Opinion Dynamics of Stubborn Clustered Networks

Networks Final Project Report

Raquel Taborga & Jack Defay

December 11, 2021

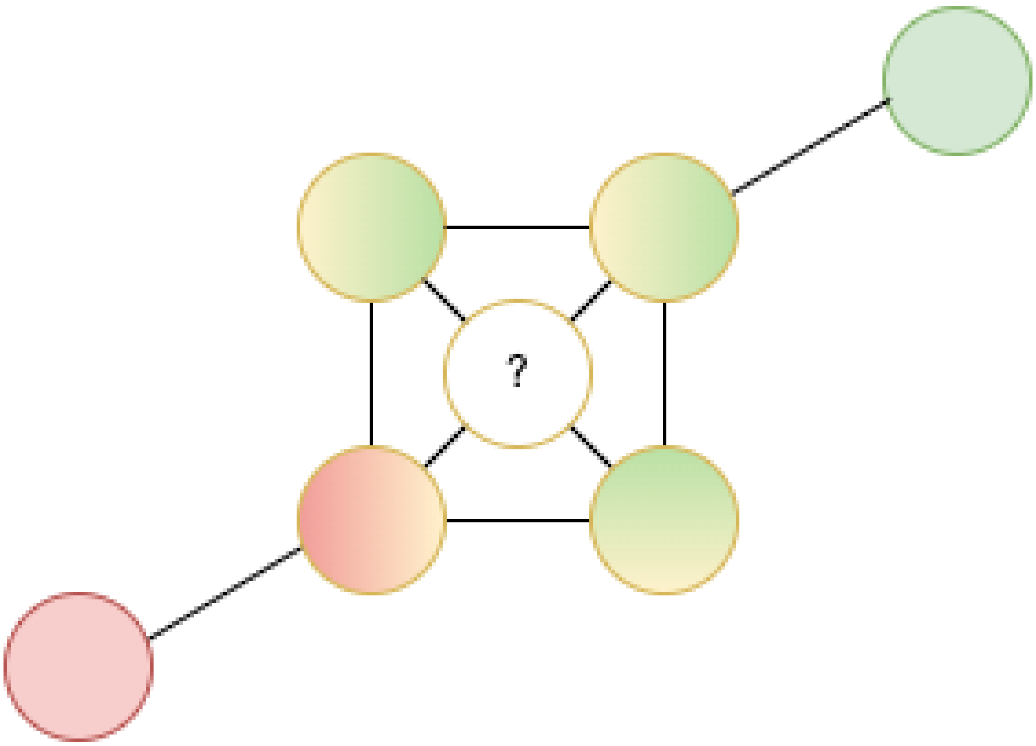

1 Introduction

Opinion dynamics is the study of networks where agents, the nodes of the network, hold some opinion, and are influenced by the agents they are connected to, the edges of the network. The weight of the edge can correspond to how strong a tie these agents have. A cluster is a subgroup of nodes within a network that share some sort of common characteristic.

In real-world opinion dynamics, we aren't just influenced by those around us, but also by our own willingness to accept new opinions and information. Some of us are simply more stubborn than others. Thus, a stubborn clustered network is a term we came up with to classify a network that contains a cluster of nodes with very low openness to changing their opinion.

This project explores under what conditions can a topological intervention change opinion dynamics in stubborn clustered networks. Although it may be difficult to make someone less stubborn, through topological change an individual can implicitly become more influenceable.

2 Background

We began our final project with a literature review on opinion dynamics and echo chambers. An Echo chamber is a social network where individuals are selected for their shared view points, intentionally or not. Echo chambers tend to polarize opinions by positive feedback and insulation from dissenting ideas. A common place we find echo chambers are in social media, in our friend groups, cults or other groups isolated from society at large, or even our own neighborhoods.

2.1 Echo-Chamber Effect

The concept of homophily in our opinion dynamics networks was analyzed through the lens of social media by Cinelli et al. where they create a model that calculates for a bias in information toward like-minded peers [2]. This makes sense, as you are probably more likely to be persuaded by a friend that already has opinions similar to you than a friend who you constantly butt heads with.

Most of the literature confirms that individuals are likely to find themselves in echo chambers, and thus, intensify their existing beliefs/opinions. However, it is possible that individuals recognize that they are in an echo chamber and intentionally choose to diversify their opinion. In a 2018 article by Dubois and Blank, these researchers found that those who are interested in politics tend to avoid echo chambers, seeing as social media allows them to diversify their opinion if they choose [4]. Another group used a "Social Mirror" where they allowed Twitter users to explore the politically-active parts of their social network to see how they would update their social network [5].

We built upon the work of existing echo chamber opinion dynamics models and findings to create our simulations. We then used our research on what agents do when they realize they are in an echo chamber to create a post-processing step which we called "intervention."

2.2 Opinion Dynamics Models

Opinion dynamics are an example of discrete time averaging systems. An averaging system is one with the equation:

$$x(k+1) = Ax(k)$$

The above linear averaging model of opinion dynamics is known as the DeGroot model [1]. In this model, the adjacency matrix, A has only non-negative entries and is row-stochastic. In our simulations, we used the *equal-neighbor* model instead of the *metropolis-hastings* model to assign edge weights in our network. This divided each row in our adjacency matrix A by the row sum evenly. In averaging system opinion dynamics, agents can reach consensus or they reach disagreement.

In order to incorporate stubbornness into the network simulation, we required an opinion dynamics model that took this into account. The Friedkin-Johnson model mentioned in the textbook *Lectures on Network Systems*, has an *influence matrix*, Λ [1]. This discrete time averaging system takes the form

$$x(k+1) = \Lambda Ax(k) + (I_n - \Lambda)x(0)$$

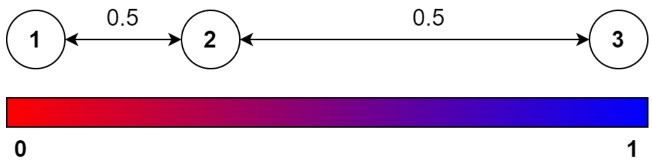

Figure 1: Illustration of the importance of the Weighted Median approach

where A is once again the adjacency matrix with dimensions $n \times n$ and Λ is a square matrix of dimensions $n \times n$ with the openness level $\lambda_{1,2 \ldots n} \in [0, 1]$ on the diagonal and zeros elsewhere. This means each agent i has a stubborness of $1 - \lambda_i$ from changing their original opinion $x(0)$.

2.3 Weighted Median

After our experiments with echo chambers according to the Friedkin-Johnsen model, we wanted to explore how different opinion dynamics models handled these extreme situations. We read about different models, but they all seemed similar, with increasing complexity that we didn't understand. Evenutally, we found a recent paper by Dr. Francesco Bullo that seemed like the perfect counterpoint to the Friedkin-Johsen model, called the Weighted Median model [6]. This model calls into question a central assumption of all other opinion dynamics models, the weighted average. Although it seems to make intuitive sense, the weighted average is not a given way for opinions to spread. The Weighted Median avoids the major pitfalls of the weighted average, and its application has interesting consequences.

3 Methods

We modularized our code to make running simulations with a simple parameter change easier. Our code is publicly available at this git repo: CLICK HERE [3]

3.1 Varying Stubbornness

Using the same network topology for each repetition, we varied the stubbornness, $1 - \lambda$ for each agent. The topology we used had the most stubborn node was also the most well connected. The opinion of the stubborn node was randomly generated from [0.9,1]. Everyone else, referred to as the masses, had opinions randomly chosen from a uniform distribution in [0,0.5]. Each agent in the masses was only connected to the most stubborn node, or cluster of nodes. With a network size of $n = 10$ nodes and a cluster size of $cs = 1$, the network is pictured in Figure2. The stubborn values were also randomly generated, where we had λ_i randomly chosen without repetition. The masses had $\lambda \in [0, 0.5]$; they had very low attachment to their original opinions, $x(0)$. The stubborn nodes had $\lambda \in [0.9, 1]$. For the core portion of this scenario's simulation, we used a network of $n = 100$ nodes with a cluster size of $cs = 1$ stubborn node.

3.2 Varying Percentage of Radical Agents

Using the same network topology as in Section 2.1, we ran simulations to determine what opinion the network converges to if we vary the percentage of a network that is stubborn. The mass had opinions and λ-values $\in [0, 0.5]$, and the cluster of stubborn nodes had initial opinions $x(0) \in [0.5, 1]$ and λ-values $\in [0.9, 1]$. The network had $n =$ 100 nodes, and we tested trials that had 1%, 10%,20%, 30%, 40%, 50% and 60% of the nodes classified as stubborn.

3.3 Intervention Between Polarized Groups

In this scenario, we wanted to give a network a bump to change the network's topology and compare the opinion dynamics before and after

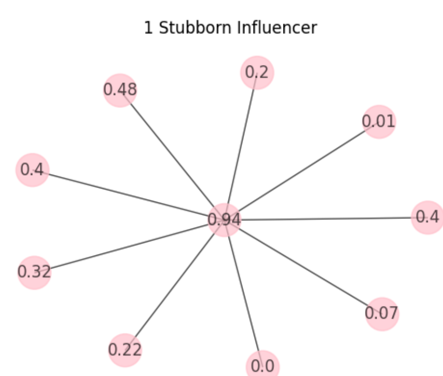

Figure 2: 10 nodes, 9 nodes comprise the mass, cluster size of 1 stubborn node. Opinions of the mass range from [0,0.5] and the stubborn node ranges from [0.9,1]

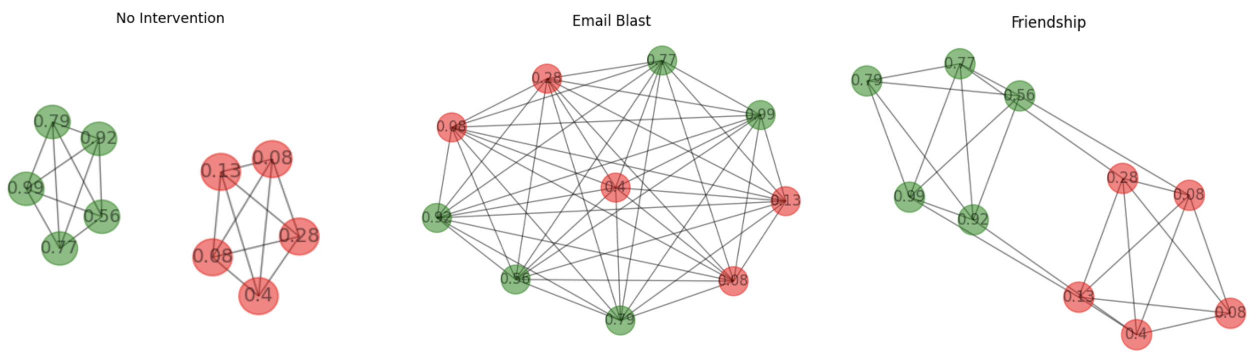

Figure 3: How intervention in polarized networks changes network topology

this post-processing. This is motivated by the question of what happens when an individual realizes they are in an echo chamber. Agents can choose to make many weak connections to agents in an opposite cluster, or they can form fewer, stronger connections in the opposite cluster. We named the weak-connections approach "email blast," much like a passive way someone might choose to diversify their reading interests by subscribing to a newsletter. We named the stronger connection approach "friendship," since new relationships and friendships have a great impact on our susceptibility to accept new opinions. The topology change is shown in Figure 3. The total amount of influence distributed across the network was the same with both methods of intervention. That is, if an agent was bumped with `friendship`, their connection with two nodes was increased by 0.5 each. If an agent was bumped by `email_blast`, they gained an edge weight of $\frac{1}{(n/2)}$ for each connection made with the opposing cluster.

Cluster opinions were created in different extremes. The default level of extremity, how radical an opinion is on a scale of 0-1, was 0.5. This meant that one cluster would have initial opinions $x(0) \in [0, 0.5]$ and the other would have $x(0) \in [0.5, 1]$, each with stubbornness $\lambda \in [0.9, 1]$.

3.4 Weighted Median Simulations

The scenario we considered for these experiments was to take the two-cluster network in Figure 3 and apply the two interventions email blast and friendship. This allowed us to investigate both how the different interventions behaved under different assumptions, and give us some insight into the consequences of the Weighted Median model.

Generally, we found that the Weighted Median model converged under looser topological constraints, especially to a well connected cluster. We also noticed that this model seems to have converted continuous opinions into categories, which doesn't quite line up with our understanding of opinions in real life. This is also discussed in the paper, and we are excited to see how this work develops in the coming years. We also noticed that this model often results in these cycles where agents swap opinions in some sort of temporary steady state. Due to the random nature of how we implemented the Weighted Median however, with enough time these almost always died out.

4 Results

4.1 Varying Stubbornness

Figure 4 shows the results of varying stubbornness λ of one agent, $cs = 1$, in a network of $n = 100$ nodes, with the topology shown in Figure 2. The stubborn node had a starting opinion in [0.9,1].

4.2 Varying percentage of radical agents

Figure 5 shows the standard deviation distribution of final opinions for a network of 100 nodes varying the percentage of stubborn nodes. We used the metric of standard deviation to categorize how diverse the convergent

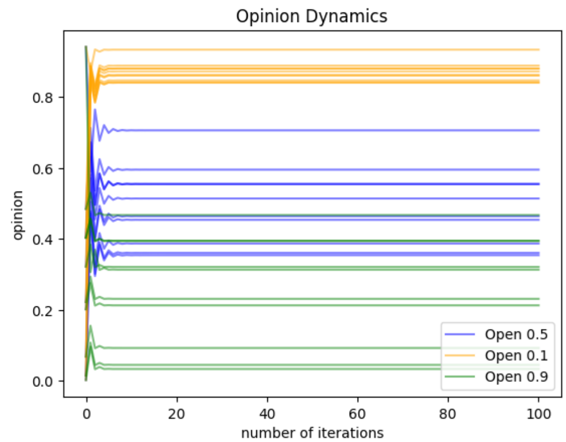

Figure 4: Resulting opinion dynamics from model with network of size n=1000. Each color represents new trial with different level of openness of stubborn cluster.

opinions of a network are. The results show that starting at around 20% influence, you get similar distributions of opinions until the tipping point of 50%.

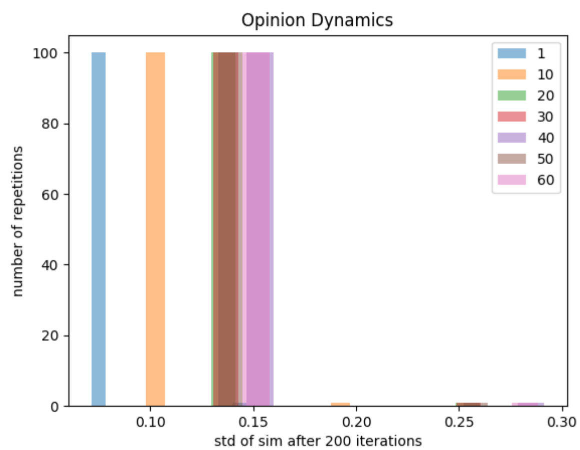

Figure 5: Standard deviation distribution for trials varying the percentage of stubborn agents in a network

4.3 Varying interventions

The results of this simulation can be found in Figure 6. Under the conditions of this simulation, the email blast had the greater harmonizing effect between the two groups. The friendship intervention was also effective at this level, but not quite as effective as the email blast.

4.4 Weighted Median Simulations

A single repetition of the Weighted Median simulation is shown in Figure 7A and the bulk standard deviation data can be found in Figure 8B.

5 Deterministic Weighted Median Simulations

To address the different implementations of the Weighted Median model, we reran the experiments presented on the randomized Weighted Median model with a deterministic version. The change to the code is shown in

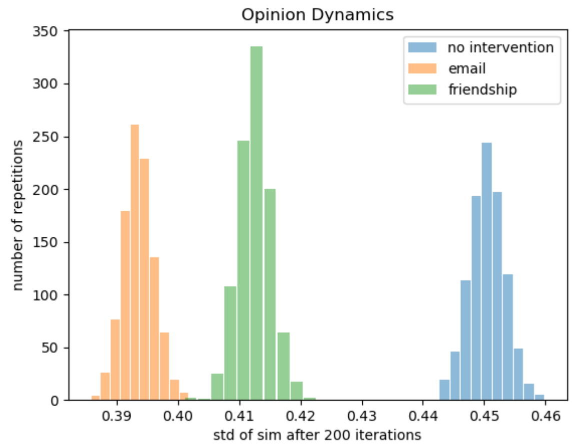

Figure 6: Cumulative Standard Deviations values for 100 repetitions of the Friedkin-Johnsen Simulation under different interventions

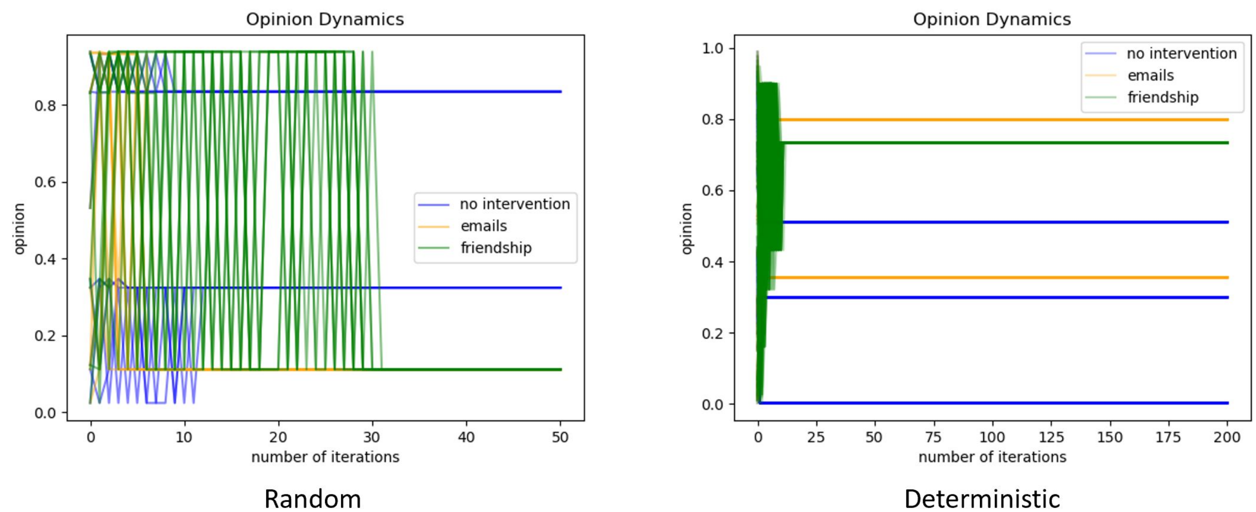

Figure 7: A single repetition of a random vs deterministic Weighted Median Simulation under different interventions

Figure 9 in the appendix. The weighted median is calculated by first sorting each column of the adjacency matrix A by opinion. Then we take the cumulative sum and find the point where the cumulative sum reaches the midpoint at 0.5. Finally, we return the indices where each cumulative sum column is closest to 0.5. In the simulation loop, we simply call this function to find the indices, and reassign the opinions based on each agent's weighted median value. Interestingly, when written this way the code is easier to vectorize, so these simulations ran about an order of magnitude faster.

The results of the simulation are shown for a single repetition in Figure 7 and the standard deviations of many repetitions in Figure 8. These simulations were each ran with 100 agents, for 200 iterations, over 100 repetitions on the two-cluster network shown earlier in Figure 3. These simulations are for a total intervention of 100, which was found to be a reasonable middle value. While both versions share similar properties like the categorical nature of the variables and the oscillation we saw in the random model, the bulk data shows different trends. At this intervention level, the Friedkin-Johnsen model showed the email intervention to perform the best and the random Weighted Median model showed emails and friendship to perform similarly. Now, with the deterministic Weighted Median model the friendship intervention dramatically outperforms the email intervention. We suspect that this is because the random model gave the weak connections of the email biased importance because of the winner-take all nature of a random selection. Over time, this led to higher rates of

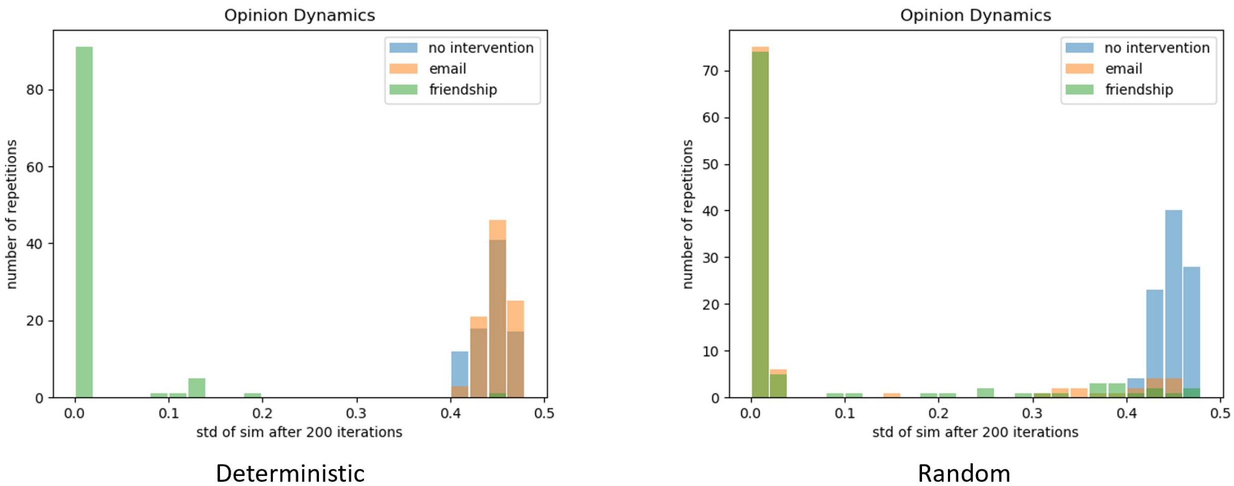

Figure 8: Cumulative Standard Deviations values for 100 repetitions of the Weighted Median Simulation under different interventions

convergence. Now, the deterministic model will almost always pick a higher valued friendship connection at the middle of the opinion spectrum. While we cannot generalize to all levels of intervention off of this experiment, this case shows that friendship converges for lower rates of intervention than email in the deterministic model.

6 Conclusion

This project's main aim was to observe opinion dynamics in various network topologies of stubborn clustered networks, and to draw conclusions about intervention techniques given the data.

For the weighted averaging simulations, we found that just one well connected stubborn agent out of 100 is enough to drastically influence the convergent opinions of the masses. Likewise, if one agent holds very little affinity to its initial opinion, the masses can easily sway them in the opposite direction of an opinion spectrum. We also concluded that when varying percentage of stubborn agents, in order to ensure a diversity of opinions (with a standard deviation of $\sigma = 0.15 \pm 0.05$) at least 20% of the network must be inexorable; assuming we have a uniform distribution of initial opinions. When it came to an agent choosing to diversify their network, the best intervention technique was the many weak ties approach. That is, if an agent finds themselves in an echo chamber, the best technique to de-radicalize themselves is to form many ties to agents with a wide range of opinions, and weigh the new opinions lightly and critically.

Repeating the experiments with a different opinion dynamics model gave very interesting results, as we realized how important the specifics of the underlying model are to the precise conclusions one draws from the data. We found stark difference between the effects of the different models at the specific level we examined, and even that different interventions performed the best using different models. Although it started out as an oversight, the comparison of the random Weighted Median model and the deterministic Weighted Median model proved to be very insightful. Even this slight change completely changed the effectiveness of the email blast intervention, which further illustrates the point.

Although we learned a lot about opinion dynamics and simulation, our biggest takeaway was the importance of the underlying model selected.s

References

[1] Francesco Bullo. *Lectures on Network Systems*, chapter 5: Discrete-time Averaging Systems, pages 65–84. Kindle Direct Publishing, 1 edition, September 2021.

[2] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9), 2021.

[3] Jack Defay and Raquel Taborga. Stubborn networks. `https://github.com/jackdefay/stubborn_networks.git`, 2021.

[4] Elizabeth Dubois and Grant Blank. The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information, Communication & Society*, 21(5):729–745, 2018.

[5] Nabeel Gillani, Ann Yuan, Martin Saveski, Soroush Vosoughi, and Deb Roy. Me, my echo chamber, and i: Introspection on social media polarization. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, page 823–831, Republic and Canton of Geneva, CHE, 2018. International World Wide Web Conferences Steering Committee.

[6] Wenjun Mei, Francesco Bullo, Ge Chen, and Florian Dörfler. Occam's razor in opinion dynamics: The weighted-median influence process. *CoRR*, abs/1909.06474, 2019.

7 Appendix

```
def run_sim(A, x0, sim_length=200):

    # initialize the list of the state of the system
    x = [x0]

    # simulation on a for loop
    for i in range(0,sim_length):
        ids = weighted_median(x[i],A)
        x_temp = x[i][ids]
        # print(x_temp)
        x.append(x_temp)

    return x

# data = x[i], weights = A
def weighted_median(data, weights):
    # take the argsort of the opinions which will be used to order the weights for cumsum
    id_sort = np.argsort(data)
    # rearrange the weights by column
    weights = weights[id_sort,:]
    # find the cumulative sum of the weights
    cs_weights = np.cumsum(weights,axis=0)
    # subtract 0.5 from each value
    cs_weights -= 0.5
    # then take the abs so that the minimum value is the one closest to 0.5
    cs_weights = np.abs(cs_weights)
    # finally, take the argmin, double check the axis
    id = np.argmin(cs_weights,axis=0)
    return id
```

Figure 9: Code for the deterministic Weighted Median simulation