The role of recommender systems in the formation of internet echo chambers

by

Daniël Salomons
d.salomons@students.uu.nl

Under supervision of
Dr. Gabriele Keller
g.keller@uu.nl

Second evaluator
Dr. Domink Klein
d.klein@uu.nl

2 July 2021
7.5 EC Bachelor's Thesis
Artificial Intelligence
Department of Humanities
Utrecht University

Contents

1 Abstract

On social media, users are free to engage with anybody they feel inclined to. However, echo chambers may be undesirable for internet discourse. *Echo chambers* are situations where users are in an environment where the same opinion gets reinforced by repeating it without due criticism or exposure to the opposite opinion. The research question this thesis addresses is how echo chambers form on a hypothetical social media website. The hypothesis is that a stronger opinion reinforces a grouping tendency with those of the same opinion. This group then reinforces the opinion which in turn reinforces the grouping tendency, keeping users stuck in an echo chamber unless external factors take these users out or dissolve these groups entirely. This paper models user behaviour on a hypothetical social media website to research which situations lead to internet echo chamber formation. This model differs from others in that it specifically models user movement and grouping behaviour. Experiments on the model included simulating social mobility and utilizing the recommender system to serve content of the opposing opinion. The results show that both social mobility and serving the opposite opinion work as measures against echo chamber formation.
Keywords: *Agent-based modelling, Opinion dynamics, Echo chambers, Polarisation*

2 Introduction

Social media can function as a platform for people to share their thoughts and ideas to the world. This freedom does not come without some observed drawbacks. For instance, users are free to exclusively interact with news sources and other users that share their own opinion. In that situation, the user can be considered to be in an *echo chamber*: an environment where the same opinion gets reinforced by repeating it without due criticism or exposure to the opposite opinion. Some platforms facilitate this phenomenon more than others (Cinelli et al. 2021), though the reinforcing effect lessens if users sufficiently vary whom they talk to and where they get their news from (Dubois and Blank 2018). Group conversation already polarises the opinion of group participants to more extremes in an effect called *group polarisation* (Sunstein 1999).

Additionally, an effect called *biased assimilation* also plays a role in the formation of echo chambers. People are likely to accept evidence that conforms to their own opinion. People consider opinion conforming evidence to be more convincing than evidence against their opinion, scrutinizing opposing evidence with more critical evaluation than conforming evidence. Opinions become more polarised if the people are sufficiently biased in this manner (Lord, Ross, and Lepper 1979).

Biased assimilation occurs more frequently on social media platforms due to the websites using recommender systems. Many modern content serving platforms utilize recommender systems in order to boost user engagement with the platform (Wu et al. 2017). The algorithms of these recommender systems create *user profiles* that consist of data regarding the user's demographics, and also more specific data such as what content the user engaged with on the website (Li and Kim 2004). In turn, the recommender systems improve and can serve content that the user is more likely to interact with. The content served by this system can polarise a user. When the user is polarised, they are more likely to interact with such polarising content because they care about it, which in turn influences the algorithm, creating a feedback loop (Jiang et al. 2019).

This paper explores the interactions between recommender systems and agent polarisation. The research question is as follows: "How are echo chambers formed on a social media website?". The relevance to AI as a field is to research the ramifications of the usage of recommender systems in terms of the formation of echo chambers. The hypothesis is divided in two parts. (1) Once an agent cares more about a subject, they are going to interact with others who care more about the subject, therefore those agents will group up. Recommender systems play into this as described (Jiang et al. 2019). As a result, agents will be stuck in an echo chamber. (2) External influences may reduce that effect, remove the agent from the echo chamber or dissolve the group entirely.

The hypothesis is explored with a computational agent-based model created with NetLogo, which I call the *user grouping model.* Agent-based models are not new in the field of social sciences (Klein, Marx, and Fischbach 2018). Agent-based modelling allows for the opportunity to research emergent group behaviour by giving simple instructions to individual agents and observing the

effects when these instruction parameters change.
An existing model of opinion dynamics is the psychologically motivated model of opinion change (Thomson 2019), where agents do not always express their true opinion, have a variety of weights to affect the influence of incoming opinions, and have a static group of other agents to talk to. It is the third mentioned property of the model that can be explored further: on the internet, agents do not necessarily need to talk to the same group of agents. By reducing the complexity of the first two mentioned properties of the model, the user grouping model seeks to isolate effects of movement and grouping dynamics. Pilditch's opinion cascade model holds the same assumption that networks remain static after formation (Pilditch 2017).
Another influential model in opinion dynamics is the bounded confidence model (Hegselmann and Krause 2002). In that model, agents share their true opinion with each other, and have various weights attached to the influence of incoming opinions of other agents. In the bounded confidence model, agents only consider opinions that do not differ from their own opinions within a certain range. The user grouping model is similar to the bounded confidence model because agent interaction in the user grouping model is indirectly determined by opinion distance. However, the user grouping model still allows for the stochastic discovery of opposing opinions, as opposed to the cut-off thresholds of the bounded confidence model. Similar assumptions are found in relative agreement model as it was built off the bounded confidence model (Deffuant et al. 2002)
Other influential agent-based models on opinion dynamics include (Salzarulo 2006; Jager and Amblard 2005; Mark 1998; Carley 1991; DeGroot 1974).

In this paper I made a distinction between the term 'agent' and 'user'. Agents are part of the model, while the term user is used to refer to assumptions about real internet user behaviour.

3 Model Assumptions

This section highlights the assumptions that went into creating the user grouping model. The section is divided into two subsections. The first subsection talks about the assumptions of the base model to research the first part of my hypothesis: a grouping tendency, paired with the influence of a recommender system, will get agents stuck in an echo chamber. The second subsection talks about the assumptions of the extensions built into the model to consider additional realistic situations. It also explores the second part of the hypothesis, that echo chambers persist until disrupted by external factors. The base model assumptions and functionality remain the same unless specified otherwise.

3.1 Base model assumptions

The base model assumes that an agent's opinion is influenced by whomever the agents talk to. There are no varying weights for incoming opinions, to provide a new perspective to opinion dynamics.

Opinions. This model has agents on a social media platform with an opinion o_t^i at time t about a singular topic. Through discourse with other agents, an agent's opinion shifts stronger towards the opinion of the agents around them. A stronger opinion of the interlocutor increases the grouping tendency of the agent, which makes that agent more likely to interact with like-minded agents. All agents have their own measure of susceptibility s_i. This susceptibility limits how much an agent's opinion may shift from their initial opinion o_0^i. Not every user radicalises nor will necessarily care about the subject as much. Therefore, in the model it is rarer to be inclined towards more extreme opinions than more neutral opinions.

Grouping. An increase in opinion makes an agent more inclined to group up. Next to that is a variable g_g set by the modeler. g_g is an abstraction of the interaction between agents and the site's recommender system. This model opted for simplicity as it perfectly guides the agent to their respective groups. In a realistic situation, users are able to ignore website recommendations and move around the website as they please. A more practical study on the effect of recommender systems on user-base polarisation is to change the currently used algorithms and measure changes in polarisation through surveys. The social media website Facebook is known for conducting studies on their user-base in this manner (Kramer, Guillory, and Hancock 2014). However, due to time limitations and the fact that these algorithms are mostly proprietary, this falls beyond the scope of the model of this thesis.

On the internet, agents generally know that a certain group shares their own opinion, either through notoriety of a group, or more simply, by its name (For instance, a group with 'MAGA' in its name supports Donald Trump) and the user can steer clear or engage accordingly. It is within these groups that polarisation occurs. Once surrounded by those who share their opinion, an agent's opinion degenerates to a more extreme variant rather quickly, modelling the findings of the law of group polarisation (Sunstein 1999). Agents prefer to seek out groups that hold a similar opinion to their own. Users with extreme opinions presumably care a lot about the subject and might naturally be more selective about who they talk with. Moderate users might shy away from the intensity of the opinions held by extremists and as such avoid them, naturally selecting their own crowd.

3.2 Model extension assumptions

All these assumptions mentioned so far were for the base model. Next, I will discuss the assumptions of the model extensions and experiments.

Opinion-altering events. Opinions about a subject do not solely exist within discourse: they are inspired by real-world events actively happening around us. To represent this, two buttons exist on the interface of my NetLogo program that provide the option to introduce an "opinion shock", to observe how the situation is affected by positive or negative news about the subject.

Social mobility. In daily life, users may not be able to completely wall themselves off from those who do not share their own opinion. They may meet and be forced to interact with people with differing opinions at school or at work. However, as geography plays a large part in homophily (McPherson, Smith-Lovin, and Cook 2001), it is likely that even those you go to school and work with share your opinion, reducing the exposure to opposing opinions.

A more concrete example for what this experiment intends to simulate is as follows: Geographically, the United States is a very politically segregated country. Gallup (Gallup 2017) found that within their sample size that the District of Columbia has a 70% Democratic leaning and 11% Republican leaning. A user who rarely leaves this state would be less exposed to differing opinions than one who would regularly visit Wyoming with a 56% Republican leaning and 27% Democratic leaning. In the model, this phenomenon is modelled by having agents 'teleport' varying distances around the website.

Site-wide preventative measures. Additionally, websites may be pressured by advertisers or investors to not facilitate as a platform for extremism. One of the most well-known situations of this occurring was in 2017 known as YouTube's 'adpocalypse', as companies complained that their advertisements were showing up next to offensive content. YouTube in turn changed its policies regarding moderating its content (Thomson 2019). To simulate these tactics against hate speech and extremism, two options are included in the model. One simple solution is to outright ban users exhibiting hateful or extremist views.

Another option for a website to curb radicalisation in its user-base is to utilise the website's recommender system to guide extreme users towards discussions of other users that do not share their opinion.

4 User Grouping Model

In the model a social network website is represented as a torus where agents move around as they browse the website. Space on this torus represents a topic available on the website or a sub-community. The torus is the standard topology supported by NetLogo.

4.1 Base model

Opinion. An agent's location on the torus determines which other agents they interact with, as they only interact with other agents within a certain constant radius $r = 3$ in euclidean distance. Agent i at time-step t has an opinion o_t^i, represented by a real number bounded between -1 and 1. Opinions close to 0 are neutral or undecided, while opinions close to either bound are extremes.

An agent's opinion is updated every time-step of the model:

$$o_{t+1}^i = \begin{cases} o_t^i + 0.1\overline{O}_t^i & \text{if } |o_{t+1}^i - o_0^i| \leq s_i \\ o_t^i & \text{otherwise} \end{cases}$$

where $\overline{O}_t^i$ the average opinion of other agents within the range $r = 3$ in euclidian distance around agent i at time t. The 0.1 is a weight attached to this average

Description	Variable	Value	B	M	ND
Opinion of agent i at time t	o_t^i	$[-1, 1]$	✓		✓*
Global agent interaction range	r	3	✓		
Global opinion weight	-	0.1	✓		
World size in NetLogo patches	-	16x16	✓		
Grouping inclination of agent i	g_i	$g_g\|o_t^i\|$	✓		
Susceptibility of agent i	s_i	$[0, 2]$	✓		✓
Average opinion around agent i	$\overline{O}_t^i$	$[-1, 1]$	✓		
Average positive, negative opinion	$\overline{O}_p, \overline{O}_n$	$[-1, 1]$	✓		
Average model polarisation	P	$\overline{O}_p - \overline{O}_n$	✓		
Global agent grouping odds	g_g	$[0, 0.5]$	✓	✓	
Global agent grouping distance	-	$g_i + r$	✓		
Percentile of socially mobile agents	t_u	$[0, 1]$		✓	
Global 'teleportation' odds	t_g	$[0, 1]$		✓	
'Teleportation' distance	-	$[0, 16]$			✓
Global opposite recommender odds	o_g	$[0, 0.5]$		✓	

Table 1: A quick overview of relevant variables in the user grouping model. **B** stands for whether the variable is part of the base model. **M** stands for whether the variable is varied for experiments by the modeler. **ND** stands for whether the value is normally distributed. *Agent opinions are initialised with a normal distribution but iterate as described in section 4.1.

opinion, as to not change an agent's opinion too quickly. The opinion of agent i only changes if the new opinion o_{t+1}^i is within the susceptibility of an agent: $o_0^i - s_i \leq o_t^i \leq o_0^i + s_i$, where s_i is a certain susceptibility value inherent to agent i. In the model this manifests as s_i being an absolute normally distributed value.

Movement and grouping. Agent i groups up with a rate of $g_i = g_g|o_t^i|$, where g_g is a global grouping variable between 0 and 0.5 set by the modeler before each experiment. Each time-step, the agent either takes a 1 euclidean distance step in a random direction at odds $1 - g_i$ or a step towards a similarly opinionated group at odds g_i. The grouping implementation is that each agent looks a distance $g_i + 3$ ahead and measures the average opinion of that location. The agent then moves towards the location where the average opinion differs the least from their own. $|o_t^i|$ is representative of an agent's inherent bias towards homophily, while g_g represents the suggestions of a recommender system, amplifying the inherent bias. A high g_g indicates a large influence of recommender systems pushing them towards similarly inclined groups.

4.2 Model extensions

Opinion-altering events. When a shock happens, 50% of the population updates their opinion up to ± 0.5 (less if the resulting opinion would exceed -1 or 1). This percentage was chosen to emulate that not necessarily everyone has heard of, or cares about the event.

Social mobility. The model has an option to teleport a subset of agents t_u to a different location every time-step. Upon population generation, a subset of agents can be specified to be part of the group with a high social mobility. This is a slider from 0 to 1, where 1 represents 100% of the population being highly mobile. Each time-step, each agent belonging to the mobile group decides whether they teleport at odds t_g determined by the modeler. This would place the agents away from their group and possibly make them interact with those of the opposite opinion. In the model, the location they are teleported to is random: an agent faces a random direction and moves a distance forward. The distance the agent moves is unevenly distributed making lower distances more likely than higher distances, to emulate the geographic homophily.

Site-wide preventative measures. In the model the option is given to outright remove agents with $|o_t^i| \geq 0.95$, representing banning radicalised agents breaking website terms of service by posting extremist content.
Additionally, there exists an option that inverses the agent's opinion for the grouping algorithm at odds o_g, sending the agent to the location with the least corresponding average opinion, rather than most. This represents a recommender system suggesting content of the opposite opinion rather than a corresponding opinion, aiming to deradicalise the user-base or to break up existing echo chambers.

5 Results and discussion

The model measures the average opinion of all agents whose opinion is $o_t^i > 0$, or $\overline{O}_p$, and the average opinion of all agents whose opinion is $o_t^i < 0$, or $\overline{O}_n$. The difference between these measures can be seen as a measure of polarisation within the website: $P = \overline{O}_p - \overline{O}_n$. This measure is the dependent variable of all the tests. In the following graphs, the simulation is run with a certain parameter set for 1000 time-steps, 100 times per parameter set. Each run the final value of P is recorded. The 100 values per parameter set are then depicted in a boxplot to visualise and discuss the results.

In figure 1, the baseline results of the user grouping model are depicted: the change in polarisation P in regard to the change in grouping g_g. It can be observed that a higher tendency to group up leads to a higher polarisation level between the two opinions. In figure 2, these results are replicated within a single run as g_g increases linearly over time. This polarising phenomenon occurs because agents are quickly led towards those who share their own opinion. By interacting, they exchange their similar opinion with each other. As a result, they intensify their own opinion in that direction. Neutral agents are either drawn to these opinionated groups due to their own slight leaning towards that opinion or move towards these agents at random and then adopt their opinion, becoming part of the group. This shows the findings that if interactions occur

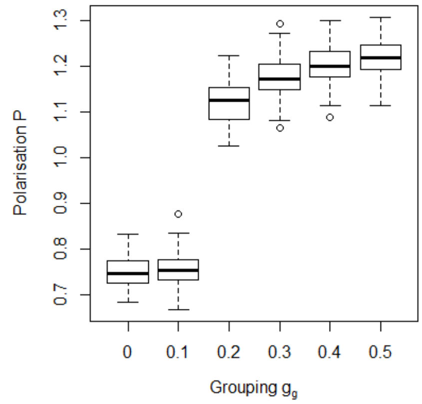

Figure 1: Boxplot of P over 100 runs of various values for g_g

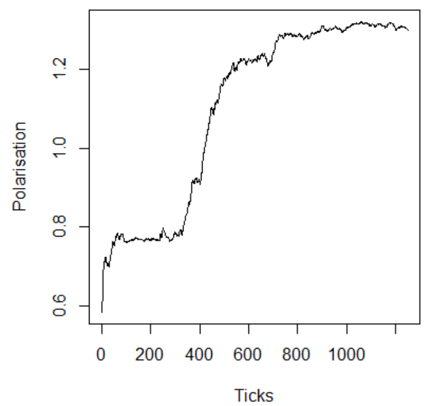

Figure 2: Increasing grouping g_g linearly with time during 1 run

within a group, the opinions of the group participants shift to a more extreme variant of the original (Sunstein 1999).

In the user grouping model, whenever an agent's opinion is intensified because of a group, that agent is likely to stay in that group. This behaviour mirrors the phenomenon where recommender systems keep users within their separate communities on the same website.

One interesting phenomenon is that it does not take much grouping-up inclination to create a high polarisation. At $g_g = 0.2$, clear groups already form within the community and polarisation increases dramatically.

This result confirms the first part of the hypothesis. Agents decide whether they want to group up at rate $g_i = g_g|o_t^i|$. At sufficiently high g_g, the value determining the influence of a recommender system, an increase in opinion increases their grouping which in turn increases their opinion, keeping them stuck into an echo chamber. Without external factors, these echo chambers persist. The radical difference between $g_g = 0.1$ and $g_g = 0.2$ in terms of polarisation is that in this simple situation, it comes down to whether echo chambers form or not within the allotted time. At a low g_g, echo chambers do not get a chance to form, keeping the opinion of the population moderate. At a high g_g, echo chambers form and remain, drastically increasing polarisation.

5.1 Extensions

I will now present and discuss the results of the extensions to the base model as described earlier. This pertains to the second part of my hypothesis about disruptions to echo chambers. In the experiments of the model extensions, the base functioning of the model remains the same with $g_g = 0.5$. Additionally, the effect of each of these extensions are tested one at a time.

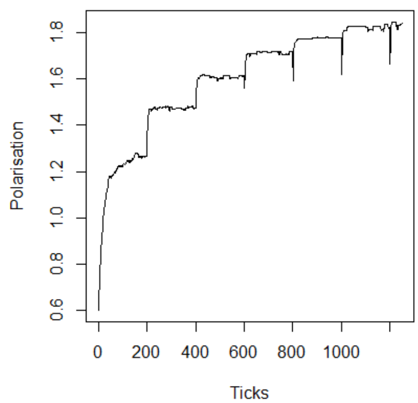

Figure 3: Polarisation as a negative shock happens every 200 time-steps

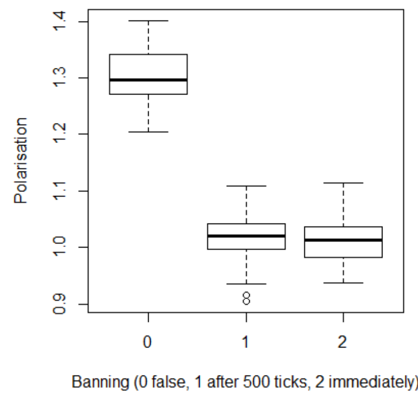

Figure 4: The effect of banning radical agents ($|o_t^i| \geq 0.95$) at $g_g = 0.5$

Opinion-altering events. Figure 3 shows polarisation over time. In this special run, half the agents (chosen at random) are shocked with a negative opinion every 200 time-steps. The resulting behaviour is that neutral agents immediately turn negative about the subject and as such flock towards the now-majority group. However, agents that were previously rooted in a very positive opinion absorb this news but are quickly pulled back to their original positive opinion by their surrounding peers. This mirrors real behaviour how users would discredit news that goes against their own paradigm. As the run goes on, negative agents absorb the negative shocks and further move towards a stronger negative opinion. The positively opinionated agents however remain rooted in their positive position. This increases the polarisation. In the second half of the plot, the spikes in reduction of polarisation are visible as the positive agents absorb the negative shock, but the correction immediately follows to a higher polarisation level than the level before the shock. Ultimately, shocks fail to disrupt echo chambers but rather instigate polarisation.

Radical agent banning. Figure 4 displays three boxplots of 100 runs each. The first boxplot depicts the average polarisation at $g_g = 0.5$ without banning radical agents, which are agents whose absolute opinion exceeds 0.95 ($|o_t^i| \geq 0.95$). The second boxplot shows the impact of banning these radical agents after letting them interact for another 500 time-steps. The third boxplot displays a run where radical agents were banned as soon as they reach or exceed the 0.95 threshold. These bans are permanent as the agents are removed from the simulation entirely. While a notably large difference in polarisation is visible, the insignificant difference between banning radical agents straight away or with a delay suggests that this model does not support the investigation of this issue. The average difference only lies in removing these agents from the calculation altogether. Investigating this issue might be interesting for future work.

Social mobility. Figure 5 depicts the results of an experiment where agents teleport at multiple global teleportation odds t_g at $g_g = 0.5$. In figure 5 specifically, the mobile group is 100%, or $t_u = 1$. The polarisation significantly reduces when these odds increase. At high teleportation odds ($t_g \geq 0.2$), user grouping inclinations become significantly less important as movement in the model is reduced to a random walk.
The assumption that all users are highly mobile is arguably inaccurate. In figure 6, 7 and 8 the results are shown of experiments where the number of agents that teleport is varied from 0% to 100%. In these experiments the odds of teleporting consists of 10%, 20% and 30% respectively. These figures show that when a larger population is increasingly mobile, polarisation reduces in the model. In the model itself, when a sufficiently large percentage of the population is socially mobile and 'teleport' frequently, a lonely few remain within echo chambers holding strong opinions. The teleportations however fail to break up these echo chambers. The echo chambers do have a lower group size than when t_u is high. Additionally, all members of these remaining echo chambers are not part of the mobile group of agents. The resulting conclusion is that those who do not have a high mobility within society fail to be exposed to sufficient different opinion to break up their echo chamber.

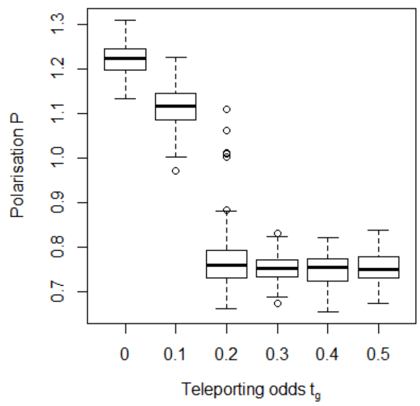

Figure 5: Average polarisation for various values of teleporting t_g

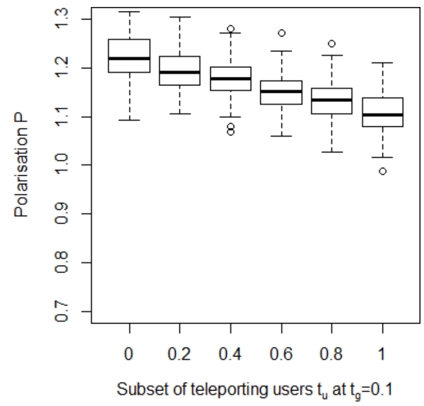

Figure 6: Increasing the mobile group t_u at $t_g = 0.1$

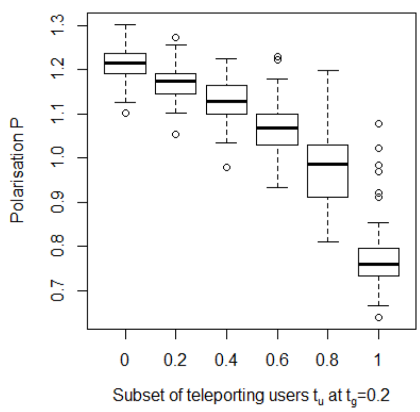

Figure 7: Increasing the mobile group t_u at $t_g = 0.2$

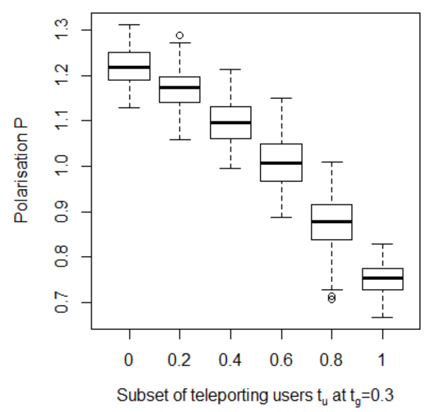

Figure 8: Increasing the mobile group t_u at $t_g = 0.3$

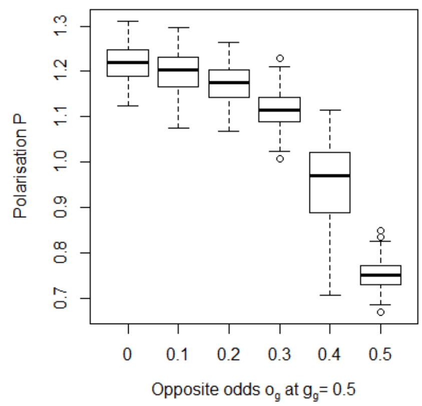

Figure 9: Opposite pathfinding with odds o_g with grouping $g_g = 0.5$

Recommending opposite content. Figure 9 shows the results of the model's implementation of utilising the recommender system to send agents towards those who hold the opposite opinion of their own. As the global odds of sending an agent of the opposite opinion o_g increase, polarisation P on the website decreases. In the model, this adjustment to agent's behaviour causes them to move around the edges of their groups rather than the centre, increasing their exposure to those who do not share their own opinion. Aside from this, polarisation reduces for two other reasons.
The first: once enough agents of one opinion enter the echo chamber of the opposite opinion and share their beliefs, the echo chamber dissolves, causing a spontaneous drop in polarisation. The odds of echo chamber dissolution are higher if o_g is higher.
The second: at high levels of o_g agent behaviour is reduced to a random walk despite every agent's inclination to group up. Random walks completely prevent the formation of echo chambers, but in a real-world application that would happen at the expense of the user's experience as the recommender algorithm recommends a lot of content irrelevant to the user's interests. This makes it an undesirable implementation for both website and user. As such, if websites seek to utilise this alternative situation, their recommender system should only recommend content with the opposite opinion 20% to 30% of the time to still benefit from the polarisation reducing effect while keeping irrelevant suggestions to a minimum.

5.2 Future work

The user grouping model provides a very simplistic view of user interaction on a single website. This simplicity was chosen in order to create and study emergent behaviour of the users especially in relation to echo chamber formation.

However, there are certain things that could be expanded upon for future works. In the user grouping model agent interaction is solely defined by location. Figure 4 depicts the results of an experiment that the user grouping model failed to show. In that experiment, agents with extreme opinions ($|o_t^i| \geq 0.95$) were removed from the simulation, as they were banned from the hypothetical social media website. However, this produced a polarisation drop only because these agents' opinions were simply removed from the average calculation. A future work could implement a certain interaction success rate where more extreme agents have less of a chance to spread their extremist views as they attempt to subvert a website's content moderation system.
Additionally, for simplicity, the user grouping model studies the effects on a singular opinion rather than a multitude of opinions. Future works could employ an opinion vector that would increase the depth of the model, as in real life users generally express an opinion on more than one topic during an interaction.
Furthermore, grouping in the user grouping model is very one-dimensional because agents can only be part of a single group at a time. The model can be expanded so that agents can be part of multiple groups at the same time. This relates with the previous idea that different groups could be centred around different opinions.

6 Conclusion

In this study I modelled the parameters of a hypothetical social media website and the interaction of its users. The research question was how echo chambers form on a hypothetical social media website.
Agent-based computer simulations found an increase in opinion resulted in closed groups that reinforce the same opinion. External factors, such as social mobility and an implementation of the recommender system to serve content differing in opinion cause a reduction of polarisation and sometimes a dissolution of echo chambers.
My conclusion of this study is if grouping based on opinion occurs in users with low social mobility, combined with recommender systems that serve content of the same opinion will result in echo chambers on social media websites. The implication of this research is that websites should be careful about aggressively utilising their recommender system. While it may boost user engagement, the results might not be desirable in terms of avoiding user polarisation. If websites facilitate a certain level of stochastic exploration within a user's news feed, the websites are less likely to create echo chambers. Users themselves would be wise to seek a healthy amount of diversification in both content they consume and with whom they interact with. This can be done by increasing social mobility if possible or simply seeking out users of the opposing opinion to have healthy, constructive discussions with.

References

[1] Kathleen Carley. "A theory of group stability". In: *American sociological review* (1991), pp. 331–354.

[2] M Cinelli et al. "The echo chamber effect on social media". In: *Proceedings of the National Academy of Sciences* 118.9 (2021).

[3] Guillaume Deffuant et al. "How can extremism prevail? A study based on the relative agreement interaction model". In: *Journal of artificial societies and social simulation* 5.4 (2002).

[4] Morris H DeGroot. "Reaching a consensus". In: *Journal of the American Statistical Association* 69.345 (1974), pp. 118–121.

[5] E Dubois and G. Blank. "The echo chamber is overstated: the moderating effect of political interest and diverse media". In: *Information, communication & society* 21.5 (2018), pp. 729–745.

[6] Gallup. *2017 U.S. Party Affiliation by State*. 2017. URL: https://news.gallup.com/poll/226643/2017-party-affiliation-state.aspx.

[7] R. Hegselmann and U. Krause. "Opinion dynamics and bounded confidence models, analysis, and simulation". In: *Journal of artificial societies and social simulation* 5.3 (2002).

[8] Wander Jager and Frédéric Amblard. "Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change". In: *Computational & Mathematical Organization Theory* 10.4 (2005), pp. 295–303.

[9] R. Jiang et al. "Degenerate feedback loops in recommender systems". In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics and Society* (2019), pp. 383–390.

[10] Dominik Klein, Johannes Marx, and Kai Fischbach. "Agent-Based Modeling in Social Science, History, and Philosophy: An Introduction". In: *Historical Social Research* 43.1 (2018), pp. 7–27.

[11] Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. "Experimental evidence of massive-scale emotional contagion through social networks". In: *Proceedings of the National Academy of Sciences* 111.24 (2014), pp. 8788–8790.

[12] Q. Li and B. M. Kim. "Constructing user profiles for collaborative recommender system". In: *Asia-Pacific web conference* (2004), pp. 100–110.

[13] CG Lord, L Ross, and MR Lepper. "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." Dutch. In: *Journal of Personality and Social Psychology* 37.11 (1979), pp. 2098–2109. DOI: 10.1037/0022-3514.37.11.2098.

[14] Noah Mark. "Beyond individual differences: Social differentiation from first principles". In: *American Sociological Review* (1998), pp. 309–330.

[15] M. McPherson, L. Smith-Lovin, and J. M. Cook. "Birds of a feather: Homophily in social networks". In: *Annual review of sociology* 27.1 (2001), pp. 415–444.

[16] Toby D Pilditch. “Opinion cascades and echo-chambers in online networks: A proof of concept agent-based model”. In: Cognitive Science Society. 2017.
[17] Laurent Salzarulo. “A continuous opinion dynamics model based on the principle of meta-contrast”. In: *Journal of Artificial Societies and Social Simulation* 9.1 (2006).
[18] C. R. Sunstein. “The law of group polarization.” In: *John M. Olin Law & Economics Working Paper* 91 (1999).
[19] P. Thomson. *Understanding YouTube Demonetization and the Adpocalypse*. June 14, 2019. URL: https://learn.g2.com/youtube-demonetization.
[20] Q. Wu et al. “Returning is believing: Optimizing long-term user engagement in recommender systems”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (2017), pp. 1927–1936.

7 Appendix A: NetLogo code

The user grouping model can be downloaded from my GitHub.
The model is programmed for, and runs best in NetLogo 6.2.0.