributed to overlapping zones of influence. However, even in the absence of stubborn groups, the standard HK model always gives consensus when $\varepsilon = 0.4$ and 0.5, with a single central group merging at a value close to 0.5. In most cases, the two stubborn groups are unable to prevent this from happening, except in confined regions where one of the stubborn groups can attract most or all of the normal agents. However, typically if only one stubborn group had been present, it would have acquired all the normal agents as followers, so the presence of the second stubborn group helps restore the central consensus and so has a moderating effect on the society. Furthermore, although not shown here, there are some cases where neither stubborn group attracts followers, but they prevent a single central consensus from forming by causing it to be fragmented into several central groups which are close together. For example, this typically occurs when both R_1 and R_2 have extreme values (R_1 close to 1 and R_2 close to 0).

Figure 5 shows the difference between the average number of agents who go to R_1 and the average number who go to R_2 as a function of 1 "R_1 and R_2. These results correspond directly to those presented in Fig. 4 and enable us to identify when one group does better than the other in attracting followers. Note that all of the results presented are anti-symmetric: if R_1 attracts more followers than R_2 for a particular pair of values of $1 - R_1$ and R_2, then the result will be reversed if we swap the values of $1 - R_1$ and R_2. So far we have considered the number of followers acquired by group one and the difference between the number acquired by group one and group two. However, rather than trying to maximize its number of followers, we could also think of group one trying to minimize its RMSD (figure S4, supplementary material).

In summary, we see from Figs. 4 and 5 that overall, high confidence does not help either stubborn agent group in the current case where $\#R_1 = \#R_2 = 40$. As ε increases from 0.2 to 0.5, the stubborn groups' scope to attract followers is reduced considerably, with neither group acquiring any followers for most of the space when $\varepsilon = 0.5$. However, where they do acquire followers, one group typically acquires all of them. When $\varepsilon = 0.2$ or 0.3, group one's best strategy for attracting followers will be to become less extreme provided group two is sufficiently extreme. However, if group two is less extreme, group one will gain no followers if it moves too close to the centre. When $\varepsilon = 0.4$ or 0.5, group one's best option is to move away from the extreme position at 1 and position itself between 0.6 and 0.7 when $\varepsilon = 0.4$ or between 0.8 and 0.9 when $\varepsilon = 0.5$. However, its success depends critically on R_2 being below about 0.15 when $\varepsilon = 0.4$ or about 0.05 when $\varepsilon = 0.5$.
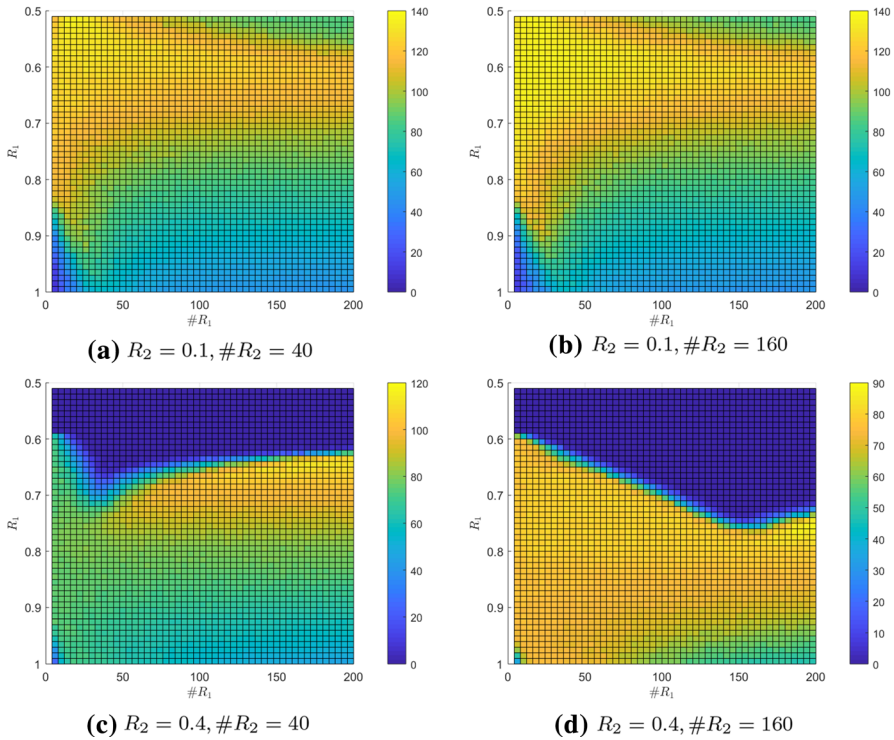
(a) $R_2 = 0.1, \#R_2 = 40$

(b) $R_2 = 0.1, \#R_2 = 160$

(c) $R_2 = 0.4, \#R_2 = 40$

(d) $R_2 = 0.4, \#R_2 = 160$

Fig. 6 The average number of normal agents who go to R_1 as a function of the number of agents $\#R_1$ in group one and the opinion R_1 of group one for the different values of R_2 and $\#R_2$ specified. Note that the results are not all on the same scale

4 Selecting Optimal Strategies

In the previous section, we changed the size of the stubborn groups for particular opinions of the two groups (i.e. the extreme values of zero and one) and then changed their extremeness for particular sizes of the two groups (40 stubborn agents). This enabled us to identify the optimal strategy for a group by changing either its size or its extremeness, but here we investigate how to select the optimal strategy when it has the flexibility to vary both parameters. In particular, we consider fixed values for the size of group two, $\#R_2$, and the opinion of group two, R_2, and then obtain results for different values of $\#R_1$ and R_1 to determine group one's optimal strategy.

We consider four pairs of values for $(\#R_2, R_2)$: $(40,0.1)$, $(160,0.1)$, $(40,0.4)$ and $(160,0.4)$. Figure 6 presents results for the number of followers acquired by group one as $\#R_1$ and R_1 are varied for each of these four pairs in (a), (b), (c) and (d) respectively. When $R_2 = 0.1$ in (a) and (b) we find that the optimal strategy for group one is to have a small group size (15–20) and be non-extreme (close to the most central value of $1 - R_1 = 0.49$). This is true irrespective of whether group two is small

($\#R_2 = 40$) or large ($\#R_2 = 160$) and in both cases group one is able to attract about 135–140 followers. Hence when $R_2 = 0.1$, group one's strategy does not appear to depend on the size of group two. The reason for the small group size is based on a point noted earlier that if group one is too large it can draw in its immediate neighbours too quickly, so that they fail to bring other normal agents along with them. This is particularly relevant when group one is close to the centre because it allows normal agents with opinions above 0.7 or so to form their own separate group. However, if group one is too small it can result in group two attracting some more followers.

Now consider Fig. 6c where $R_2 = 0.4$ and $\#R_2 = 40$. Here group one's optimal strategy is to have a large group size (close to the largest value of 200) and be relatively non-extreme ($1 - R_1$ around 0.35). Notice, however, that this point is very close to a boundary so that if group one gets too close to the centre ($1 - R_1 \geq 0.39$) then it will attract no followers. This is due to the phenomenon we discussed in the context of Fig. 4a whereby if the difference between the two stubborn groups is very similar to the confidence level, then a group of normal agents will converge to their midpoint. Suppose that group one had a fixed size of 40 (as was the case in Fig. 4a) so that it could only improve its performance by selecting its opinion. In that case it could only attract a maximum of about 90 followers. However, by also adjusting its group size to around 200 as is the case here, it can attract around 110 followers.

When we consider Fig. 6d where $R_2 = 0.4$ and $\#R_2 = 160$, group one's optimal strategy is again to have a large size (200), but this time it needs to be a bit more extreme (around $1 - R_1 = 0.24$). As in (c), this point is close to a boundary. So when $R_2 = 0.4$, group one's optimal strategy is to increase its group size to around 200 irrespective of the size of group two and to adjust its opinion to be close to the relevant boundary.

Why are the optimal strategies so different for $R_2 = 0.1$ in (a) and (b), where group one should have a small group size and non-extreme opinion (close to 0.5), in comparison to those for $R_2 = 0.4$ in (c) and (d), where it should have large group size and somewhat more extreme value? In effect, we have already addressed the extremeness question. When $R_2 = 0.4$ neither group will acquire any followers if group one moves too close to 0.5. So it is beneficial for group one not to be too extreme, but it must not get too close to the centre either. What about the large group size? When $R_2 = 0.1$ this is a disadvantage, so why is it an advantage when $R_2 = 0.4$? Because the neighbourhoods of group one and group two overlap, there will be normal agents who are directly influenced by both groups. Hence, if group one is larger it will exert greater influence on these normal agents. Furthermore, because group one is not too close to the centre (unlike the case where $R_2 = 0.1$) it can still draw in normal agents who have high opinion values. This combination of factors favours a large group size when $R_2 = 0.4$.

So far we have assumed that group one's goal is to maximize the number of followers it acquires, but Fig. 7 presents corresponding results if group one is trying to maximize the difference between its followers and the followers of group two. From (a) and (b), we can see that when $R_2 = 0.1$, group one's strategy remains unchanged: small group size and non-extreme. In (c) where $R_2 = 0.4$ and $\#R_2 = 40$, the optimal strategy is again very close to its position in Fig. 6c. However, in (d) where $R_2 = 0.4$
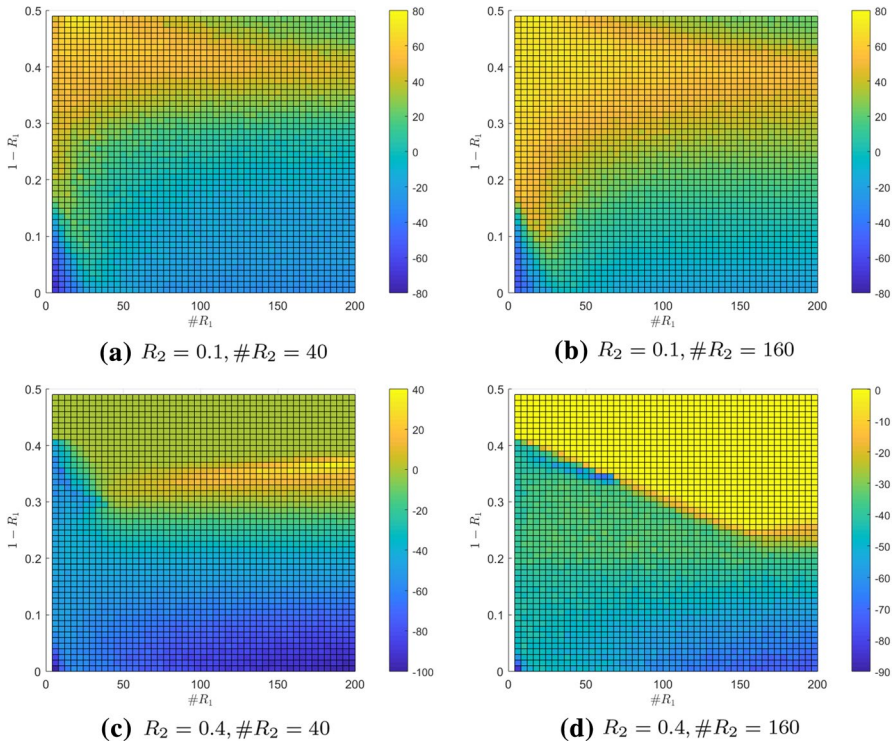
(a) $R_2 = 0.1, \#R_2 = 40$

(b) $R_2 = 0.1, \#R_2 = 160$

(c) $R_2 = 0.4, \#R_2 = 40$

(d) $R_2 = 0.4, \#R_2 = 160$

Fig. 7 As for Fig. 6, but now the results are for the difference between the average number of normal agents who go to R_1 and the average number who go to R_2

and $\#R_2 = 160$, things are very different. Now group one's optimal strategy is to select any point above a boundary, which corresponds to the region where group one attracts no followers. The reason for this surprising result is that by adjusting its size and opinion, group one is unable to attract more followers than group two so all it can do is aim for parity, which is achieved by ensuring that neither group attracts any followers.

Figure 8 presents corresponding results for a third optimization strategy for group one. This time the goal is to optimize the extent to which normal agents converge to positions closer to group one than to group two as measured by the difference between the RMSD for the two groups. We will refer to this goal as *maximizing the group's social influence* since it is concerned with the influence on the opinions within the society irrespective of whether the group acquires any followers as defined in Sect. 2. Given this goal, the optimal strategy for group one when $R_2 = 0.4$ in figures (a) and (b) is to have a large group size (close to 200) and be non-extreme ($R_1 = 0.51$). Note the contrast in terms of group size with the previous two optimization goals. How can this difference be explained? Recall that the reason why a large group size for group one did not enable it to attract a lot of followers was because it drew in followers too quickly allowing a separate
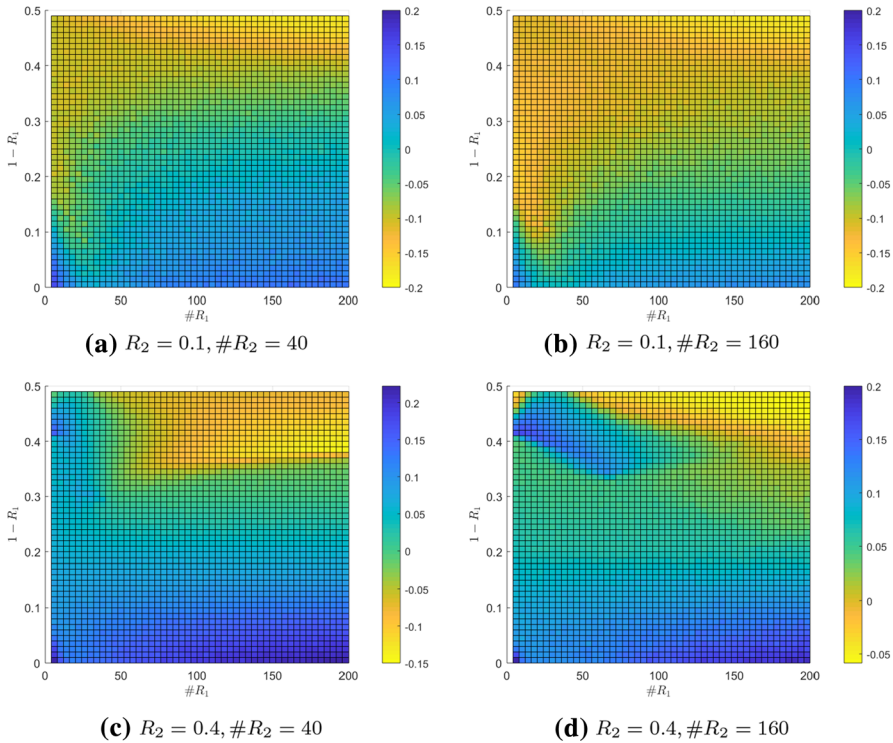
(a) $R_2 = 0.1, \#R_2 = 40$

(b) $R_2 = 0.1, \#R_2 = 160$

(c) $R_2 = 0.4, \#R_2 = 40$

(d) $R_2 = 0.4, \#R_2 = 160$

Fig. 8 As for Fig. 6, but now the results are for the difference between the RMSD for group one and the RMSD for group two

group of normal agents to form with an opinion greater than 0.7. However, now this third group becomes an advantage for group one because although they reduce its RMSD score since they are some distance away from it ($R_1 = 0.51$), they are even further away from group two ($R_2 = 0.4$) and so reduce its RMSD score even further.

In Fig. 8c where $R_2 = 0.4$ and $\#R_2 = 40$, the optimal strategy is to have a large group size (200) and a relatively non-extreme opinion (around $1 - R_1 = 0.38$), which is just above the boundary mentioned earlier and would result in group one acquiring very few followers. It might seem surprising that group one has no followers in the optimal case but results from individual runs (supplementary figure S5) show that all of the normal agents form a consensus group very close to group one, although not close enough to be counted as followers. In (d) where $R_2 = 0.4$ and $\#R_2 = 160$, the optimal strategy is also to have a large group size (200), but the opinion should be closer to the centre (around $1 - R_1 = 0.46$).

These results are summarized qualitatively in Table 1 along with corresponding results for $R_1 = 0.25$ and all of the results repeated again but with $\varepsilon = 0.4$. Results for the optimal values of R_1 are characterized as follows: Central (C) if $R_1 \leq 0.6$, Moderate (M) if $0.6 < R_1 < 0.9$, and Extreme (E) if $R_1 \geq 0.9$. Similarly,

Table 1 Optimal strategies for group one's opinion, R_1, and size, $\#R_1$, for various values of ε, R_2 and $\#R_2$ given three distinct goals: maximizing the number of followers (Goal 1); maximizing the difference between the number of followers of group one and group two (Goal 2); maximum social influence—the difference between the RMSD score of group one and group two (Goal 3)

ε	R_2	$\#R_2$	Goal 1		Goal 2		Goal 3	
			R_1	$\#R_1$	R_1	$\#R_1$	R_1	$\#R_1$
0.2	0.1	40	C	S	C	S	C	L
0.2	0.1	160	C	S	C	S	C	L
0.2	0.25	40	C	M	C	M	C	M
0.2	0.25	160	M	L	C	L	C	L
0.2	0.4	40	M	L	M	L	M	L
0.2	0.4	160	M	L	C, M	S, M, L	C	L
0.4	0.1	40	M, C	M, L	M, C	M, L	M	M
0.4	0.1	160	M	L	M	L	C	L
0.4	0.25	40	M	M, L	M	M, L	M	M
0.4	0.25	160	E	M	E, M, C	S, M, L	M	L
0.4	0.4	40	E, M	M, L	E, M	M, L	E	M
0.4	0.4	160	E	M	E, M, C	S, M, L	M	L

Values for R_1 are Extreme (E), Moderate (M) or Central (C), while $\#R_1$ has values Small (S), Medium (M) or Large (L)

optimal values of $\#R_1$ are characterized as: Small (S) if $\#R_1 \leq 40$, Medium (M) if $40 < \#R_1 < 160$, and Large (L) if $\#R_1 \geq 160$. If there is a range of optimal values spanning different regions, all relevant regions are noted. For example, we saw in Fig. 7d that if $\varepsilon = 0.2$, $R_2 = 0.4$ and $\#R_2 = 160$, then R_1 has optimal values of R_1 for Goal 2 (maximizing the difference between the followers of group one and group two) in both the Central and Moderate regions, while $\#R_1$ has optimal values in the Small, Medium and Large group sizes, so all of these are specified in the relevant line in Table 1.

These results show that the optimal strategy varies a lot depending on the confidence, characteristics of group two, and the goal. We note several general points. While it might be expected that having a central rather than extreme opinion would be most beneficial (see Chen et al. 2016), it turns out that this is only true in certain cases. In particular, it applies in almost all cases when confidence is low ($\varepsilon = 0.2$), and group two is either extreme ($R_2 = 0.1$) or $R_2 = 0.25$. In other cases group one's optimal strategy is usually to be neither too extreme nor too central, but have a moderate opinion. In our earlier discussion, we saw that a reason for this is that if group one is too central this can result in its acquiring no followers. However, in some cases, having an extreme opinion can be the optimal strategy. For example, when confidence is high ($\varepsilon = 0.4$) and group two is central $R_2 = 0.4$, having an extreme opinion can be to group one's advantage. Given the central position of group two and the high confidence value, the only way for group one to attract followers is to move to an extreme opinion.

We have already seen that a small group size can be the optimal strategy in some cases, but in the case of Goal 1 (maximizing followers), this only applies when confidence is low ($\varepsilon = 0.2$ and group two is extreme ($R_2 = 0.1$). Small group size is not optimal for Goal 3 (maximizing social influence) in any of the cases considered. Our discussion of Fig. 8 explains why this is the case. Note also from Table 1 that for

Goal 3 the optimal strategy for group one is always to have a large group size whenever group two is large. As we have already noted, a possible disadvantage of having a large group size is that it can draw in followers too quickly, but this need not be a problem when Goal 3 is in mind as discussed in the context of Fig. 8c. However, in cases where the two stubborn groups have overlapping neighbourhoods so that both have the same direct influence over some normal agents, it is important to have a larger group size to exert a greater influence.

5 Discussion

We have explored opinion dynamics mainly in the context of competing stubborn agent groups. However, our results may be interpreted in a number of different ways. We will briefly consider our results in terms of opinion leaders, truth and campaigns.

It is not surprising that stubborn opinion leaders or groups can lead to polarization of the society. However, our findings also show the potential power of opinion leaders who are fixed in their views to moderate extremism. For example, two opposing opinion leaders can bring about a central consensus that would have been unlikely to occur if neither had been present. Also, even in cases where a central consensus would have occurred in the absence of any stubborn opinion leaders or groups, suppose that in fact there is a single stubborn group that acquires all the normal agents as followers. We find cases of this kind where a suitable opinion leader would bring about a central consensus. In these types of scenarios, opinion leaders who are not too extreme help moderate extremists. Interestingly, we have even found cases where an opinion leader with extreme views could also have a moderating effect. All of the cases just noted occur in scenarios where the opinion leader gains no followers, but nevertheless exerts a powerful social influence.

Our findings have shown that increasing the size of a stubborn group, especially where there is a desire to maintain a certain level of extremeness, may lead to a worse outcome for the group. This carries over straightforwardly to opinion leaders: a greater reputation, so that more influence is exerted on those who are close to the leader's opinion, can result in fewer followers.

As noted in Sect. 2, Hegselmann and Krause (2015) discuss an interpretation of their model in terms of truth. Given this interpretation, our simulations can be seen as modelling a scenario where group one represents the truth and group two represents an opposing opinion, which could be a social media source of fake news, for example. As such, our simulations enable us to determine the extent to which social learning would be affected by this alternative source. Our results show that if the confidence of the society is high, then successful social learning could take place with everyone in society ending up finding the truth provided the source of fake news does not have too much influence. Alternatively, everyone could end up with the false belief of the fake news source, but this can be avoided provided the truth has sufficient influence.

The results in Sect. 3.2 can be interpreted as exploring scenarios where the truth does not lie at the edge of the opinion space. The results again suggest that successful social learning could take place provided the source of fake news is extreme and

the value of truth itself is neither too extreme nor too central. Otherwise no-one will end up finding the truth and in most cases they will end up at a position intermediate between truth and fake news. For lower values of confidence, some people can find the truth even if the fake news is not so extreme, though there can also be a danger of no-one finding the truth.

The handling of a public health campaign is another possible interpretation of our model. In this interpretation, the size of the group could represent the intensity of the campaign, while the opinion of the stubborn group would now represent where the campaign positions itself on the spectrum of opinion. Our findings could be useful to help identify strategies for campaign organizers to achieve maximum social influence. For example, if the confidence of a society is low (people are resistant to change) on a given issue, and the campaign is opposing an extreme position that has a lot of influence, our results suggest that the campaign organizers should position themselves in the centre of the spectrum of opinion and organize a high-intensity campaign. On the other hand, if the society is more open to persuasion on an issue and the opposing message is not a dominant one, the campaign organizers will maximise influence by delivering a lower-profile but more extreme message.

6 Conclusions

We have investigated the influence of competing stubborn agent groups on the opinion dynamics of normal agents through the use of computer simulations based on the HK model. We explore the parameter space in considerable detail to investigate the impact of varying both the extremeness and size of competing stubborn agent groups, identifying optimal strategies for maximising follower numbers and social influence.

Intuitively, it might be expected that groups which are large and central would be more successful, but our results show that this is often not the case. We find some scenarios where a small group size or extreme opinion can be part of the optimal strategy for a stubborn group. There are many cases where having neither too small nor too large a group size can lead to the best outcome for a group. Similarly, there are many cases where having an opinion that is neither too central nor too extreme can lead to the best outcome for a group. It might also be thought that competing stubborn groups would lead to a more polarized society, but that is not always the case. We have seen that competing groups can bring about a central consensus among the normal agents (with neither stubborn group acquiring any followers) that would be unlikely to occur in the absence of the stubborn groups. We also find cases where neither stubborn group acquires any followers, but if only one of them had been present, it would have acquired all the normal agents. In such cases, the introduction of the second stubborn group has a moderating influence on the society. We have also identified cases where small changes in the sizes of the two stubborn groups can bring about a sharp transition from one group acquiring all the followers to the other group's doing so.

We acknowledge that any conclusions are tentative. Our model depends on a number of simplifying assumptions and so there is scope for further developments in

several respects. For example, as mentioned earlier, heterogeneous confidence levels and degrees of stubbornness could be introduced, though this would make it more difficult to explore the full parameter space. It is clear that future research in this area would also benefit from connections with other disciplines such as psychology and political science as well as from further empirical studies to establish more fully which models/model features are most appropriate for a given context (Sobkowicz 2009; Takács et al. 2016). In addition to these suggestions for extending the work, we hope that the framework adopted here, which involves systematic investigation of the parameter space to identify non-trivial dynamics, will prove useful in other studies of the social influence of groups and opinion leaders.

References

Acemoglu, D., Como, G., Fagnani, F., & Ozdaglar, A. E. (2013). Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, *38*, 1–27.

Acemoglu, D., & Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, *1*(1), 3–49.

Bala, V., & Goyal, S. (1998). Learning from neighbours. *The Review of Economic Studies*, *65*(3), 595–621.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, *107*(3), 797–817.

Banerjee, A., & Fudenberg, D. (2004). Word-of-mouth learning. *Games and Economic Behavior*, *46*(1), 1–22.

Barriere, V., Clavel, C., & Essid, S. (2017). Opinion dynamics modeling for movie review transcripts classification with hidden conditional random fields. In *Proceedings of Interspeech, 2017* (pp. 1457–1461).

Bernardes, A. T., Stauffer, D., & Kertész, J. (2002). Election results and the Sznajd model on Barabasi network. *European Physical Journal B*, *25*, 123–127.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*(5), 992–1026.

Bindel, D., Kleinberg, J., & Oren, S. (2015). How bad is forming your own opinion? *Games and Economic Behavior*, *92*, 248–265.

Callen, E., & Shapero, D. (1974). A theory of social imitation. *Physics Today*, *27*(7), 23–28.

Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, *81*, 591–646.

Chen, S., Glass, D. H., & McCartney, M. (2016). Characteristics of successful opinion leaders in a bounded confidence model. *Physica A: Statistical Mechanics and Its Applications*, *449*, 426–436.

Chen, X. D., Wu, Z., Wang, H., & Li, W. (2017). Impact of heterogeneity on opinion dynamics: Heterogeneous interaction model. *Complexity*, *2017*, 5802182:1–5802182:10.

Choi, S. (2015). The two-step flow of communication in twitter-based public forums. *Social Science Computer Review*, *33*(6), 696–711.

Deffuant, G. (2006). Comparing extremism propagation patterns in continuous opinion models. *Journal of Artificial Societies and Social Simulation*, *9*(3), 8.

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, *3*, 87–98.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, *69*(345), 118–121.

DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, *118*(3), 909–968.

Di Mare, A., & Latora, V. (2007). Opinion formation models based on game theory. *International Journal of Modern Physics C*, *18*(09), 1377–1395.

Dong, Y., Ding, Z., Martínez, L., & Herrera, F. (2017). Managing consensus based on leadership in opinion dynamics. *Information Sciences*, *397–398*, 187–205.

Dong, Y., Zha, Q., Zhang, H., Kou, G., Fujita, H., Chiclana, F., & Herrera-Viedma, E. (2018). Consensus reaching in social network group decision making: Research paradigms and challenges **(in press)**.

Douven, I., & Riegler, A. (2010). Extending the Hegselmann–Krause model I. *The Logic Journal of the IGPL*, *18*(2), 323–335.

Duggins, P. (2017). A psychologically-motivated model of opinion change with applications to American politics. *Journal of Artificial Societies and Social Simulation*, *20*(1), 13.

Ellison, G., & Fudenberg, D. (1995). Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, *110*(1), 93–125.

Etesami, S. R., & Başar, T. (2015). Game-theoretic analysis of the Hegselmann–Krause model for opinion dynamics in finite dimensions. *IEEE Transactions on Automatic Control*, *60*(7), 1886–1897.

French, J. R. P. (1956). A formal theory of social power. *Psychological Review*, *63*, 181–94.

Galam, S. (1997). Rational group decision making. A random field ising model at $t = 0$. *Physica A: Statistical Mechanics and Its Applications*, *238*, 66–80.

Galam, S. (2002). Minority opinion spreading in random geometry. *The European Physical Journal B*, *25*(4), 403–406.

Galam, S. (2017). The Trump phenomenon: An explanation from sociophysics. *International Journal of Modern Physics B*, *31*, 1742015.

Galam, S., Feigenblat, Y. G., & Shapir, Y. (1982). Sociophysics: A new approach of sociological collective behaviour. I. mean-behaviour description of a strike. *The Journal of Mathematical Sociology*, *9*(1), 1–13.

Gargiulo, F., & Mazzoni, A. (2008). Can extremism guarantee pluralism? *Journal of Artificial Societies and Social Simulation*, *11*(4), 9.

González, M. C., Sousa, A. O., & Herrmann, H. J. (2004). Opinion formation on a deterministic pseudo-fractal network. *International Journal of Modern Physics C*, *15*, 45–57.

Harary, F. (1959). A criterion for unanimity in french's theory of social power. In D. Cartwright (Ed.), *Studies in social power* (pp. 168–182). Oxford: University of Michigan.

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, *5*(3). http://jasss.soc.surrey.ac.uk/5/3/2.html.

Hegselmann, R., & Krause, U. (2005). Opinion dynamics driven by various ways of averaging. *Computational Economics*, *25*, 381–405.

Hegselmann, R., & Krause, U. (2006). Truth and cognitive division of labour first steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, *9*(3). http://jasss.soc.surrey.ac.uk/9/3/1.html.

Hegselmann, R., & Krause, U. (2015). Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model. *Networks and Heterogeneous Media*, *10*, 477–509.

Holley, R. A., & Liggett, T. M. (1975). Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, *3*(4), 643–663.

Jadbabaie, A., Molavi, P., Sandroni, A., & Tahbaz-Salehi, A. (2012). Non-Bayesian social learning. *Games and Economic Behavior*, *76*(1), 210–225.

Katz, E. (1957). The two-step flow of communication: An up-to-date report on an hypothesis. *The Public Opinion Quarterly*, *21*(1), 61–78.

Klamser, P. P., Wiedermann, M., Donges, J. F., & Donner, R. V. (2017). Zealotry effects on opinion dynamics in the adaptive voter model. *Physical Review E*, *96*, 052315.

Krause, U. (2000). A discrete nonlinear and non-autonomous model of consensus formation. *Communications in Difference Equations*, *2000*, 227–236.

Liang, H., Yang, Y., & Wang, X. (2013). Opinion dynamics in networks with heterogeneous confidence and influence. *Physica A: Statistical Mechanics and its Applications*, *392*(9), 2248–2256.

Lorenz, J. (2010). Heterogeneous bounds of confidence: Meet, discuss and find consensus!. *Complexity*, *15*(4), 43–52.

Martins, A. C. R. (2008a). Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics C*, *19*, 617–624.

Martins, A. C. R. (2008b). Mobility and social network effects on extremist opinions. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, *78*, 036104.

Martins, A. C., de Pereira, C. B., & Vicente, R. (2009). An opinion dynamics model for the diffusion of innovations. *Physica A: Statistical Mechanics and Its Applications*, *388*(15), 3225–3232.

Martins, A., & Kuba, C. D. (2010). The importance of disagreeing: Contrarians and extremism in the CODA model. *Advances in Complex Systems*, *13*(5), 621–634.

Mobilia, M., Petersen, A., & Redner, S. (2007). On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, *08*, P08029.

Mobius, M., & Rosenblat, T. (2014). Social learning in economics. *Annual Review of Economics*, *6*(1), 827–847.

Molavi, P., Tahbaz-Salehi, A., & Jadbabaie, A. (2018). A theory of non-Bayesian social learning. *Econometrica*, *86*(2), 445–490.

Moussaïd, M., Kämmer, J., Analytis, P., & Neth, H. (2013). Social influence and the collective dynamics of opinion formation. *PloS One*, *8*, e78433.

Oster, E., & Feigel, A. (2015). Prices of options as opinion dynamics of the market players with limited social influence. ArXiv e-prints p arXiv:1503.08785.

Pineda, M., & Buendía, G. (2015). Mass media and heterogeneous bounds of confidence in continuous opinion dynamics. *Physica A: Statistical Mechanics and Its Applications*, *420*, 73–84.

Roch, C. (2005). The dual roots of opinion leadership. *Journal of Politics*, *67*, 110–131.

Ruf, S. F., Paarporn, K., Pare, P. E., & Egerstedt, M. (2017). Dynamics of opinion-dependent product spread. In *2017 IEEE 56th annual conference on decision and control (CDC)* (pp. 2935–2940).

Schulze, C. (2003). Advertising in the Sznajd marketing model. *International Journal of Modern Physics C*, *14*(01), 95–98.

Smith, L., & Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica*, *68*(2), 371–398.

Sobkowicz, P. (2009). Modelling opinion formation with physics tools: Call for closer link with reality. *Journal of Artificial Societies and Social Simulation*, *12*(1), 11.

Stauffer, D. (2003). How to convince others? Monte Carlo simulations of the Sznajd model. In J. E. Gubernatis (Ed.), *The Monte Carlo method in the physical sciences* (Vol. 690, pp. 147–155). American Institute of Physics Conference Series.

Sun, R., & Mendez, D. (2017). An application of the continuous opinions and discrete actions (CODA) model to adolescent smoking initiation. *PloS One*, *12*, e0186163.

Sznajd-Weron, K. (2005). Sznajd model and its applications. *Acta Physica Polonica B*, *36*, 2537.

Sznajd-Weron, K., & Sznajd, J. (2000). Opinion evolution in closed community. *International Journal of Modern Physics C*, *11*, 1157–1165.

Sznajd-Weron, K., & Weron, R. (2003). How effective is advertising in duopoly markets? *Physica A: Statistical Mechanics and Its Applications*, *324*(1), 437–444.

Takács, K., Flache, A., & Maes, M. (2016). Discrepancy and disliking do not induce negative opinion shifts. *PloS One*, *11*, e0157948.

Verma, G., Swami, A., & Chan, K. (2014). The impact of competing zealots on opinion dynamics. *Physica A: Statistical Mechanics and its Applications*, *395*, 310–331.

Waagen, A., Verma, G., Chan, K., Swami, A., & D'Souza, R. (2015). Effect of zealotry in high-dimensional opinion dynamics models. *Physical Review E*, *91*, 022811.

Weidlich, W. (1971). The statistical description of polarization phenomena in society. *British Journal of Mathematical and Statistical Psychology*, *24*(2), 251–266.

Weidlich, W. (1972). The use of statistical models in sociology. *Collective Phenomena*, *1*(1), 51–59.

Yildiz, E., Ozdaglar, A., Acemoglu, D., Saberi, A., & Scaglione, A. (2013). Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, *1*(4), 19:1–19:30.