COS 521 Final Project

Arjun Mani, Harshvardhan Babla

January 2020

1 Introduction

An individual's opinions are based on their own biases as well as those formed from their interaction with the community. Furthermore, opinions evolve over time through continued interaction. A voter may pick a candidate in an election through their own aligned interests, but may also be swayed by the opinions of their friends. In most scenarios there exists a correct opinion, or a global ground truth. Even if the individuals are well-informed (i.e. they're likely to initially favor the correct opinion) certain dynamics or network structures may lead to the majority holding an incorrect opinion. Particularly, situations where individuals forgo their own opinion in favor of their friends' opinions (known as an *information cascade*) obstruct information aggregation.

In the present day, information cascades are especially present in the proliferation of misinformation and fake-news through online social media platforms such as Facebook, Instagram, and Twitter. Ease of communication and high levels of connectivity accelerate the propagation of distorted opinions. We are especially motivated by the presence of "influencers" on these platforms. These are users who are connected to an unusually large number of other users, paid to advertise products, services, and ideas. This leads to our main motivating question: does the presence of "influencers" in a social network who possess incorrect beliefs lead to the proliferation of these beliefs throughout the network?

We construct a family of graphs such that for any time after zero, the graph holds a majority incorrect opinion despite each node being initially biased towards the correct opinion. This construction is primarily motivated by intuitions about network structures that cause specific nodes to have a higher than usual influence on the majority opinion. This and related work forms the body of our theoretical contribution. We then verify several hypotheses empirically on real life social networks, which provide us insight into the workings of majority dynamics and social networks as a whole. Interestingly, running a simulation of majority dynamics on the graph of prominent Florentine families returns the Medici as the most influential family, accurately reflecting their historic rise to prominence.

Our paper focuses on the model of *synchronous majority dynamics*. Initially, all actors have private beliefs about a binary state of the world, biased towards the correct opinion with probability $\frac{1}{2} + \delta$. In each time step, all individuals synchronously update their vote to match the majority opinion of their friends, breaking ties with their personal opinion. After several time-steps, a winning opinion is declared by performing a majority vote on the population. This is clearly a naïve model, as a truly rational actor would update their opinion through a weighted, or perhaps Bayesian, majority. However majority-dynamics are considered as a more reliable model of relatively uninformed actors (such as voters), whereas Bayesian dynamics is preferred for models of fully rational actors (such as financial traders) [4].

2 Related Work

Our work is related to a line of literature regarding information aggregation and retention in social network graphs. Below we briefly describe most related works, limiting our scope to models that consider independent initial beliefs biased towards the correct opinion with probability $\frac{1}{2} + \delta$.

Banerjee [5] and Bikchandani, Hirshleifer, and Welch [6] first analyze social networks dynamics in a Bayesian setting. This is a model where agents are fully rational and sequentially update their opinions based on the opinions of their neighbors. Their seminal work first identifies information cascades in Bayesian dynamics. Subsequent works provide models that avoid information cascades [7], analyze structures that result in correct consensus [8, 9], etc. While the goals of Bayesian dynamics align with our work on a high-level, they are technically mostly unrelated as we only consider non-Bayesian dynamics.

An alternative line of study is concerned with non-Bayesian, heuristic information aggregation. The study of Voter and DeGroot dynamics considers voters who update their opinion by copying a random neighbor [12, 13] or through a weighted average [10, 11] of their neighbor's reports. The work of Golub and Jackson [14] is somewhat parallel to ours as they find that consensus occurs in a social network occurs if and only if the most influential (i.e. highest degree) node becomes vanishingly less influential as the population size diverges to infinity. Moreover, the work of Berger [15] resembles our work on adversarial graphs, where they show that it is possible for the opinion of a small constant sized majority to dominate the global opinion of the graph. However, unlike our model these works consider a continuous space of opinions and reports.

This paper focuses on the domain of majority dynamics where agents update their vote by declaring the majority opinion of their neighbors. This paper extends the work of Mossel, Neeman, and Tamuz [1] who investigate the graph-structures that result in "efficient aggregation of information". They give a condition on the graph structure which ensures that the correct opinion is chosen with probability approaching 1 as $n \to \infty$, and provide examples of graphs where aggregation fails. Much of our theoretical work is built off their work in this latter area. Tamuz and Tessler [2] study the retention of information and provide sufficient conditions under which the ground truth can be reconstructed from the final state of the dynamics, through any recovery method (not just a majority vote).

Our work is also heavily inspired by the results of Feldman et al. [3] and Bahrani et al. [4] who study similar themes in asynchronous majority dynamics. Feldman et al. [3] claim that for networks that are sparse and expansive, the population will converge to the correct opinion with high probability. Bahrani et al. [4] study examples of structures that stabilize to a correct majority. However, we study synchronous rather than asynchronous dynamics. As pointed out in [4], the difference between the synchronous and asynchronous models is elucidated by the complete graph with n nodes. In synchronous dynamics, everyone announces their private belief in the first step, and updates to the majority in the next step. Hence, a correct consensus occurs with probability $1 - e^{-\Omega(n)}$. However, in asynchronous dynamics the entire population copies the announcement of the first node and a correct majority occurs only with probability $\frac{1}{2} + \delta$.

3 Setting and Exploration

3.1 Majority Dynamics

Our setting is as follows. Let V be a finite set of voters, organized in an undirected social network represented by $G = (V, E)$. Each edge in G represents a connection between two voters. Denote N_v as the set of neighbors of v. We enforce that every vertex has an odd number of neighbors by creating a self-loop for every vertex that would otherwise have an even number of neighbors.

There exist exactly two opinions in this setting. We denote $X_v(t) \in \{-1, 1\}$ as the opinion of node v at time t. In this setting, 1 is considered the correct opinion and -1 is considered incorrect. At time $t = 0$, each voter is initialized with a random opinion $X_v(t) \in \{-1, 1\}$ according to probability distribution P_δ, where $P_\delta(1) = 1/2 + \delta$ and $P_\delta(-1) = 1/2 - \delta$. The initial opinions are chosen independently and identically for all $v \in V$. Observe that once the initial opinions $\{X_v(0)\}_{v \in V}$ are chosen, the process is completely deterministic.

At time $t \in \{0, 1, ..., T\}$, every voter $v \in V$ updates their opinion to the majority opinion of their friends.

$$X_v(t) = \operatorname{sign} \sum_{u \in N_v} X_u(t-1)$$

At some large time T, a majority-vote election takes place with result

$$Y_T = \operatorname{sign} \sum_{v \in V} X_v(T)$$

To avoid ties in the election, we assume that $|V|$ is odd.

Finally, following [1] we define the *wisdom* $\mu_\delta(G,T)$ of majority dynamics on G, up to time T as the probability of recovering the ground truth through a majority vote at T:

$$\mu_\delta(G,T) = Pr_\delta\left[Y_T = 1\right]$$

We also use the term "wisdom" more generally to refer to the probability of recovering the correct opinion, even in a non-majority dynamics setting.

3.2 Degree

We begin with the following simple but useful lemma.

Lemma 3.1. *Let $\mu_\delta(G,0)$ be the result of a majority vote at T=0. If $Pr[X_v(0) = 1] = 1/2 + \delta$, then $\mu_\delta(G,0) \to 1$ as $|V| \to \infty$, for all values of $\delta \in [0, \frac{1}{2}]$.*

Proof. The variables $X_v(0)$ for all v are i.i.d. Bernoulli Random Variables with parameter $p = 1/2+\delta$. Since Y_0 takes the majority of the voters at $t = 0$, $Pr[Y_0 = 1] \geq 1 - e^{-O(n)}$ by a simple Chernoff Bound. Thus, the lemma holds. □

Since each vertex updates using a similar majority vote on its neighbors, we observe from Lemma 3.1 that as the degree of a node becomes large, the probability that it is wise goes to 1. This then drives the intuition that higher average degree per node implies a higher probability of recovering the ground truth. While, this is difficult to analyze theoretically beyond $t = 1$, we verify this intuition experimentally by examining the wisdom of randomly generated d-regular graphs for varying values of d. The details and results are found in Figure 1. Although there is an upward trend, the similarity across d indicates that adding edges to a graph with fixed $|V|$ does not significantly increase the probability of recovery.

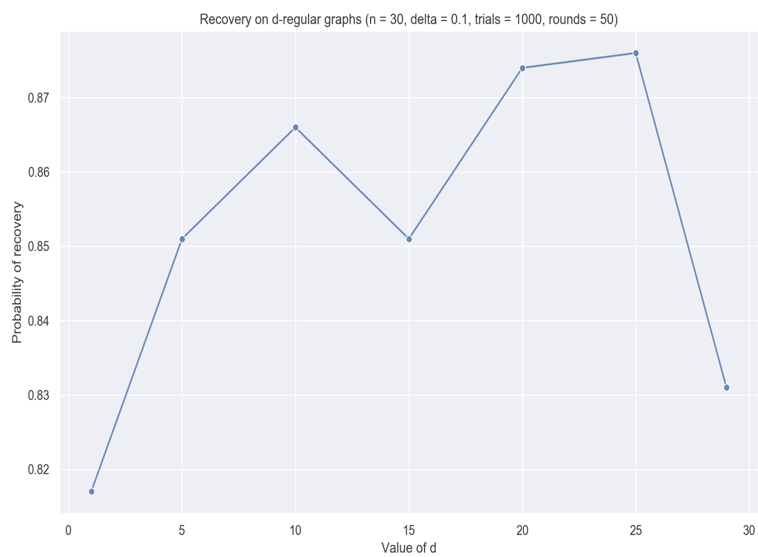

Figure 1: **Probability of recovery (or wisdom) as a function of d.** The above plot shows the probability of recovering the correct opinion from a d-regular graph using a majority vote, for various values of d. For each value of d, we ran 1000 trials. In each trial, a random d-regular graph was generated, and each node assigned an opinion randomly according to the protocol in 3.1. Synchronous majority dynamics were run for 50 iterations, following which the nodes were surveyed. The probability of recovery is simply estimated as the proportion of trials with successful recovery. As we see from the graph, there is a slightly linear trend as d increases, but all d result in a similar probability of convergence. This suggests that increasing the number of edges in a graph for fixed $|V|$ does not significantly increase the probability of recovery.

3.3 High Degree

Another basic hypothesis is that nodes with high degree have a disproportionate "influence" on the probability of recovery. Intuitively, this means that if v is a high-degree node, a higher than average set of nodes probe v's opinion. Thus, the opinions that high-degree nodes are seeded with are likely to significantly affect the outcome of recovery. Our experiment, detailed in Figure 2, confirms this hypothesis. We see that changing the initial opinions of high-degree nodes from correct to incorrect has much more of an effect on the graph's wisdom than changing the initial opinions of low-degree nodes. This idea is explored in more detail in Section 5.3.

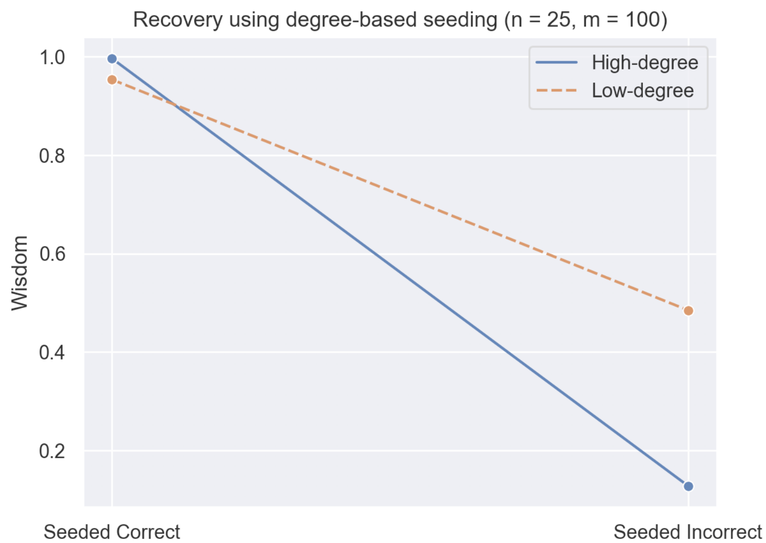

Figure 2: **Degree as a proxy for influence.** In each trial, we take a random $G_{n,m}$ graph, where $n = 25$ and $m = 100$. The expected degree of this graph is $(\sum_v deg(v))/|V| = 8$. We first seed the opinions of high-degree nodes (≥ 10) to 1 and then -1, and measure the drop in wisdom. We then separately the seed the opinions of low-degree nodes (≤ 6) to 1 and then -1, and measure the drop in performance. Although generally each graph contains more low-degree than high-degree nodes, the drop in wisdom is much more significant for high-degree nodes, indicating their influence. Trials=1000, Iterations=50, $\delta = 0.1$.

3.4 Convergence

By Lemma 3.1, a majority vote on any graph at $t = 0$ is likely to recover the correct opinion with high probability. We would thus expect that the wisdom of most graphs would be high for any T (assuming large enough n). However, for certain types of graphs, the nature of synchronous majority dynamics causes the majority opinion of some graphs to oscillate between -1 and 1. Figure 3 shows this experimentally on a set of standard graphs all with 256 nodes.

Examining Figure 3, we see oscillatory behavior for star and wheel graphs and convergence for all other graphs. Theorem 3.2 below is a classical result from McCulloch and Pitts [16].

Theorem 3.2. *For finite graphs, in the synchronous majority dynamics model, each voter's opinion either converges or, after some time, oscillates between -1 and 1 with period two.*

The behavior of star graphs can be explained more generally via Lemma 3.3, which traces through majority dynamics:

Lemma 3.3. *Let G_n be a complete bipartite graph with sets S and T, where $|S| = c$ and $|T| = f(n)$ for some strictly monotonically increasing function in n. As $n \to \infty$, each voter's opinion oscillates with a period of 2 with non-zero probability.*

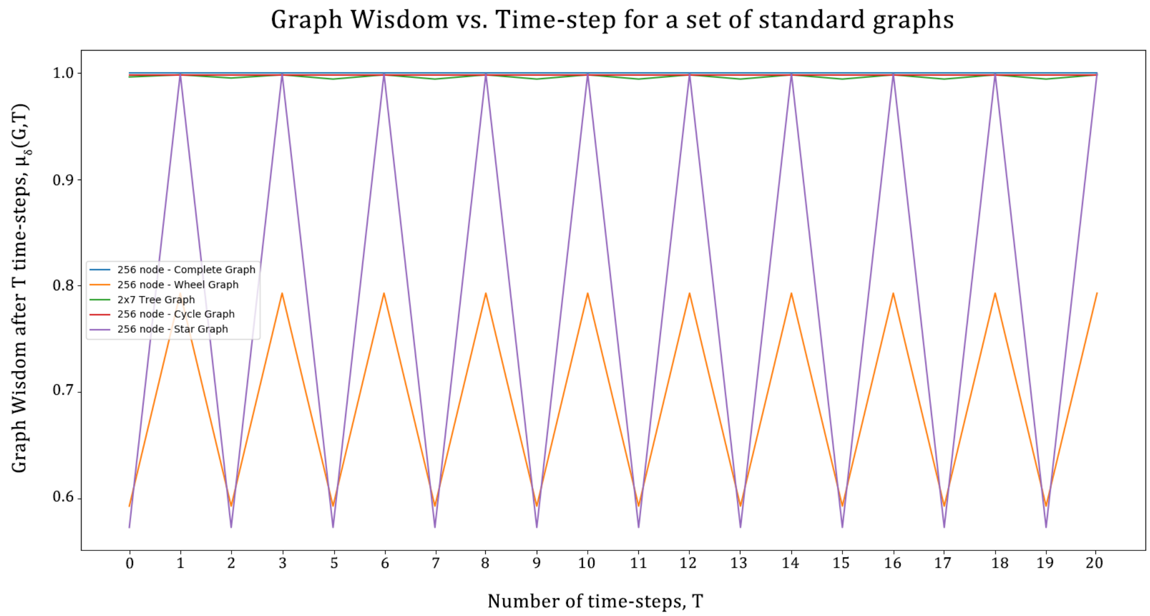

Figure 3: **Wisdom as a function of T.** This figure shows the probability of recovering the correct opinion across T for different types of graphs. The five graph types are the 256-node complete graph (blue), 256-node wheel graph (orange), height-7 binary tree (green), 256-node cycle graph (red), and 256-node star graph (purple). For each graph, we ran 1000 trials. A majority vote was taken after every time-step, for a total of 20 time-steps (initial seed with $\delta = 0.1$. We can see from the graph that the complete, tree, and cycle graphs converge immediately with high probability. However, the wisdom for the wheel and star graphs oscillate with a period of 2.

Proof. Condition on the event that $X_v(0) = -1$ for all $v \in S$. This happens with probability $(1-p)^c$. We observe that since the graph is bipartite, $X_v(1) = -1$ for all $v \in T$. In addition, since the graph is fully-connected, by Lemma 3.1 each vertex in S will be 1 at $t = 1$. Thus, each vertex flips its opinion. At time $t = 2$, since every vertex in S is 1 in the previous iteration and vice versa for T, each voter's opinion will flip again. More generally, each voter's opinion flips every iteration such that $X_v(t) = X_v(t+2)$.

Since $n \to \infty$, $f(n) > c$. Thus, since all vertices in T are incorrect at $t = 1$, the incorrect opinion will be recovered. More generally, we observe that the graph will be majority correct for even t and majority incorrect for odd t. In other words, $Y_T = Y_{T+2}$. Since the probability that $X_v(0) = -1$ for all $v \in S$ is non-zero and independent of n, the majority opinion will oscillate with non-zero probability. □

Note that the star graph is a special case of Lemma 4.2 where $|S| = 1$ and $|T| = n-1$. Thus, majority dynamics oscillates with probability $1 - p$.

4 Adversarial Graphs

An interesting question to ask is whether there exist a sequence of graphs whose size goes to infinity, for whom the probability of recovery is strictly less than 1 for all $T \geq 1$. Note that it takes careful effort to design a graph such that this theorem holds, as the graph's wisdom at $T = 0$ converges to 1 as $n \to \infty$. Theorem 4.1 below extends the work of Mossel et al. [1], as explained below.

Theorem 4.1. *For any $\delta > 0$, there exists a sequence of graphs G_n, where the size converges to infinity, such that*

$$\sup_n \sup_{T \geq 1} P_\delta[Y_T = 1] < 1. \tag{1}$$

Proof. We offer a constructive proof. Let $p = \frac{1}{2} + \delta$ be the probability of any node holding the opinion 1. Let $G_n^{(c)} = (A \cup B, E)$ be a family of graphs defined by some constant c, where $|A| = c/(1-p) - (c-1)$ and $|B| = n(c/(1-p)+1)$. For simplicity, we assume that $1/(1-p)$ is an integer.

Following [1], let each vertex in A be connected to each vertex in B. No nodes in A are connected to each other. The nodes of B are arranged in n cliques, each of size $c/(1-p)+1$. There are no edges between the cliques, and each vertex in B has a self-loop. Intuitively, this graph is set up such that the nodes in A act as "influencers" for the population in B, that is the number of vertices that probe A's opinion is large relative to its size.

For the rest of the proof, let us condition on the event that all the nodes in A are initially wrong, i.e. $X_a(0) = -1$ for all $a \in A$. This happens with probability $(1-p)^{|A|}$.

Consider a clique C in B. Each vertex in C is connected to $c/(1-p)+1$ vertices in C and $c/(1-p)-(c-1)$ vertices in A. Thus, if $\lfloor c/2 \rfloor + 1$ vertices in C vote -1 at $t = 0$, then every voter in the clique will vote -1 at $t = 1$. This is because each $v \in C$ is connected to c more nodes in C than in A. If $\lfloor c/2 \rfloor + 1$ nodes in C have the wrong opinion then a strict majority of v's neighbors will be wrong at $t = 0$.

It is now incumbent upon us to find the probability that $\lfloor c/2 \rfloor + 1$ nodes in C are wrong at $t = 0$. We do this through the following lemma.

Lemma 4.2. *For any clique $C \subseteq V$, let E be the event that $\lfloor c/2 \rfloor + 1$ nodes in C are wrong at $t = 0$. Then $P[E] \geq 1 - 3/e^2$.*

Proof. We observe that the probability of this event not occurring is

$$\sum_{i=0}^{\lfloor c/2 \rfloor} \binom{\frac{c}{1-p}+1}{i} (1-p)^i p^{\frac{c}{1-p}+1-i}. \tag{2}$$

Initially, this seems like a difficult expression to make progress on. However, consider the observation that $p^{1/(1-p)} \to 1/e$ as $p \to 1$, approaching from below. Thus, we reason that the above probability grows with p. We therefore examine the limiting behavior as $p \to 1$. As we shall see, this provides us with a rigorous argument for the function growing with p, since the middle term, the only term that does not grow with p, cancels.

Consider each term in the sum. As $p \to 1$, the right sub-term goes to $1/e^c$. Meanwhile, the first two sub-terms approaches the following:

$$\begin{aligned}\binom{c/(1-p)+1}{i}(1-p)^i &= \frac{(1-p)^i}{i!}\prod_{j=0}^{i-1}\frac{c}{1-p}+1-j \\ &\to \frac{1}{i!}c^i.\end{aligned}$$

Note, as mentioned above, that the term $(1-p)^i$ disappears. Thus, as $p \to 1$, the sum converges to

$$\sum_{i=0}^{\lfloor c/2 \rfloor} \frac{c^i}{i! e^c}. \tag{3}$$

Since the denominator dominates the numerator, this goes to 0 very fast as a function of c (this is verified analytically). Thus, the maximum value of this expression is for $c = 2$, where it equals $3/e^2$. Since examining the limiting behavior as $p \to 1$ gives an upper bound on the probability of $\neg E$, $P[\neg E] \leq 3/e^2$ and therefore $P[E] \geq 1 - 3/e^2$ for any c. $\square$

With Lemma 4.3, we can now complete the proof. The probability that any one clique will vote all incorrect at $t = 1$ is at least $1 - 3/e^2 \approx 0.594$. Thus, the number of cliques that will incorrectly vote at t=1 dominates the distribution $\mathrm{Binom}(n, 0.59)$. By Hoeffding's inequality, the probability that a majority of the cliques will vote -1 at $t = 1$ is at least $1 - \exp(-0.01n)$. Once this happens, those cliques will all vote -1 for

$t \geq 2$. In addition, every node in A will also vote -1 for $t \geq 2$. Therefore, for all $t \geq 2$, a majority vote will result in the incorrect opinion.

Finally, observe that the event that all of A initially votes -1 is independent from the event that a majority of cliques have enough voters that vote -1. Thus, the probability that both events occur is $(1-p)^{|A|}(1-\exp(0.01))$, which is positive and independent of n. Thus, information does not aggregate with a non-zero probability independent of the size of the graph, and the theorem is proved for the family of graphs $G_n^{(c)}$ parametrized by c.

□

The authors in [1] prove this theorem for the same underlying topology of graphs, with $c = 1$. Part of the reasoning is similar, but since we prove the theorem for all c, our proof significantly diverges. One might ask, other than generality, why proving for all c is useful. Our answer is that in [1], the authors restrict the size of A, the "influencers", to be approximately the size of each clique. However, a more realistic model of social networks is where you have a smaller group of influencers targeting a number of larger communities. By extending the proof to all c, we show that if the number of influencers is any constant number less than the size of each community, the graph will still reach the incorrect opinion with non-zero probability. Our class of graphs more closely approximates the topology of a real-life social network.

The graph described in the proof for Theorem 4.1 highlights an interesting method to drive the majority opinion of a graph away from the ground truth by seeding the opinion of a small group of "influencers" to be false. The below corollary begins to pave the way for designing a larger set of adversarial graphs. First, we prove the following lemma:

Lemma 4.3. *For any clique $C \subseteq G_n^{(c)}$, let E be the event that $k \leq c$ nodes in C are wrong at $t = 0$. For $c \leq 100$, $P[E] > 0.5$.*

Proof. Proceed similarly to the proof in Lemma 4.2. Now, the probability of failure is

$$\sum_{i=0}^{c-1} \frac{c^i}{i!e^c}. \tag{4}$$

This sum is difficult to analyze, but experimentally it takes on values < 0.5 for large c. For example, when $c = 100$, the value of the sum is 0.487. Evaluation beyond this point is numerically difficult.

□

In the following proof, we assume the above lemma holds for all c. We are guaranteed that it holds for large c (≤ 100).

Corollary 4.3.1. *Modify the structure of the graph $G_n^{(c)}$ from Theorem 4.1, by attaching a subgraph H_c to each clique $C \in B$ such that $|D_C| \leq c-1$, where D_c is the set of vertices in H_c with at least one neighbor in C. If each node in D_C is connected to more vertices in C than not, then all vertices in $A \cup C \cup D_C$ hold the incorrect opinion for $t \geq 2$ with non-zero probability independent of $|V|$.*

Proof. With the condition that every vertex $v \in D_C$ is connected to more nodes in C than, elsewhere in the graph, if all vertices in $N_v \cap C$ hold the same opinion at time t, then all vertices in D will update to hold the corresponding opinion at the next time step. Therefore, WLOG we can say that D only has neighbors within C.

Following the proof for Theorem 4.1, we condition on the event that all vertices in A and k nodes in C hold the wrong opinion (i.e. -1). Let us consider the worst case where the vertices in C and D_C have maximal mutual influence while those in A and D have minimal mutual influence - i.e. each vertex in D_C is connected to every vertex in C and none in A. This is the worst case since each node in D_C could be seeded to 0, and we condition on the event that all vertices in A and k nodes in C hold the wrong opinion (i.e. -1).

At $t = 0$, the members of C have at least $|A| + k = |C| - c + k$ neighbors that vote -1 and at most $|C| - k + |D_C|$ neighbors that vote 1. Therefore, to ensure that all members of C vote -1 at the next time step ($t = 1$), we must have that $2k > |D_C| + c$. At $t = 2$, the members of C vote -1 if $|C| > |A| + |D_C| = |C| - c + d \implies |D| < c$. Together, both inequalities give that $k > |D_C|$.

Finally, we use Lemma 4.3 that says that for $k \le c$, $P[E] > 1/2$ (E as defined in the lemma). Thus, by Hoeffding's Inequality the majority of cliques will be all -1 at time t = 1 with non-zero probability independent of $|V|$. Thus, all of A will be -1 at time $t = 2$. We have previously argued that all of C and D_C will be -1 at time $t = 2$. Thus, the lemma holds. □

Corollary 4.3.1 proposes a compelling method to "attack" the majority opinion of the graph, by seeding the initial opinion of only a small constant sized set of nodes. Section 4.1 walks through an example of such an attack.

4.1 An Example: Trees

Consider a hierarchical social network modeled as a tree with height h; organized such that the leaf nodes belonging to each parent form a clique of size $2/(1-p)+1$. Given a constant sized set of nodes A with $|A| = 2/(1-p) - 1$ seeded with an initial opinion of -1 (i.e. the opposite of the ground-truth), Corollary 4.3.1 implies that any clique of leaf nodes C and their corresponding parent D_C are likely to entirely hold the opinion -1 after two time-steps, if A is fully connected to all C. Hence, by connecting the adversarial set A to all the leaf nodes in the social network, we can say w.h.p. that the two bottom-most layers of the tree hold the false opinion after two time-steps.

Once this happens, the false opinion rapidly propagates up the tree with high probability, as each parent node with more one child adopts the opinion of its children if they all agree. Therefore, after only $h+1$ time-steps, the root node would be infected with the false opinion w.h.p. This emphasizes a way to not only shift the bias of a majority vote away from the ground-truth within only a few time-steps, but also infect the "most-powerful" node in the hierarchy (analogous to the root node of a tree).

5 Real Social Networks

In the following sections, we run several experiments on well-known real-life social networks. Each experiment is centered around a different idea and provides unique insight and intuition. We use the following four graphs in our experiments:

- **Karate Club Graph**: A social network representing a university karate club. 34 nodes, edges document links between pairs of members who interacted outside the club.
- **Florentine Families Graph**: 15 nodes, each represents a prominent Florentine family. Edges indicate marriage links.
- **Davis Southern Women Graph**: 18 nodes, each represents a Southern woman. Edges indicate co-occurrence at social events. The graph is bipartite.
- **Connected Caveman Graph**: Not a real-life social network, but a realistic model of a real-life social network. Defined by n, k, graph contains n cliques of size k. For each clique, an edge is selected at random and reconnected to another clique.

The topologies of each of these graphs are included in the appendix for reference.

5.1 Thresholding

After T iterations, majority dynamics attempts to recover the correct opinion of the graph by taking the opinion held by the majority of voters. Mossel et al. [1] proposed the idea of establishing higher or lower thresholds for agreement. Here, the recovered opinion is 1 if for some threshold $\alpha \in [0,1]$, $\sum_{v \in G} I_v \ge \alpha|V|$, where I_v is an indicator variable indicating whether the opinion of the v-th node is 1. In other words, at least an α proportion of the nodes must be correct to recover the correct opinion. Observe that majority dynamics is simply a special case with $\alpha = 0.5$.

Most graphs will have reasonably high wisdom under majority dynamics, because they are already reasonably wise at $t = 0$. Studying the probability of recovery here is interesting because it reveals the extent to which a graph achieves consensus. If it has achieved complete consensus on the correct opinion, then no matter the value of α it will recover correctly. Thus, the sensitivity of wisdom to α is a good measure of consensus.

Figure 4 shows how recovery probability varies as a function of α for four social network graphs. Recovery probability predictably drops, but for the three social networks it is still ≈ 0.5 even for $\alpha = 0.99$, indicating that consensus is reached half the time. However, recovery probability drops drastically for the connected caveman graph. One likely theory is that the less easily a graph can be divided into communities (clusters of nodes with high internal degree and low external degree), the more likely that it will achieve consensus. Small communities could be all or largely wrong, in which case the rest of the graph cannot correct them. This hypothesis is bolstered by the fact that a complete graph of size 100 has a wisdom of 0.974, while a 10,10-connected caveman graph has a wisdom of 0.007. Examining patterns of disagreement in the connected caveman graph reveals that disagreement occurs at a clique-level (entire cliques agree or disagree), which is what we would expect.

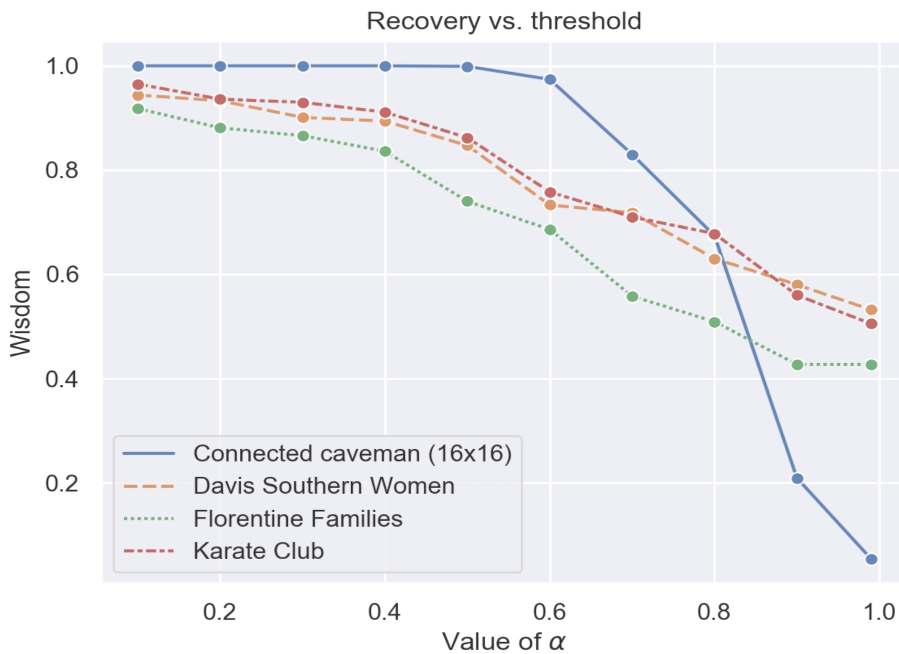

Figure 4: **Probability of recovery as a function of α.** The above plot varies the threshold α and displays the wisdom of all four graphs for each value of α. Interestingly, while the three real-life social networks perform reasonably for high α, the performance of the connected caveman graph plummets. This may be well-understood by the fact that the connected caveman graph contains a high degree of community structure. Nodes in a community tend to converge to the same opinion independent of the rest of the graph. As always, 1000 trials, 50 iterations, and $\delta = 0.1$.

5.2 Seeding

We return to the ideas of Section 3.3, that the opinions of high-degree nodes have more influence than those of low-degree nodes. We experiment with seeding nodes according to different probability distributions than P_δ. In particular, we compare P_δ to two PDs that favor placing the correct opinion on high-degree nodes and low-degree nodes respectively. These are defined as the following:

$$P_{hd}(v_{op} = 1) = \sigma(\deg(v) - m(G)), \tag{5}$$

$$P_{ld}(v_{op} = 1) = \sigma(m(G) - \deg(v)), \tag{6}$$

where σ is the sigmoid function and $m(G)$ is a function that computes the median degree of a graph. Figure 5 shows the results. The temperature of the sigmoid was adjusted to create roughly equal numbers of initially correct nodes for each graph - however, the number of nodes initially correct is still a confounding factor (P_δ tends to set more initially correct). Despite this, we see that low-degree seeding does much worse than random or high-degree seeding, and high-degree setting generally does better than random seeding. This indicates both the importance of the initial seeding, as well as the usefulness of high-degree as a proxy for influence.

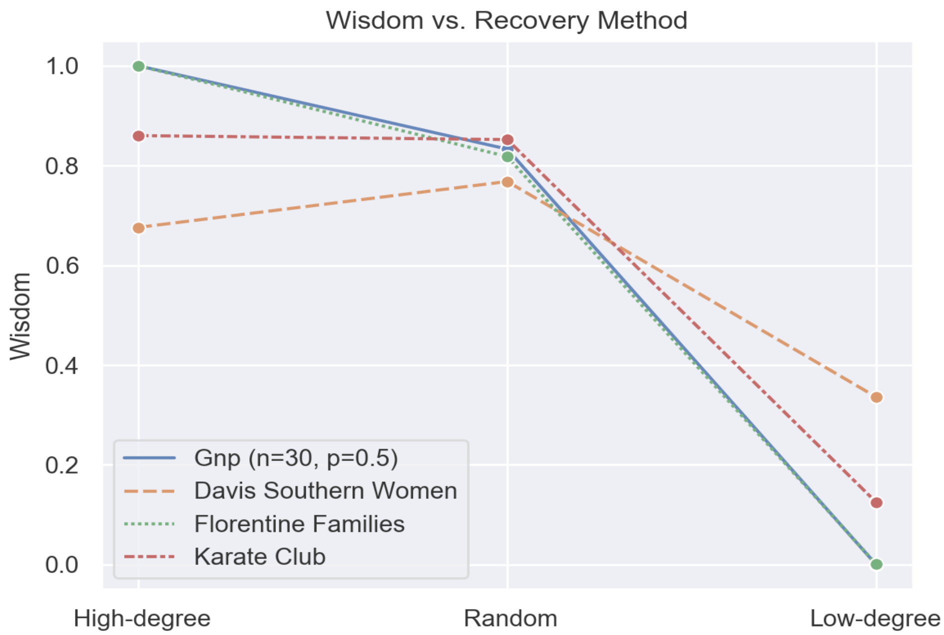

Figure 5: **Probability of recovery for different seeding strategies.** The above plot shows how the wisdom of each graph changes as the seeding strategy changes from favoring high-degree nodes to P_δ to favoring low-degree nodes. Both the former polling strategies do significantly better than favoring low-degree nodes, and generally high-degree node seeding does slightly better than random seeding (the noise is due to the random method usually seeding more nodes to 1). Trials=1000, Iterations=50, $\delta = 0.1$.

5.3 Influence

The process of seeding in the paper has been guided by "influential" nodes. These are nodes whose opinions have a significant effect on the opinions on the rest of the network. Our proxy for influence thus far has been degree. However, we can measure the influence of a node through the following method: seed the node's opinion to -1, and then to 1, and observe how the wisdom changes. A high-influence node should see a significant difference, with the caveat that this approach works most effectively for smaller graphs where the effect of a single node can be significant.

To make this approach more concrete, for each trial we consider every vertex v. We set $X_v(0) = 1$ and attempt recovery, and then set $X_v(0) = -1$ and do the same, with all other nodes seeded as usual. The wisdom is then computed across all the trials. Finally, the influence of a node is the percent change in the wisdom, or

$$I_v = \frac{P[Y_T = 1|X_v(0) = 1] - P[Y_T = 1|X_v(0) = -1]}{P[Y_T = 1|X_v(0) = -1]}. \tag{7}$$

5.3.1 Florentine Families Graph

The Florentine Families Graph is a particularly interesting case for this approach given its history. The Medici family, although not the most wealthy or possessing the most seats in the legislature, eventually rose to power. This is perhaps well-understood through this social network, which reveals that the Medici Family were central in the graph and had the highest number of connections among any of the families. The question is whether the process of majority dynamics can successfully identify the Medici Family as the most influential node in the graph.

Examining Figure 6 (which contains more details) reveals that the method was successful. Nodes are colored by influence from dark to bright, and the Medici family stands out as significantly more influential than any other family. In the graph, the Medici node not only has a high degree, but its neighbors have comparatively low degree relative to median degree in the graph. This implies that the Medici node's opinion will have a high level of influence on the opinions of several nodes. This idea suggests that influence in majority dynamics is more complex than high degree. One possibility is that a high-influence node v must have a low-degree path to a high number of vertices w ($\prod_{u \in P(v,w)} deg(u) \leq k$). Similar ideas have been formalized in [4] for asynchronous dynamics.

The success of majority dynamics here suggests that "influence" in majority dynamics might indicate the importance of a node in a social network. It is intuitive that nodes which influence the opinions of other nodes might be considered more important. Although there are other measures of importance (e.g. betweenness centrality, which also does an excellent job on this graph), majority dynamics might prove useful for such an analysis. This could be a future direction of investigation.

This same process was also run on the Karate Club graph. Since the analysis and conclusions are similar, we choose not to include the graph.

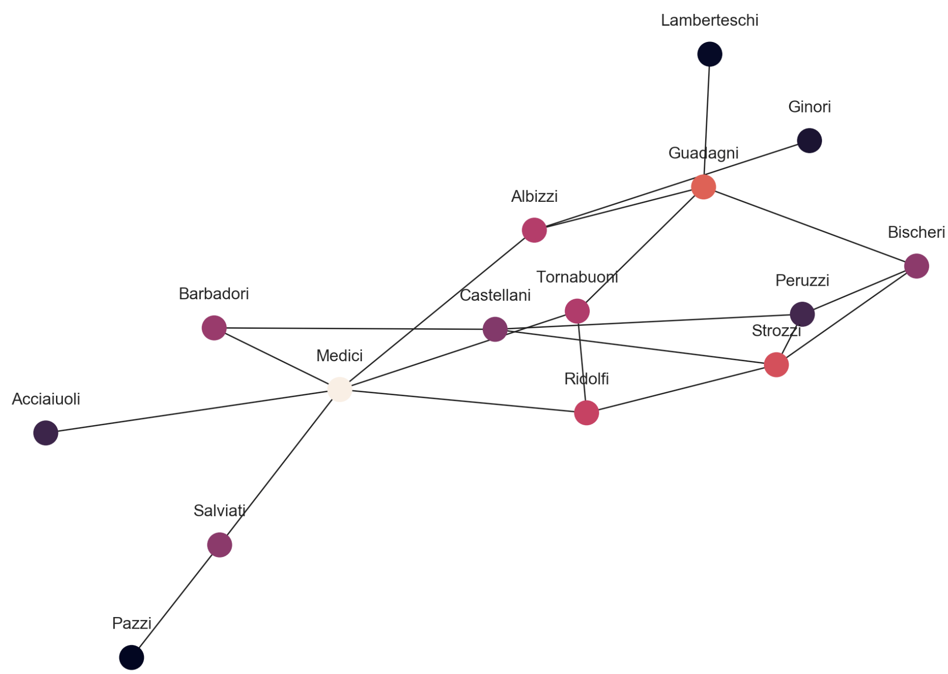

Figure 6: **Graph of Florentine families.** Each node is colored by its influence (dark to light). The Medici family has a much higher influence any other family, which might explain its rise to prominence.

References

[1] Elchanan Mossel, Joe Neeman, and Omer Tamuz, Majority dynamics and aggregation of information in social networks, Autonomous Agents and Multi-Agent Systems 28 (2014), no. 3, 408–429.

[2] Omer Tamuz and Ran Tessler. Majority dynamics and the retention of information. 2013.

[3] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. Reaching consensus via non-bayesian asynchronous learning in social networks. In Approximation, Random-ization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain, 2014.

[4] Maryam Bahrani, Nicole Immorlica, Divyarthi Mohan, S. Matthew Weinberg. Asynchronous Majority Dynamics in Preferential Attachment Trees. arXiv:1907.05823

[5] Abhijit V. Banerjee. A simple model of herd behavior. The Quarterly Journal of Economics, 107(3):797–817, 1992.

[6] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change in informational cascades. Journal of Political Economy, 100(5):992–1026, October 1992.

[7] Lones Smith and Peter Sorensen. Pathological outcomes of observational learning. Econometrica, 68(2):371–398, March 2000.

[8] Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. Games and Economic Behavior, 46(1):1–22, January 2004.

[9] Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. Review of Economic Studies, 78(4):1201–1236, 2011.

[10] Morris H. DeGroot. Reaching a consensus. Review of Economic Studies, 69(345):118–121,1974.

[11] Mohsen Ghaffari and Merav Parter. A polylogarithmic gossip algorithm for plurality consensus. In ACM Symposium on Principles of Distributed Computing (PODC), 2016.

[12] Peter Clifford and Aidan Sudbury. A model for spatial conflict. Biometrika, 60(3):581–588,1973.

[13] Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. The Annals of Probability, 3(4):643–663, 1975.

[14] Benjamin Golub and Matthew O. Jackson. Naıve learning in social networks and the wisdom of crowds. American Economic Journal: Microeconomics, 2(1):112–149, 2010.

[15] Eli Berger. Dynamic monopolies of constant size. J. Comb. Theory, Ser. B, 83(2):191–200,2001.

[16] Warren McCulloch and Walter Pitts, A logical calculus of the ideas immanent in nervous activity, The Bulletin of Mathematical Biophysics 5 (1943), no. 4, 115–133.

Appendix: Diagrams of the Real Life Social Network Graphs

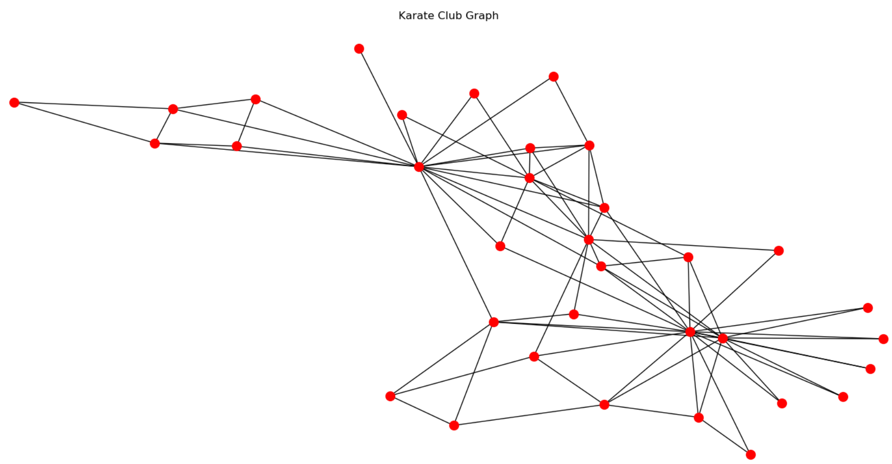

Figure 7: **Karate Club Graph** This diagram represents a social network of members at a university karate club. The graph has 34 nodes, and the edges between nodes correspond to members who interacted outside the club.

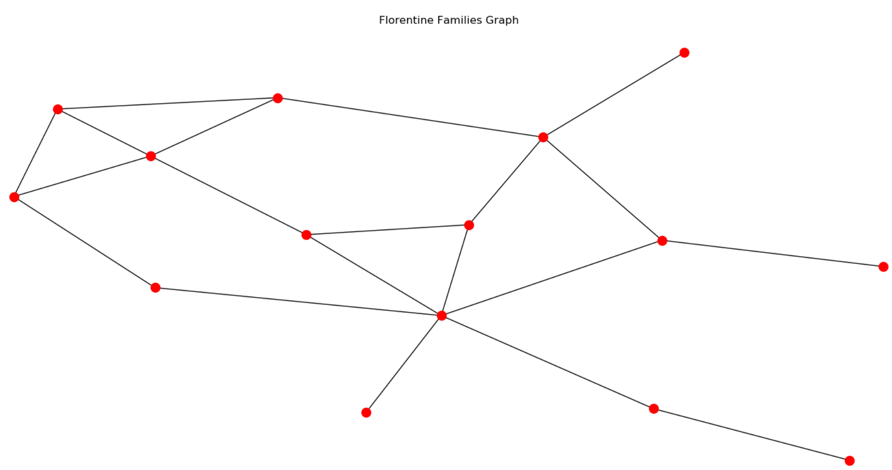

Figure 8: **Florentine Families Graph** This diagram shows the full Florentine Families Graph. The graph has 15 nodes, each corresponding to a prominent Florentine family. The edges indicate marriage links.

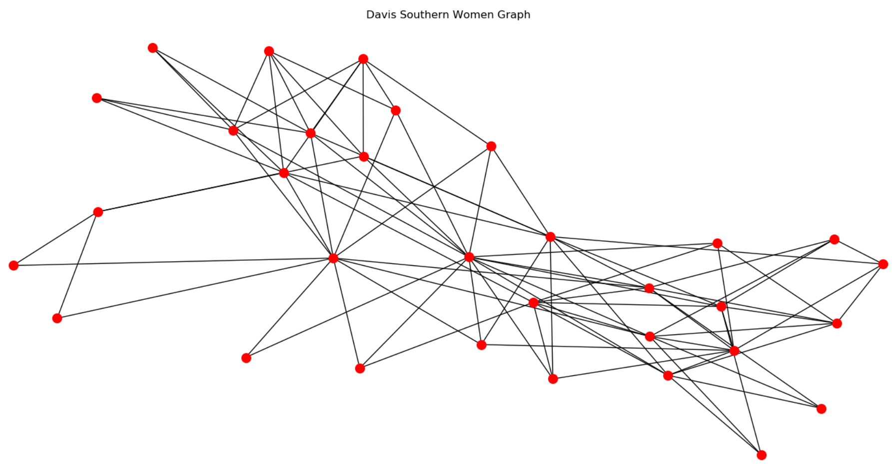

Figure 9: **Davis Southern Women Graph** This diagram shows the full Davis Southern Women Graph. The graph has 18 nodes, each of which corresponds to a Southern woman. The edges indicate co-occurence at social events. This is a bipartite graph.

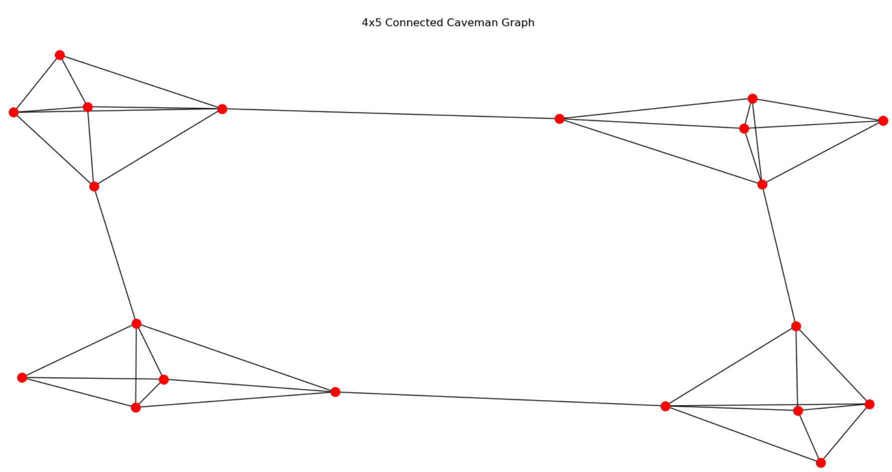

Figure 10: **4x5 Connected Caveman Graph** This is a simplified version of the connected-caveman graph graph used in section 5. This topology describes 4 cliques with 5 vertices each. An edge has been removed from each clique and used to connect to a neighboring clique along a central cycle, such that all 4 cliques form an unbroken loop. Connected caveman graphs are deterministic and perfect.