ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed
Computing

A debatable way to gain participation in voting

Bachelor's Thesis

Robin Bosshard

borobin@student.ethz.ch

Distributed Computing Group
Computer Engineering and Networks Laboratory
ETH Zürich

Supervisors:
Dr. Darya Melnyk, Pál András Papp
Prof. Dr. Roger Wattenhofer

September 8, 2020

Acknowledgements

I would first like to thank my supervisors, Dr. Darya Melnyk and Pál András Papp, for their expertise and support during my Bachelor thesis. The discussions from the weekly zoom meetings helped when I was faced with different problems. When I was unsure how to proceed they gave me perspective.
I also want to thank Prof. Dr. Roger Wattenhofer and the Institute for giving me the opportunity to do something almost from scratch.

Abstract

Usually, democratic decision-making processes suffer from low participation ratio. In order to engage more people in the debate voters advertise their own vote to the neighborhood. In turn, previously not participating people are persuaded to cast a vote as soon as they receive enough external votes. Either they agree with the most frequent incoming vote or they might have reasons to counteract.

We examine some scenarios in the voting with two alternatives, and with three alternatives in the one-dimensional political spectrum. We simulate the voting dynamics on multiple different graph topologies, such as the cycle graph, the regular graph and the scale-free Barabási-Albert random network graph, with the initial opinions of nodes chosen at random. We measure the final deviation of the resulting voting ratio from the ground truth. It gives an indication of how advocation of votes can change the outcome of an election.

Contents

CHAPTER 1

Introduction

Some topics may not concern a lot of people at the time of the vote, thus letting only a small fraction of the community decide on the final outcome. But what drives citizens to cast a vote? It is probably not pure self-interest since one single vote seems to have vanishing impact on the overall result. As the effort of voting normally does not payout the resulting benefits, political scientists see also the social aspect of voting and declare it as a main driver of raising participation. The political commitment and voting behaviour in the social environment is essential to get involved in a debate [1]. When a campaign progresses people vote more likely if they see close friends and family have voted [2].

Imagine we allow officially to advertise the own vote to non-participating agents aimed at persuading them to cast a vote as well. Undecided agents receive thereby external votes from their social neighborhood, offering a local picture of the ongoing decision process. Those votes might provide incentives whether it is worthwhile to vote or not or for what, especially when there is a strategical scope. The basic idea originates from the vote process Vote-Vote [3], where you can cast your vote also for someone you choose and this person abstains.
The concrete model incorporates an activation threshold for each agent. That means an agent casts a vote if a certain voting participation ratio in the neighborhood is attained. The model also assumes that purely unbiased agents blindly adopt the most frequent incoming vote while biased agents can also intervene according to their (hidden) preferences. How many agents finally cast a vote is dependent on the initial voters, the presupposed (individual or global) activation threshold and the allowed maximum number of synchronous advertisements rounds.

At the expense of the rather simple mechanism we expect a tightrope walk on presuming higher voting turnout but risking disproportionate voting outcomes. Only a subset of the additional voters cast a vote according to their predisposed opinion. The other agents do not reveal their first preference when casting a vote. Together they produce participation gain but also induce a deviation from the initial state. How much does the final voting outcome deviate from the initial state and the ground truth? Does the opinion supported by the initial majority remain the majority also in the final voting state?

CHAPTER 2

Related Work

An important goal in social choice theory is to understand the dynamics in multi-agent decision making processes. The bare final outcome of a democratic voting hides the previous interactions between agents. For example, researchers try to explain how opinion dynamics unfold in social networks under various diffusion models. In [4] they provide an overview of prominent models, involving the interplay of individual predisposition and the influence of positive and negative peer interaction. Given a binary decision process, a simple opinion formation progress model is for instance the Majority Rule (MR) model [5]. Under (MR) model, agents are assumed to naively adopt the opinion corresponding to the majority opinion held in the neighborhood. So-called swing voters tend to submit to the predominant party or the perceived most influenceable one. This intuitive behaviour is motivated by the bandwagon effect, a phenomenon that definitely appears in voting [6].

Plurality voting is perhaps the most commonly applied option when it comes to aggregating the preferences of multiple voters [7]. For the sake of convenience, we also stick to plurality voting in this thesis. Even when the individual hidden preferences are fixed, the perception on which selectable alternatives or election candidates are likely to be in contention for victory can change during a campaign. This comes along with a dynamic behaviour of strategic vote aligning and thus to potential misreporting of the true preference [7]. In this field researchers study equilibrium dynamics and convergence, see for example [8]. Convergence in this context means that no voter has an incentive to change the vote to make the outcome more favorable from his point of view. Voters who vote not according to their first preference are considered to be *manipulative*. Note that voters might be manipulative not primarily intending to forge the outcome maliciously but rather to avoid a voting outcome which they really do not prefer.

In the research of opinion dynamics the threshold model is used to study *cascades* of opinion adoptions within a network triggered by a so-called initial *seed* set. Threshold models can be used to model collective behaviour. For example, in the model pioneered by Granovetter [9] an agents' behaviour depends on a certain fraction of other individuals in the neighborhood which are already engaging in a certain behaviour. Depending on the individual threshold values and network topology there are different spreads observable: some that are doomed to fail propagating globally already in an early state while others succeed to propagate globally. In this thesis we introduce a threshold value that corresponds to the attainable participation ratio of the neighborhood and thus intuitively also to the perceived importance of the voting. Once the threshold is attained, a non-participating agent is persuaded to cast a vote as well.

Another indirect related topic is liquid democracy [10]. In this system voters have the opportunity to delegate their vote to a chosen set of other participants. Upon voting abstention these participants form together a surrogate vote. However we do not pursue a voting model where we actively choose other participants that relieves us of the responsibility to make our own decisions. Outsourcing the mindset might be a bad idea in the long run. Instead, finding yourself in an unbiased debate of antagonizing forces might be for many non-participating people a reason to recover one's poise. Or in some cases they would like to intervene rather than just passively entrust their voting power to other people.

An active area under investigation in social choice theory is finding an appropriate aggregation of individual preferences into a collective preference outcome. The ranked preferences are collected in preferential ballots and then combined. Famous setups where variants of ranked-choice voting are used are awarding Oscars to people from the film industry or when the best football player is awarded with the FIFA Ballon d'or. However, there is no all-purpose voting method. The Arrow's Impossibility theorem [11] discloses that no voting method is spared from arising paradoxical results. Already with 3 voters and 3 alternatives it is easy to come up with individual preference profiles, yielding in a cyclic collective preference such that there is no unbeaten alternative [12]. This problem was discovered by Condorcet already in 1785. When making a particular domain restriction of preferences, then the median voter theorem gives a consistent answer to finding the most preferred voting outcome. Duncan Black showed that in one-dimensional issue spaces *single-peaked* preferences are a sufficient guarantee for the existence of a median voter [13].

The effect of *vote splitting* can arise in elections where multiple similar candidates stand for election and people vote with a single choice. A similar candidate who eats up votes from another similar candidate is called *spoiler* candidate. In the extreme case, a dissimilar candidate wins the election thanks to vote splitting due to a spoiler candidate who does not withdraw the candidature early on. One of the most mentioned example in the literature is the US presidential election in the year 2000 when Ralph Nader took votes from Al Gore enabling George W. Bush to win [14]. There are voting methods which address the problem of vote splitting, for instance, under runoff voting the spoiler effect is mitigated.

CHAPTER 3

Voting Model

3.1 General Setting

We start in a voting state where some agents already have cast their vote, also called the initial seed set. The ongoing voting is conducted in a given underlying social network $G = (V, E)$ with V, the set of agents entitled to vote with $|V| = n$ and E, the set of (bidirectional) connections between the agents with $|E| = m$. The initial voting state is expressed by a vector v, holding the information whether and what the agents voted for. Some of the voters are decided and either have one top alternative or a ranking.
Non-participating agents can be biased or unbiased. Biased non-voters do have an opinion but, as the name indicates, they do not vote initially. Unbiased non-voters are totally indecisive and therefore also evenly susceptible to any influence from outside without bias. Note that non-participating agents and non-voters are used interchangeably but mean the same.
To conclude, we only consider fixed mindsets which we call the *ground truth*. This ground truth is expressed correspondingly by a vector t and does naturally not contradict vector v.

Non-participating agents (which are split up in biased and purely unbiased) possess a threshold value θ. If the neighborhood's ratio of participation attains θ, the non-voters are persuaded to cast a vote. Which vote to cast is ruled by a conversion policy, depending on the specific model under consideration.
Furthermore, we make a fundamental model restriction on the voting dynamics: Once agents have voted, they must stay with that decision and thus are not allowed to change their vote during the campaign. Therefore it comes to a natural convergence to a final state since the number of non-participating agents decreases in each (synchronous) round until the voting state does not change anymore, i.e. there is no new vote cast.

Based on the initial voting state v, the final voting outcome d is established by multiple synchronous "advertise-to-your-non-participating-neighbors"[1] [3] rounds. Depending

[1] Each voter advertise his vote to all of his non-participating neighbors in each round. This model could be extended such that we allow for weighted advertisements or selecting only a subset of the neighborhood.

on the underlying network and conversion policy, we compare how the final dynamic voting outcome d, the initial voting state v and an assumed known ground truth t deviate from each other with respect to a certain metric. In the following, we use grey nodes to represent all nodes that are not voting currently, i.e. both the purely unbiased nodes and the nodes that have an opinion but are currently non participating.

3.1.1 Implementation

The simulations on the dynamic voting process are executed in a Jupyter Notebook. For the different underlying social networks we use network generators from NetworkX [15] which is a free-software Python library. NetworkX is used for social network analysis and allows to generate and draw synthetic graphs.
An interactive interface allows to simulate different scenarios by setting parameters, such as the number of participating agents and non-participating biased and unbiased agents. Each scenario can then be run multiple times. For each instance the agents are randomly distributed[2] over the generated network graph anew. Then the corresponding recordings and the detected worst case are displayed. Here is a list of some (averaged) recordings that are accumulated during the N runs:

- participation gain in total
- participation gain per color
- participation gain standard deviation per color
- voting ratios of d (in percentage)
- relative frequency which color wins, determined by plurality rule
- required rounds until convergence (when no further vote casts are triggered)
- relative frequency of a switch, i.e. when the initial majority cannot prevail in the final outcome
- deviation from the ground truth, i.e. how much the ratios in the initial and final state deviate from the ground truth ratios
- relative frequency of an improvement, i.e. when d's ratios with respect to t are closer to those of v's with respect to t

[2] In general, we place agents randomly over the network and neglect real world properties like network homophily for example. Network homophily is a principle stating that similar minded persons are more likely connected in a network than dissimilar persons.

In the following, we list some graphs that can be selected as the underlying social network and state why they might be of interest. We only consider bidirectional network graphs for the sake of convenience.

- **Cycle graph**
  Starting with a simple graph structure helps to understand the basic mechanism and to track individual conversions. It also helps towards finding bugs in the implementation.

- **Regular random graph**
  The idea is to distribute the individual advertisement power in a fair way accomplished by letting all nodes have the same degree. However, we assume that the specific topology was determined before the beginning of the voting process. In the simulation we simply generate a random k-regular graph.

- **Watts-Strogatz random graph**
  By analogy with the small-world phenomenon (also known as six-degree of separation), the random graph generated by the Watts-Strogatz model is called small-world network because it incorporates such small-world properties as high clustering and small characteristic path lengths [16].

- **Barabási-Albert random graph**
  The topology and evolution of real networks is governed by robust organizing principles [17]. Networks occurring in the real world are approximately scale-free. That means they possess power laws in the degree distribution and contain so-called hubs which we can observe in real social networks. A well-known model in order to synthetically generate such human made networks is the Barabási-Albert (BA) model. The individual advertisement power is no longer neutrally distributed but more related to the reality. It is motivated by the possibility that we can bind the advertisement power of the nodes to some preexisting social network. The (BA) graph is used for the result chapter.

- **Erdős-Rényi random graph**
  The model of Erdős and Rényi requires two fixed quantities to generate random graphs: the number of nodes and either the number of edges or the probability of edge creation. Then a graph is chosen uniformly at random from the set of all graphs specified by the two quantities. When we simulate the dynamic voting within such a virtual random graph, then the network intrinsically has little to do with the way people interact in the real world with each other. It is just another approach by randomly assigning positions in the network towards a fair proceeding to determine who is allowed to advertise its vote to whom.

3.2 Binary Voting

Assume there are only two alternatives to vote for, e.g. blue and red. This might be an election with two opposing candidates, a voting issue with two different policies of a referendum or simply a binary decision on yes or no.

For each simulation instance we start with $v \in \{-1, 0, 1\}^n$, where $\{-1, 0, 1\}^n$ stands for $\{\text{blue}, \text{grey}, \text{red}\}^n$. Depending on the initial random placement of the seed set we might obtain different final voting outcomes $d \in \{-1, 0, 1\}^n$. We also assume to have the ground truth $t \in \{-1, 0, 1\}^n$ at hand which is obtained by the agents' underlying preferences. For example, a node is labeled in t with -1 (colored in blue) whenever he prefers blue over red, i.e. blue $>$ red. Or labeled with 1 (colored in red) whenever red $>$ blue holds and labeled with 0 (colored in grey) whenever he is purely unbiased.

We consider a conversion policy which includes an intervention rule that takes effect as soon as the threshold θ is attained and the own preference does not conform to the majority of the incoming votes. Whenever biased agents are encouraged to cast their vote then they decide to put their support behind their first preference. See Algorithm 1 for the corresponding conversion policy or read the next paragraph to follow the details.

Algorithm 1 Conversion policy for binary voting:

procedure UPDATE($v, rounds$) ▷ synchronous update round
    $\mathcal{I} \leftarrow getInfluenceableAgents(v)$ ▷ v = temporary voting state
    if $rounds = 0 \vee |\mathcal{I}| = 0$ **then** ▷ converged?
        return v
    for each agent $a \in \mathcal{I}$ **do**
        if activation threshold $\theta(a)$ attained **then**
            $proposed \leftarrow mostFrequentVote(a)$
            if a is biased **then**
                if $proposed = firstPreference(a)$ **then**
                    $color \leftarrow proposed$ ▷ [confirmation]
                else if $proposed = secondPreference(a)$ **then**
                    $color \leftarrow firstPreference(a)$ ▷ [intervention]
            if a is unbiased **then**
                if incoming votes sufficiently polarized **then**
                    $color \leftarrow proposed$ ▷ [agreement]
                else
                    $color \leftarrow 0$ ▷ [abstention]
        else
            $color \leftarrow 0$ ▷ [abstention]
        $v[a] \leftarrow color$ ▷ update v accordingly
    return UPDATE($v, rounds - 1$) ▷ recursion

The algorithm processes the updates in synchronous rounds until convergence. Synchronous means that the effective updates are carried out in the end of a round. In the loop, only influenceable agents with neighborhood's participation ratio greater or equal to the threshold θ are handled. The variable *proposed* stores the most frequent vote of the incoming votes. If the agent is unbiased then *proposed* either conforms with the first preference or not, yielding respectively either a *confirmation* or an *intervention*. If the agent is unbiased then the polarization π is checked. If the incoming votes are sufficiently polarized then it comes to an *agreement*. The polarization value was introduced because uncertain people might not be encouraged to cast a vote if there is no clear external appeal. For instance, $\pi = 2$, means that the most frequent vote must have a margin of at least 2 on the second most frequent vote.
Note that line 7 to 11 could be trivially summarized to: $color \leftarrow firstPreference(a)$ but we explicitly distinguish the two cases between confirmation and intervention for the sake of generalisation and tracking purposes.

3.2.1 Demo of some model aspects

This section is intended for becoming accustomed to some model aspects. Basic examples are designed to demonstrate different aspects, e.g. attaining high enough threshold and polarization but also stating how to measure the deviation of the voting states.

First, we specify two reasons why it might require more than only 1 round to convert a non-participating agent (NPA):

i) NPA is not adjacent to sufficient many voters (i.e. threshold θ not attained) or

ii) NPA is unimpressed due to insufficient polarisation of the incoming votes and thus not yet animated enough to cast a vote in case this is necessary for unbiased agents in the chosen model

A dummy example demonstrating i) is shown in Figure 3.1 with a chosen global threshold $\theta = 0.6$ and polarization value $\pi = 0$. The latter implies that unbiased agents are able to break ties, using random choice. Here the network consist of only 4 agents depicted as nodes connected with edges, reflecting the neighborhood links. We start initially with 2 blue voters and 2 non-voters in v. We learn from the ground truth t that the non-voter node in the middle is biased red. This agent receives 2 external votes and therefor is persuaded to cast a vote in the first round because $\frac{2}{3} \geq \theta = 0.6$. The lower right node is initially not adjacent to sufficient many voters, i.e. not attaining the threshold because $\frac{1}{3} < \theta = 0.6$. However, the previous vote cast induces a sufficient high participation also for the lower right node in the second round.

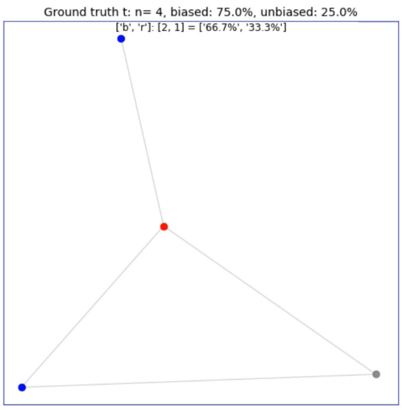

(a) Ground truth t

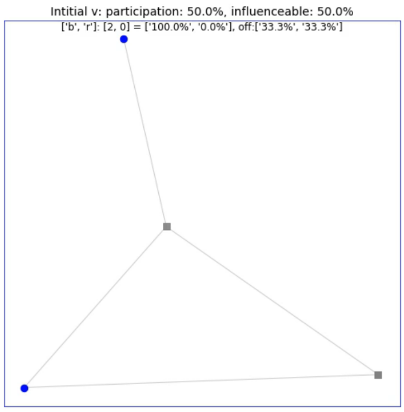

(b) Initial state v

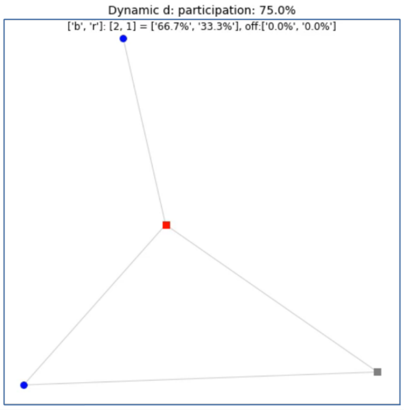

(c) After first round

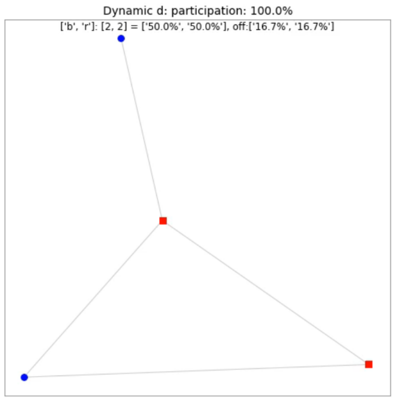

(d) Final state d

Figure 3.1

Assume now $\theta = 0$ and $\pi = 1$. The latter implies that unbiased agents require at least 1 vote advantage of the proposed color for a clear incentive to cast a vote. Otherwise they keep abstaining from voting as they are not able to break ties in this variant. Figure 3.2 demonstrates a case of ii). The agent under consideration is the unbiased non-voter on the crossing position, receiving one blue and one red vote. In the first round he remains undecided since we have a tie and $\pi = 0$. In the meantime, his right neighbor is persuaded to cast a vote for blue. Hence, the agent on the crossing position receives in the second round another blue vote, persuading him to cast a vote for blue as well.

Figure 3.3 shows an example with $\theta = 0$ and $\pi = 1$ where two out of four conversions deploy the intervention rule. Intermediate rounds are omitted in the figure from now on. Furthermore, we observe a *switch* from blue, the winner in the initial state v to red, the winner in the final dynamic state d.

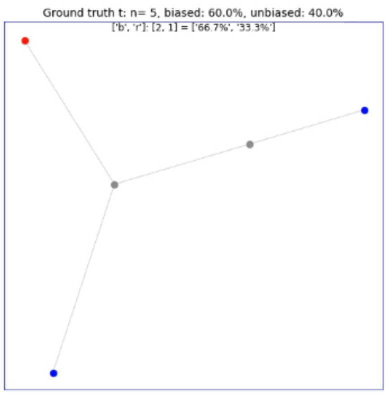

(a) Ground truth t

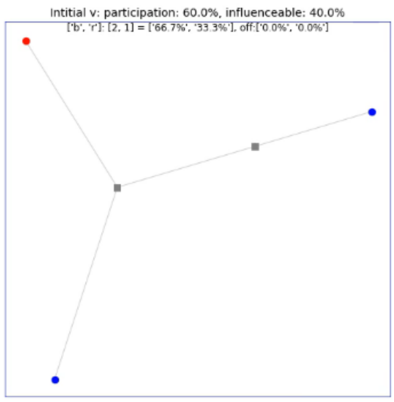

(b) Initial state v

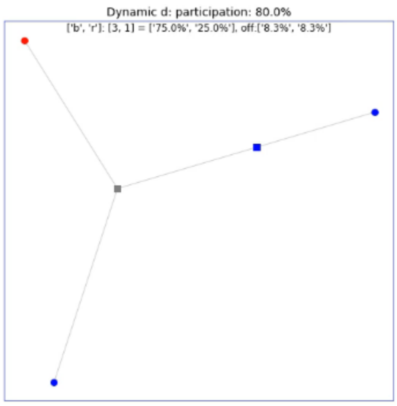

(c) After first round

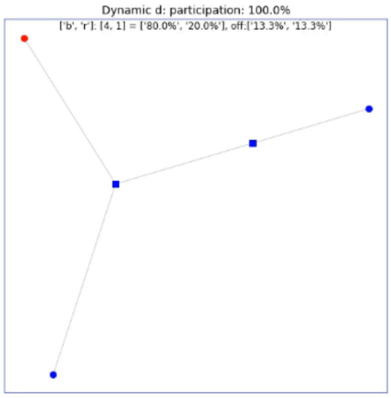

(d) Final state d

Figure 3.2

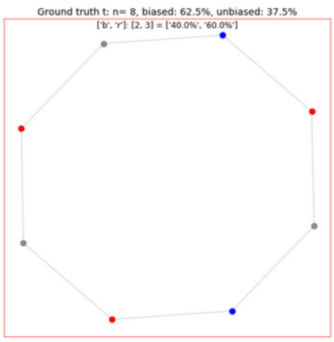

(a) Ground truth t

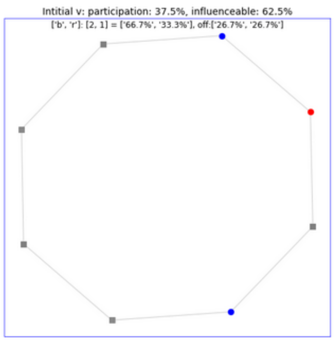

(b) Initial state v

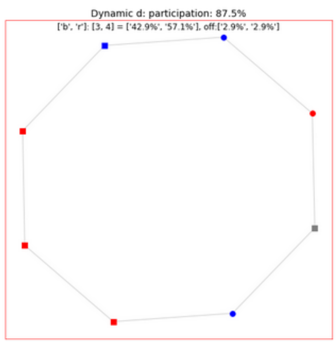

(c) Final state d

Figure 3.3

Note that the boundaries of the figures represent which color is currently in majority.

How do we measure the *deviation* of two voting states? We can read off from t, v and d how many voters or biased agents (in the case for t) per color exist. Accordingly, we define T, V and D capturing the respective ratios per color. The deviation of both V and D with respect to T is measured by taking the sum of the absolute differences for each color ratio. This can be expressed by the 1-norm.

- $v_{dev} = \|T - V\|_1 = \sum_{i=1}^{2} |T_i - V_i|$.
- $d_{dev} = \|T - D\|_1 = \sum_{i=1}^{2} |T_i - D_i|$.

If $d_{dev} < v_{dev}$, we say that the simulated instance has led to an *improvement*. This is the case in the example showed in Figure 3.3 as $d_{dev} = |40.0\% - 42.9\%| + |60.0\% - 57.1\%| < |40.0\% - 66.7\%| + |60.0\% - 33.3\%| = v_{dev}$.

There are of course also initial situations which lead to a *worse*[3] outcome. Such an example is demonstrated in Figure 3.4 with $\theta = 0.5$ and $\pi = 1$. We have again a switch from blue to red.

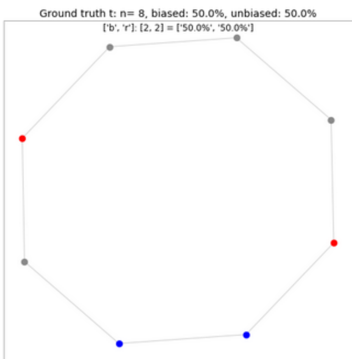

(a) Ground truth t

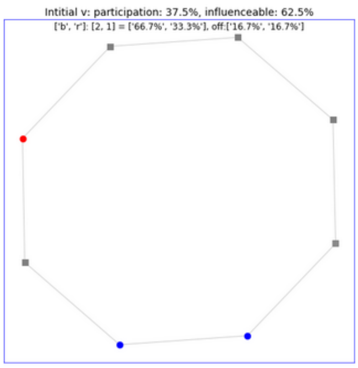

(b) Initial state v

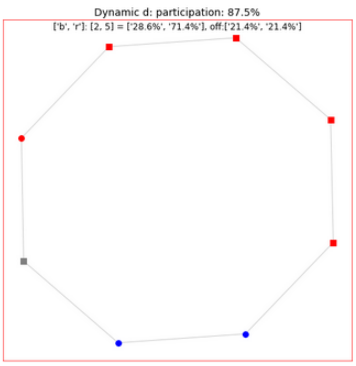

(c) Final state d

Figure 3.4

[3]If $d_{dev} > v_{dev}$

The idea that higher participation in the neighborhood encourages non-voters to participate more likely is a high-level principle and can be implemented in many ways. Here we think of larger and more social network-like graphs. Using a global threshold (the same for each agent) can trigger in some cases a cascade that spreads globally. Sometimes the cascade stops early and drains away, leaving behind only small locally spotted participation gains. Apart from the threshold, the spreading characteristics are also dependent on the network's connectivity parameter and the initial voters. Figure 3.5a shows the numbers of conversions for some fixed threshold $\theta = 0.4$. The underlying network is a Barabási-Albert random graph with some fixed connectivity parameter such that the average degree is nearly 10. The same setting is run over 1000 simulation instances, only varying the positions of the initial voters in the network. In order to obtain a striking picture we have chosen a sparse initial distribution of voters, namely only 75 voters in a network with 500 agents in total. There is a gap visible in the distribution of number of conversions. It came to a global spread only if a critical part of the network was participating.

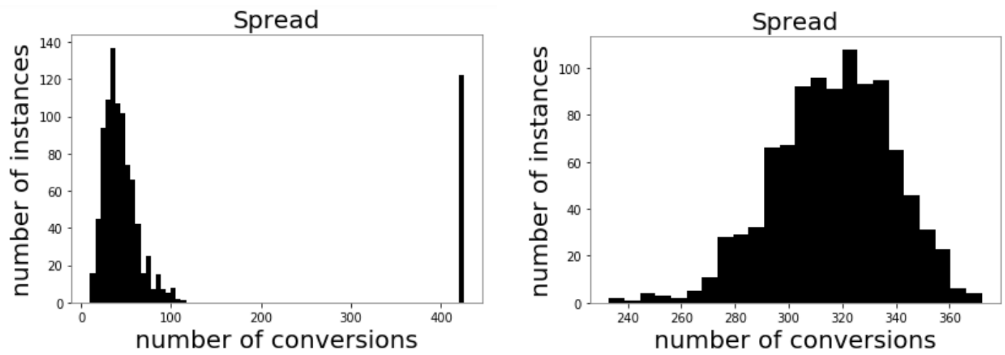

(a) (b)

Figure 3.5: Relative frequency of triggered cascades corresponding to the described setting above and with a global threshold in (a) and a normal distributed (non-global) threshold in (b).

In order to mitigate this effect of extreme spreads we introduce a more fine-grained threshold $\theta' \in \mathbb{R}^n$ which is generated from a normal distribution with mean $\theta \in [0, 1]$ and $\sigma \in [0, 1]$ such that we have different threshold values for each agent. In practice, the same setting from above and with $\sigma = 0.5$ results in expected total participation gains in the range where the gap was before, see Figure 3.5b. Hence, there are automatically agents in the network that decide to participate on their own merits when for some agent a, $\theta_a \leq 0$ and some that will never participate when $\theta_a > 1$. Note that, if we set for instance $\theta = 0.5$ and $\sigma = 0.5$, then we have about $\frac{2}{3}$ of the agents with an individual threshold in $[0, 1]$.

3.2.2 Examples on participation gain

We examine some special scenarios that can occur in terms of the initial situation. First we examine the total participation gain in the cycle graph and the 10-regular random graph depending on a fixed global threshold θ.

The underlying network is a cycle graph with $n = 100$ where we assume that all agents are biased and $p\%$ already cast a vote initially. We assume a global threshold $\theta = 1$. Thus, for a successful animation of a non-participating agent it requires that both neighbors already cast a vote. It is easy to see that only 1 round is required. We can consider n times 3 nodes in a row. It comes to an enclosed conversion if and only if the 2 outer nodes both participate and the node in between is non-participating. What is the expected participation gain? The conversion probability of the node in between is $p_{conversion} = p_{participating}^2 \cdot (1 - p_{participating})$. The expected absolute participation gain is therefore $n \cdot p_{conversion}$. Having for example 50% initial participation gives $p_{conversion} = (\frac{50}{100})^2 \cdot \frac{50}{100} = \frac{1}{2}^3 = \frac{1}{8} = 12.5\%$. The worst found instance is shown in Figure 3.6 with 13% participation gain and a lucky distribution for the blue party.

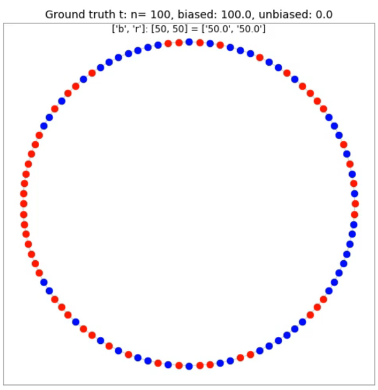

(a) Ground truth t

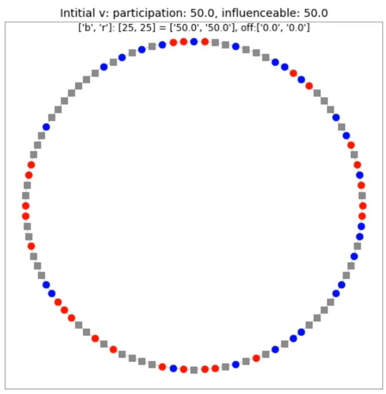

(b) Initial state v

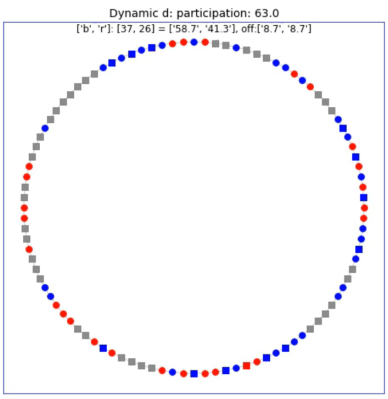

(c) After first round

Figure 3.6: Worst case detected of totally 1000 random instances

The setting is now a 10-regular random graph with $n = 100$. We assume again that all agents are biased and 50% of the agents, chosen at random, already cast a vote initially. The same setting is run for each threshold value 1000 times. First, we observe in Figure 3.7a) that a global threshold of $\theta \leq 0.6$ leads to a global spread while $\theta \geq 0.7$ causes only partial network spreads. The number of rounds peaks where the gains provoke few but enough new gains without converging early. This observation becomes more drastic for the same setting but with $n = 1000$.

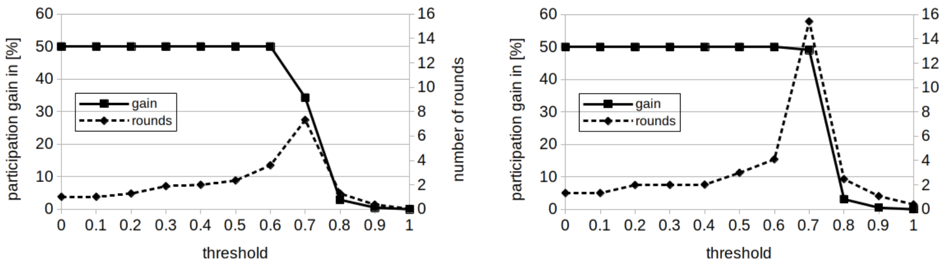

(a) $n = 100$ (b) $n = 1000$

3.3 Multiple Colors

Multiple alternatives to vote for is a mixed blessing as it can also impede the decision-making. To simplify the landscape of the voting model we impose the following restrictions:

- one-dimensional political spectrum
- agents possess single-peaked preferences

We could further simplify the agents' voting behaviour. For instance, we can assume that only the top-two ranked preferences are taken into account as candidates to cast a vote for. In other words: lower ranked preferences are on a blacklist and are never taken into consideration.

3.3.1 Left-Right political spectrum with 3 parties

As a direct enhancement of the binary voting system we add another party in the middle of the axis such that there is now room for compromise (left party = blue, middle party = green, right party = red). The single-peaked preferences enable that the (biased) agents can be classified into four types of agents with following preference profiles:

1. blue $>$ green $>$ red
2. green $>$ blue $>$ red
3. green $>$ red $>$ blue
4. red $>$ green $>$ blue

In the following, we associate a color to biased non-participating agents according to their first preference. The employed conversion policy, see Algorithm 2, is a plausible guideline of voting behaviour in this restricted model. What the conversion policy essentially incorporates when agents have (strategic) incentives to vote for their second preference:

- green agents cast a vote for either red or blue
- red and blue agents only cast a vote for green (they will never cast a vote for the opposing party)

Note that strategic decisions are made, indicated with *compromise*, since agents attempt to prevent an outcome that is their worst choice. For that purpose, the second most frequent vote comes also into consideration. Unbiased agents stick to the most frequent incoming vote when the thresholds are attained.

Algorithm 2 Conversion policy: 3 parties, one-dimensional domain and single-peaked

```
procedure UPDATE(v, rounds)                                  ▷ synchronous update round
    I ← getInfluenceableAgents(v)                            ▷ v = temporary voting state
    if rounds = 0 ∨ |I| = 0 then                              ▷ converged?
        return v
    for each agent a ∈ I do
        if activation threshold θ(a) attained then
            proposed ← mostFrequentVote(a)
            if proposed = firstPreference(a) then
                color ← proposed                             ▷ [confirmation]
            if proposed = secondPreference(a) then
                if firstPreference(a) = secondMostFrequentVote(a) then
                    color ← firstPreference(a)               ▷ [intervention]
                else
                    color ← proposed                         ▷ [compromise]
            if proposed = thirdPreference(a) then
                if firstPreference(a) = secondMostFrequentVote(a) then
                    color ← firstPreference(a)               ▷ [intervention]
                else
                    color ← secondPreference(a)              ▷ [compromise/intervention]
            if a not biased then
                if incoming votes sufficiently polarized then
                    color ← proposed                         ▷ [agreement]
                else
                    color ← 0                                ▷ [abstention]
        else
            color ← 0                                        ▷ [abstention]
        v[a] ← color                                         ▷ update v accordingly
    return UPDATE(v, rounds − 1)                             ▷ recursion
```

Again, we would like to define a model specific metric to measure the deviation between the initial, final and ground truth voting state. Once we have a metric at hand, we can decide whether the final outcome after the simulated dynamics is an improvement to the initial voting state or not. We present a straightforward way (see Algorithm 3) that maps the preference profiles to a reference triplet vector T. In the style of Borda count [18], the idea is here that T encodes the "idealized" ratios for each of the three colors in the system. How do we obtain this triplet vector?

Algorithm 3 Reference triplet vector

procedure
    $T \leftarrow [0, 0, 0]$
    for each biased agent a **do**
        if a's preference profile is of type 1 **then**
            $T_a \leftarrow [2, 1, 0]$
        else if a's preference profile is of type 2 **then**
            $T_a \leftarrow [1, 2, 0]$
        else if a's preference profile is of type 3 **then**
            $T_a \leftarrow [0, 2, 1]$
        else if a's preference profile is of type 4 **then**
            $T_a \leftarrow [0, 1, 2]$
        $T = T + T_a$
    normalize T

From v and d we can read off some triplet vectors V, D capturing the respective ratios per color. The deviation of either V or D with respect to the reference triplet vector T is measured by taking the sum of the absolute differences for each color ratio which can be expressed by the 1-norm:

- $v_{dev} = \|T - V\|_1 = \sum_{i=1}^{3} |T_i - V_i|$.
- $d_{dev} = \|T - D\|_1 = \sum_{i=1}^{3} |T_i - D_i|$.

As T implicitly entails a voting outcome that has the weight placed at the middle party, we expect higher improvements whenever agents have incentives boosting the middle party, be it for casting their first preference or due to making a compromise. When the middle party is not very popular, then we expect a race between the two wing parties, imposing an outcome that might not map well the ratios of T.

CHAPTER 4

Results

We simulate show case scenarios in Barabási-Albert random network graphs and examine the influence of the number of unbiased agents. The stated scenarios are simulated over varying non-global thresholds, fixed standard deviation $\sigma = 0.5$ and polarization value $\pi = 0$ (such that unbiased agents do break ties).

4.1 Binary Voting

We simulate 3 basic scenarios with different starting situations. First, we deal with a scenario where we have accurate voting ratios initially, then we give attention to two scenarios where respectively the illegitimate or the legitimate party is ahead by some margin.

4.1.1 Scenario 1: Accurate voting ratios initially

Let the initial voting state v be proportional to the ground truth t. Concretely, suppose there are 100 blue and red initial voters each and also 100 blue and red biased non-participating agents each. For this special case of a perfectly close head-to-head race, we are mainly interested in the deviation of the final dynamic outcome caused by different numbers of purely unbiased agents, namely $0, 100$ and 500. The corresponding results are shown in Figure 4.1a. As no party has a head start, the obtained deviations serve as a baseline, revealing the inherent deviation induced by the system's randomness. Low thresholds cause many counterbalancing vote casts while higher thresholds reduce the number of vote casts but result in a larger variance in deviation. Accordingly, with 0 unbiased agents we see up to $\theta = 1$ that the higher the threshold the more prone is d to larger deviations. As soon as unbiased agents are in the network they introduce advantageous conditions for an initially lucky distributed color. Thus, the system becomes more likely susceptible to considerably high deviations. With 500 unbiased agents, the peak of the highest deviation is not at $\theta = 0$ or $\theta = 1$ but somewhere in between. The reason for this is that the initially weaker party faces unequal difficulty to convert unbiased agents, e.g. for low thresholds the problem is presumably less substantial.

Suppose now we start with an unbalanced voting ratio of 120 blue and 80 red initial voters. Furthermore, assume there are are also 120 blue and 80 red biased non-participating agents such that the ground truth is again accurately mapped. We expect higher deviations than in the balanced scenario and anticipate that blue will win by a large margin, at least when there are many unbiased agents. Blue is given a clear head start that leads most of the time to an overshooting of the ground truth ratio, i.e. blue converts disproportionately many unbiased agents. The corresponding results are shown in Figure 4.1b.

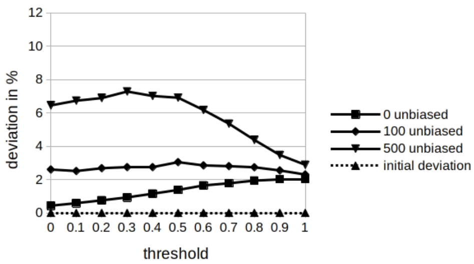

(a) balanced

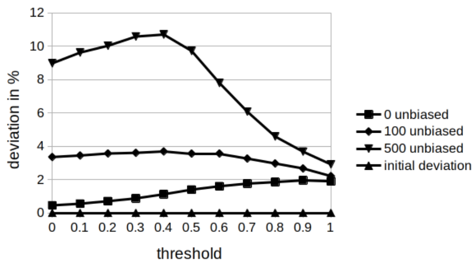

(b) unbalanced

Figure 4.1: The final deviation $\|T - D\|_1$ over (non-global) threshold values with $\|T - V\|_1 = 0$, simulated in a BA random graph with average degree 9.9.

4.1.2 Scenario 2: Illegitimate party ahead

We choose an initial voting state that does not justifiably represent the ground truth:

Initial voters: blue $= 100$, red $= 80$.
Non-participating agents: biased blue $= 70$, biased red $= 100$.

From these numbers, we learn that blue is initially illegitimately ahead. We expect from the more prevalent party red a race to catch up, at least when the number of unbiased agents is not very high. When the number of unbiased agents is high, we expect to drift away from the ground truth ratios. We also expect that this effect increases when the network connectivity becomes higher. Likewise, we show the simulated results (see Figure 4.2) for a Barabási-Albert random network graph with average degree 9.9 and 19.6, each for $0, 100$ and 500 unbiased agents. In addition, we examine the empirical switch probability, see Figure 4.3. Remember that red starts with a minor voting ratio in v and therefore, a switch happens here whenever red succeeds to obtain the majority in the final outcome d.

We conclude from the figures that the voting dynamics introduce a counterbalancing effect on the initially skewed represented ground truth t as long as the number of unbiased agents is not too high. With unbiased agents present, blue exploits its predominance in the network to convert most of them in the early rounds. It leaves red less chance to outweigh blue's rushing conversion boost, especially in the setup with the higher degree and 500 unbiased agents where it comes to an overshooting of the initial deviation. With unbiased agents present, it turns out that the higher network connectivity lowers the likelihood of obtaining a switch. From the perspective of the initially illegitimate weaker party, it tends to become harder to convert unbiased agents from the very beginning on. Except for absent unbiased agents, a switch is slightly more likely. This is an artefact of the slightly higher participation gain obtained in the simulation with degree 19.6. Furthermore, we observe that the empirical switch probability tends to 0 as the threshold becomes 1 as it simply results in too little participation gain for red to catch up seriously.

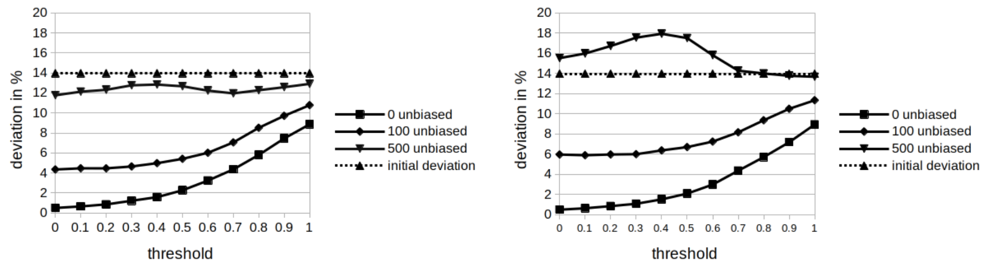

(a) (b)

Figure 4.2: The final deviation $\|T - D\|_1$ for $0, 100$ and 500 unbiased agents compared to initial deviation $\|T - V\|_1$ over (non-global) threshold values. Simulated in BA graph with average degree 9.9 in (a) and 19.6 in (b).

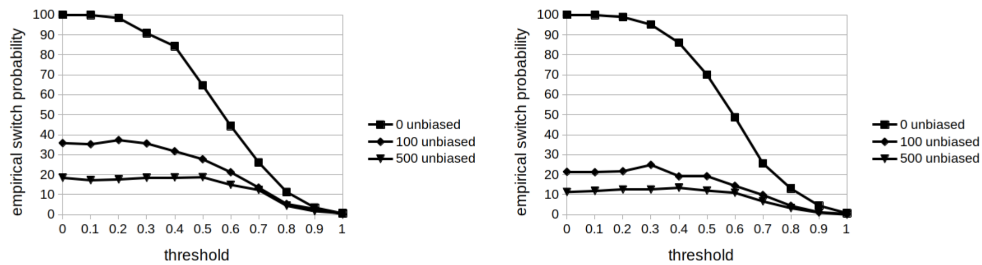

(a) (b)

Figure 4.3: Empirical switch probability corresponding to the simulations of Figure 4.1.

4.1.3 Scenario 3: Legitimate party ahead

In this scenario we are interested in the likelihood of observing an illegitimate switch, i.e. where one party is numerically superior (here it is party blue) but nevertheless loses in the final outcome d. We examine the situation where we have:

Initial voters: blue$= 100 + x$, red$= 100$
Non-participating agents: biased blue $= 100$, biased red $= 100$, unbiased $= 100$, for $x \in 1, ..., 10$.

We show the results only for $\theta = 0.5$ and $\sigma = 0.5$, see Figure 4.4.

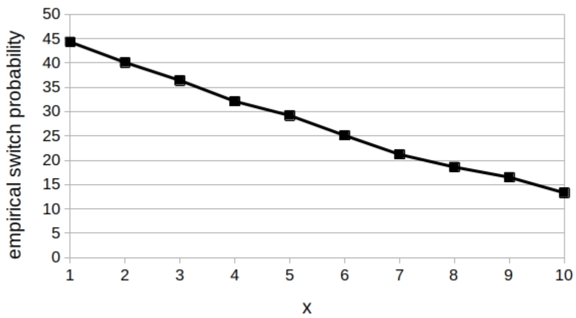

Figure 4.4: Empirical switch probability as a function of x

For instance, if $x = 10$, red wins in 13.3% of the cases and if $x = 20$, red wins in 1.7% of the cases (the latter case is not apparent in the figure).

4.2 Three parties in the one-dimensional political spectrum

We want to model a typical situation where the spoiler effect arises with three parties in the one-dimensional political spectrum, with voters holding single-peaked preferences. This setting was introduced in section 3.3.1. For a proper spoiler effect, we assume that no party takes itself out of the race in favor of another party a priori.
In the following, we do not allow agents of type 2. Remember that we have parties: blue, green and red on an axis from left to right. Absent agents of type 2 (agents with preference profile: green > blue > red) imply that green and red are ideologically closer together (similar alternatives). Moreover, blue is never supported by any biased agent holding a first preference differing from blue. In return, we let the blue party be in advance a tiny bit such that blue holds the plurality in the initial state v. We set the initial voters of green and red such that they outnumber blue when tallied together but each party alone is inferior. The same numbers are assumed for the non-participating agents. We simulate the setting in a Barabási-Albert random network graph with average degree 9.9 and a total number of 600 biased agents where half of them are initially participating.

4.2.1 Scenario 1

Initial voters: blue: 110, green: 90, red: 100
Non-participating biased agents: blue: 110, green: 90, red: 100

We examine the likelihood of winning, i.e to obtain the plurality in d, see Figure 4.5. We expect a race between blue and red.

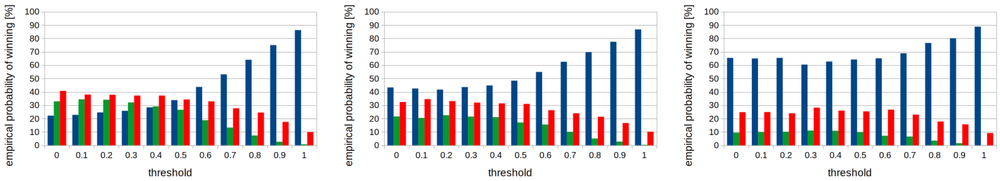

(a) (b) (c)

Figure 4.5: Empirical probability of winning in scenario 1 for (a) 0, (b) 100 and (c) 1000 unbiased agents. Simulated in BA graph with average degree 9.9.

The missing type 2 agents implicate that green and red build together a kind of coalition. Anyhow, blue is not expected to be left behind due to vote splitting. Nevertheless, since green has initially less voters than red, some non-participating biased green agents start to make a compromise in favor for red, accordingly to the conversion policy. We conclude from the figures that blue can defend its head start only when the threshold is high enough or when there is a large number of unbiased agents in the network.

4.2.2 Scenario 2

Initial voters: blue: 110, green: 100, red: 90
Non-participating biased agents: blue: 110, green: 100, red: 90

Now we swap the numbers of green and red agents from scenario 1 but we do not expect a symmetric outcome.

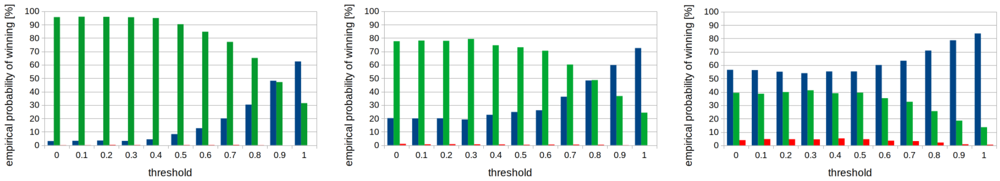

(a) (b) (c)

Figure 4.6: Empirical probability of winning in scenario 2 for (a) 0, (b) 100 and (c) 1000 unbiased agents. Simulated in BA graph with average degree 9.9

When we compare figures 4.6 and 4.5, we learn that the green party succeeds more likely to catch up than the red party at the same relative popularity. This is because green is also taken into consideration by non-participating blue agents while this is not the case for red. Red disappears in the race for victory. In return, the biased red non-voters are expected to put support behind green in an increasing degree when blue spreads to a greater extent and often turns up with the most frequent incoming vote. The race between blue and green becomes tighter as the number of unbiased agents grows, empowering blue in the first place.

The corresponding deviations from the scenarios 1 and 2 are shown in Figure 4.7. Scenario 1 has overall higher deviations imposed by the voting ratios superior represented in the wings.

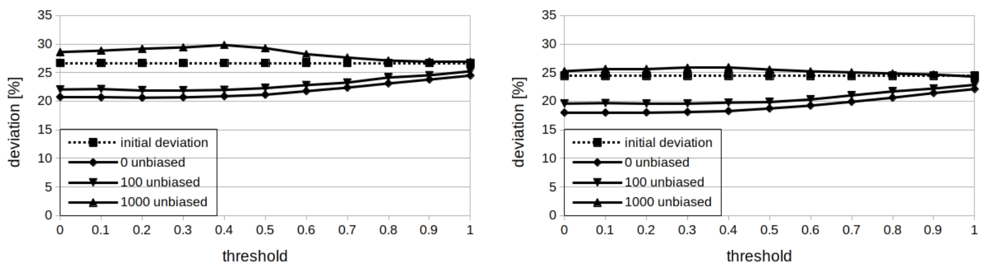

(a) scenario 1 (b) scenario 2

Figure 4.7: Deviations for scenario 1 and 2

CHAPTER 5

Conclusion and Future Work

In chapter 4, we demonstrate the dynamics of our hypothetical voting model by simulating various scenarios with different initial voting states, set up in a Barabási-Albert random network graph. First, in binary voting, we start with initially accurate voting ratios in order to exhibit the scope of the resulting inherent deviation induced by the dynamics' intrinsic randomness. Also, we examine scenarios with skewed initial voting ratios with one color either legitimately or illegitimately ahead. We assert a tendency that the final outcome approximates the ground truth due to biased agents making use of counteracting the most frequent incoming vote. On the other hand, large numbers occurring unbiased agents can lead to disproportional growing of a party. A highly unbalanced initial voting state can amplify this property.
In the second part we examine a somewhat exaggerated scenario with three alternatives predisposed to the spoiler effect. It shows that the actual winner under first preference dynamics loses popularity under the model-specific conversion policy such that one of the two similar alternatives is able to catch up in terms of voting turnout. Thereby the middle party is better off as it is chosen as a compromise from both wings.

In general, the model assumptions made are fairly out of touch with reality. However, it might be worth to adapt and extend the model, also to account for conceptual shortcomings:

1. Let the event of casting a vote be dependent on more features
2. Introduce weighted advertisement by reason that peer influence is not symmetric in general
3. Examine a system with more than 3 parties. In particular, we have to determine a suitable conversion policy that respects the system's political landscape and the agents' possible preferences.

Bibliography

[1] D. W. Nickerson, “Is voting contagious? evidence from two field experiments,” *American Political Science Review*, vol. 102, pp. 49–57, 2008.

[2] Z. Corbyn, “Facebook experiment boosts us voter turnout,” *Nature*, 2012.

[3] [Online]. Available: https://votevote.centreeasy.com/#/home

[4] A. Sîrbu, V. Loreto, V. D. P. Servedio, and F. Tria, “Opinion dynamics: Models, extensions and external effects,” *Participatory Sensing, Opinions and Collective Awareness*, p. 363–401, May 2016. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-25658-0_17

[5] V. Nguyen, G. Xiao, X.-J. Xu, Q. Wu, and C.-Y. Xia, “Dynamics of opinion formation under majority rules on complex social networks,” *Scientific Reports*, vol. 10, p. 456, 12 2020.

[6] R. Nadeau, E. Cloutier, and J.-H. Guay, “New evidence about the existence of a bandwagon effect in the opinion formation process,” *International Political Science Review / Revue internationale de science politique*, vol. 14, no. 2, pp. 203–213, 1993. [Online]. Available: http://www.jstor.org/stable/1601152

[7] R. Meir, K. Gal, and M. Tal, “Strategic voting in the lab: compromise and leader bias behavior,” *Autonomous Agents and Multi-Agent Systems*, vol. 34, no. 1, p. 31, Mar 2020. [Online]. Available: https://doi.org/10.1007/s10458-020-09446-x

[8] N. Gohar, *Manipulative Voting Dynamics*. Cambridge Scholars Publisher, 2017. [Online]. Available: https://books.google.ch/books?id=WTiywgEACAAJ

[9] M. Granovetter, “Threshold models of collective behavior,” *American Journal of Sociology*, vol. 83, no. 6, pp. 1420–1443, 1978. [Online]. Available: https://doi.org/10.1086/226707

[10] B. Ford, “Delegative Democracy,” 2002.

[11] K. J. Arrow, “A difficulty in the concept of social welfare,” *Journal of Political Economy*, vol. 58, no. 4, pp. 328–346, 1950. [Online]. Available: http://www.jstor.org/stable/1828886

[12] P. J. Coughlin, *Single-Peaked Preferences and Median Voter Theorems*. Boston, MA: Springer US, 2004, pp. 849–853. [Online]. Available: https://doi.org/10.1007/978-0-306-47828-4_188

[13] R. Congleton, *The Median Voter Model*, 01 2003, pp. 707–712.

[14] E. Sen, Amartya; Maskin, “A better way to choose presidents,” *nybooks*, 2017.

[15] A. A. Hagberg, D. A. Schult, and P. J. Swart, “Exploring network structure, dynamics, and function using networkx,” in *Proceedings of the 7th Python in Science Conference*, G. Varoquaux, T. Vaught, and J. Millman, Eds., Pasadena, CA USA, 2008, pp. 11 – 15.

[16] D. J. Watts and S. H. Strogatz, “Collective dynamics of ’small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, Jun. 1998. [Online]. Available: http://dx.doi.org/10.1038/30918

[17] R. Albert and A.-L. Barabási, “Statistical mechanics of complex networks,” *Reviews of Modern Physics*, vol. 74, no. 1, p. 47–97, Jan 2002. [Online]. Available: http://dx.doi.org/10.1103/RevModPhys.74.47

[18] P. Emerson, “The original borda count and partial voting,” *Social Choice and Welfare*, vol. 40, no. 2, pp. 353–358, Feb 2013. [Online]. Available: https://doi.org/10.1007/s00355-011-0603-9

APPENDIX A

The appendix shows the reason why the theoretic aspects of this topic were not further pursued. Analysing probabilities of events are either trivial or it becomes quickly intractable or meaningless. We consider the super special case of a blue-red party-system where all non-participating agents are unbiased. In particular, there is no intervention. The global threshold value θ and π are set to 0.

A.1 Binary voting in the complete network

Let $b > r$. What happens?

There is only 1 round. All voters advertise to each non-participating agent. Each undecided agent is challenged by these $p = b + r$ incoming votes and perceives directly the initial ratio of b to r and converts according to the majority rule to blue since we assumed $b > r$.
The party that is already ahead in v, wins also in d and the ratio of b to r of the final outcome grows. The initial state determines in a deterministic way the final voting outcome. It does not come to a switch of the voting outcome but indeed, we have in the end 100% voting participation.

A.1.1 Binary voting in the complete network but with random pick

This is a completely different model from the one studied in the thesis. Non-voters adopt a color from the neighborhood's external votes by random pick. Assume n is odd and $b > r$. Assume the unbiased non-participating agents cast a vote according to a random picked vote of the external votes.

We want to estimate the chance that the red party, who is the minority in v, is able to flip the voting outcome to its favor in the final voting outcome d. Can we make an estimate on that probability $redWins$? Note that $redWins = blueLoses$.
Let $X^{(u)}$ be the sum of u identical and independent Bernoulli random variables X_i, each with expected value $p = \frac{b}{r+b}$. In other words, $X^{(u)} := \sum_{i=1}^{u} X_i$. The exact probability would be $blueWins = \sum_{i=k}^{u} \binom{u}{i} p^i (1-p)^{u-i}$, with $k = \frac{r-b+u}{2} + 1$. This is because blue wins the voting $\iff b + i > \frac{n}{2} = \frac{r+b+u}{2}$, where i is number of undecided agents

that necessarily have to be converted to blue in order that $blueWins = 1$.

Calculating the probability $blueLoses \iff redWins$ with two well-known estimation techniques we obtain:

Using the Central Limit Theorem if u is large, we can approximate $blueLoses = \Pr[X^{(u)} \leq \frac{r-b+u}{2}] = \Pr[\frac{X^{(u)}-u\mu}{\sigma\sqrt{u}} \leq \frac{\frac{r-b+u}{2}-u\mu}{\sigma\sqrt{u}}] = \Pr[\frac{X^{(n)}-u\mu}{\sigma\sqrt{u}} \leq \frac{r-b+u(1-2\mu)}{2\sigma\sqrt{u}}] \approx \Phi(\frac{r-b+u(1-2\mu)}{2\sigma\sqrt{u}})$, where $\mu = \frac{b}{b+r}$, $\sigma = \mu(1-\mu)$

Or we can apply the Chernoff bound to estimate an upper bound on the probability of $blueLoses$ assuming $b - r > \frac{u}{2}$. Define $X := \frac{r-b+u}{2}$ and $\delta := \frac{1-X}{\mathbb{E}[X^{(u)}]}$.

Then, $blueLoses = \Pr[X^{(u)} \leq X] <= e^{-\frac{1}{2}\delta^2\mathbb{E}[X^{(u)}]} = e^{-\frac{(X-\mathbb{E}[X^{(u)}])^2}{2\mathbb{E}[X^{(u)}]}}$. Note that $\mathbb{E}[X^{(u)}] = \frac{b}{b+r}u = \mu u$.

With inputs satisfying: $b - r > \frac{u}{2}$, we observe exponential vanishing probability for $blueLoses$. For really close head-to-head ratio in the initial voting state v, the red party could be lucky and flip the outcome. We see how intricate it is to make estimation in this direction.