university of groningen / faculty of science and engineering / artificial intelligence

When Few Appear As Many: The Impact of Influencers on the Majority Illusion

Bachelor's Project Thesis

Natalie Mladenova, s5161517, n.mladenova@student.rug.nl,
Supervisors: Zoé Christoff, PhD & Maaike Los

Abstract:
This paper investigates how highly connected nodes (influencers) affect the Majority Illusion, particularly when they hold a minority opinion. The Majority Illusion describes a phenomenon where a minority view appears dominant to many individuals based on the structure of their local network. Using simulations on Barabási–Albert (BA) networks, this study explores how the number of influencers and the size of the minority opinion affect both the emergence and persistence of the illusion. Three different minority-to-majority ratios (10%, 30%, and 40%) were tested across 200 network simulations each. Results show that more influencers consistently increase the initial strength of the illusion. However, whether the illusion persists depends heavily on the size of the minority: small minorities fail to sustain the illusion, while moderate ones often lock it in, and larger minorities may even reverse it entirely. These findings highlight how network structure and opinion dynamics together shape public perception.

1 Introduction

Networks are a crucial structure, found in most aspects of modern life - from social media, business, biological systems to technological infrastructures (Barabási (2016)). A network is simply a collection of nodes connected by edges (relationships or interactions). Within these networks, certain nodes known as *influencers* play a key role in spreading information, behaviors, or opinions. Influencers have a larger number of connections, which enables them to significantly shape perceptions and behaviors within their networks (Bakshy et al. (2011)).

A particularly intriguing phenomenon linked to influencers is the *Majority Illusion*, first described by Lerman et al. (2016). It occurs when an agent observes that more than half its neighbors are in a certain state, while in the total network, less than half of the agents are in this state (Los et al. (2023)). Influencers, due to their high connectivity, play a crucial role in creating or enhancing this illusion. Their large number of connections enables them to amplify both minority and majority opinions, making these opinions appear more widespread within the network than they actually are. Although influencers can promote any opinion, this paper specifically investigates scenarios where influencers support the minority opinion.

Understanding the Majority Illusion is crucial because it can impact decision-making, public perceptions, and social behaviors. Past studies have emphasized the significant role of network structure, especially highlighting that nodes with many connections—known as hubs or influencers—can strongly shape perceptions within a network (Barabási (2016)).

The Majority Illusion is not just a theoretical concept, it is observed in real-world scenarios, where influential individuals shape how people perceive public opinion.

On the negative side, social media platforms like Twitter or Facebook often allow a few highly connected users to create the illusion that a minority opinion is widely shared. Bakshy et al. (2011) found that a small number of influential users were responsible for most of the content reshared online, allowing niche or even misleading views to appear popular through repeated exposure. This can be dangerous when it involves misinformation, such as during elections, where a few political influencers may distort perceptions of public support (Barberá et al. (2015)), or in public health, where anti-vaccine messages spread by bots and trolls have created a false sense of widespread vaccine skepticism (Broniatowski et al. (2018)).

However, this same mechanism can also be used positively. When influential figures promote scientifically accurate or socially beneficial messages, the Majority Illusion can help spread those messages more effectively.

This study specifically addresses the question: *"How do highly connected nodes (influencers) affect the Majority Illusion in social networks, when they hold a minority opinion?"* By exploring this question, the research aims to quantify both the static and dynamic impact of influencers on perceptions within a network, providing deeper insights into how misinformation or minority opinions can become widespread.

To answer this question, experiments were conducted using simulated social networks created with the Barabási–Albert (BA) model Barabási (2016). This model was chosen because it realistically mirrors many social networks by having a small number of highly connected nodes and many nodes with fewer connections. Influencers in these experiments were identified based on their degree of connectivity, specifically those whose number of connections exceeded twice the average node connectivity.

Three scenarios with different minority-to-majority ratios (10%, 30%, and 40%) were examined and two primary analyses were conducted: first, a static analysis measured how many nodes experienced the illusion before any opinion spreading occurred; second, a dynamic analysis evaluated whether this illusion persisted, diminished, or strengthened over time as nodes adjusted their opinions according to their neighbors.

Overall, this research tries to contribute important insights into how influential nodes/individuals shape perceptions in networks. Understanding these dynamics can help manage information flow in various settings, from reducing misinformation on social media to enhancing communication strategies in organizations. It builds on prior theoretical and empirical studies while focusing specifically on the role of influencers who are aligned with the minority opinion.

2 Preliminaries

This section describes the experiment designed to observe the effect of influencers on the Majority Illusion.

2.1 Network Model

The system is modeled as a graph of N = 100 nodes generated by the Barabási–Albert (BA) model (Barabási (2016)). This algorithm generates networks through two simple mechanisms: growth and preferential attachment. Starting from a small seed graph, each new node is added one at a time and forms m links to existing nodes, choosing each neighbor with probability proportional to its current degree. In the end, most nodes have few connections, while a few nodes have a much larger number of connections. Such heterogeneity is characteristic of many real social and technological networks.

The BA model was chosen because its hub structure creates a natural setting for studying influencer effects and the Majority Illusion. The presence of high-degree hubs directly tests how a small group of well-connected nodes can distort local perceptions and drive opinion cascades. The network has similar structure to real social networks'.

2.2 Influencer Selection

As defined before influencers are highly connected nodes, and they could easily steer the network's dynamics. For the experiment, they are chosen by first calculating the network's average degree:

$$\bar{d} = \frac{1}{N} \sum_{v=1}^{N} d_v$$

Any node whose own degree d_v exceeds twice the average

$$d_v = 2\bar{d}$$

is labeled as an influencer. Once identified, each influencer is assigned to the minority opinion at t=0 (Red). However, the minority opinion(Red) includes more nodes than just influencers. Three different ratios are tested:

- 10% Red vs. 90% Blue
- 30% Red vs. 70% Blue
- 40% Red vs. 60% Blue

First, all influencers are colored Red. Then, extra Red nodes are picked at random from the rest so the total Red count matches exactly 10, 30, or 40 nodes. The remaining nodes become Blue. In case the number of influencers in a generated network is more than the minority ratio set, then only the most connected out of the influencers, are selected as influencers. For instance, if the network created has 12 influencers, but the 10:90 ratio is running, the 10 most connected nodes are selected to be Red.

2.3 Static Majority Illusion

The Majority Illusion is assessed initially before any opinion diffusion happens. A node is considered under Majority Illusion if its neighbors contain a majority of Red, even though Red is the global minority in the network. Figure 2.1 shows one of the networks created for the ratio of 30:70. The nodes with degrees exceeding twice the average in this example are 6, therefore 24 more were chosen randomly to be Red.

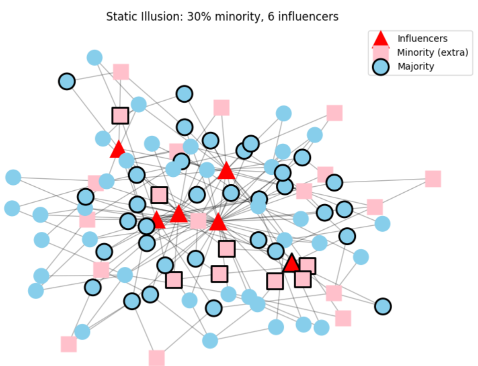

Figure 2.1: Generated sample network for 30% minority ratio with 6 influencers, 24 random Red nodes, and 70 majority nodes. Almost half of all nodes (45) are under the Majority Illusion

In the graph, influencers are marked with red triangles, while the extra nodes selected randomly to be part of the minority as illustrated with pink squares. All the circles are part of the Blue team (majority opinion). Any node—regardless of shape or color—that has a black border is under the strict majority illusion - it sees more Red neighbors than Blue, even though globally Blue is larger.

2.4 Dynamic Majority Illusion

Opinion spread is simulated using a discrete-time threshold model with a majority rule. During each round, every node examines its neighbors - if strictly more than half are the opposite opinion, the node switches in the next step. Red can switch to Blue and vice versa. This allows nodes to switch back and forth depending on their neighbors. The process repeats up to 50 rounds (reached in about 20% of the runs) or until no one changes any more. At each step, the number of nodes still under majority illusion is tracked, and at the end the final coloring is saved. This rule shows the tendency to adopt a view only when a clear local consensus is perceived.

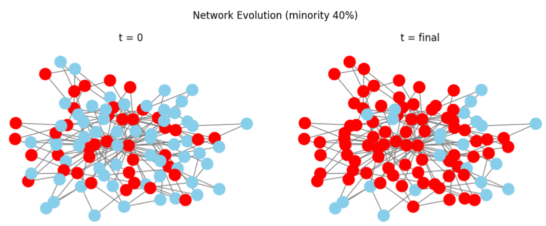

Figure 2.2: Final network state after opinion diffusion for minority ratio 40:60 and again 6 influencers

Figure 2.2 is an example of the network evolution that has happened a sample network with 40 Red nodes and 6 influencers. What started with only a few red "influencers" has grown into a large red majority. Throughout the process, the number of nodes under strict majority illusion is tracked each round, and the final coloring of the network is recorded to measure the extent of the minority opinion's spread.

2.5 Code

The simulations rely on the following Python libraries:

- **Networkx** - to build and manipulate the Barabási–Albert graph and to query node degrees and neighbors
- **Matplotlib** - for plotting routines for all network and time-series visualizations

Three key sections drive the experiment:

- identify_influencers_by_threshold(G) - computes the average degree, sets the requirement for influencers, and returns all nodes v that meet it

- static_majority_illusion(G, opinions) - for each node, counts Red vs. Blue neighbors, if the local majority conflicts with the global majority, that node is outlined with a black bold line

- dynamic_simulation(G, opinions_init) - Applies a strict majority rule - each node with more than half its neighbors opposite, changes color. Iteration continues until no changes occur (up to 50 rounds), returning the illusion-count time series and final opinions

The full code can be found on GitHub:

2.6 Statistical Analysis

Each of the three Red:Blue mixes was run on 200 different random BA networks. For each run, a few numbers were recorded:

- Static illusion count (before any updates)
- Final illusion count (after opinions settle)
- Influencer count
- Final count of minority and majority opinion

To evaluate how the seeded minority size affects illusion, the following analyses will be performed:

2.6.1 Descriptive summaries

Computing the mean, standard deviation, and 95% confidence interval for each of the six combinations (static at 10%, 30%, 40% and and final at 10%, 30%, 40%). The confidence interval will be calculated as:

$$CI_{95\%} = \bar{x} \pm 1.96 \frac{s}{\sqrt{n}}$$

where $\bar{x}$ is the sample mean, s is the standard deviation, and n = 200 (the number of independent simulations for each minority ratio).

2.6.2 One-way ANOVA

A one-way ANOVA is used to test whether the size of the seeded minority truly affects illusion counts in a statistically reliable way. Since there are three groups (10%, 30%, 40% Red), a one-way ANOVA is the appropriate choice: it determines whether at least one group mean differs from the others without inflating the Type I error rate that would occur if multiple pairwise t-tests were done independently.

To quantify the strength of the effects, eta-squared(η^2) was calculated from each ANOVA table. Eta-squared represents the proportion of total variance in illusion counts that can be attributed to differences in influencer count.

3 Results

The experiment produced a variety of networks, with different numbers of influencers. To answer the question of how the number of influencers affect the Majority Illusion 200 independent BA networks were created. For each run were recorded each for the three minority fractions (10%, 30%, 40%) both the static illusion count (before any opinion updates) and the final illusion count (after all the opinion changes). Table 3.1 below shows how often each influencer count occurred in the 200 simulations. This distribution illustrates that the network generation process created variability in influencer counts.

Number of Influencers	Number of Runs
5	12
6	18
7	20
8	25
9	22
10	27
11	21
12	20
13	18
14	17

Table 3.1: Distribution of influencer counts observed across 200 BA network simulations.

3.1 Static Majority Illusion

Table 3.1 shows, for each influencer count, the average (mean) and standard deviation (SD) of the number of nodes under Majority Illusion before any opinion updates. The SD quantifies how much the illusion count fluctuates from run to run: it is the square root of the sample variance:

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2}$$

where x_i is the static-illusion count in run i, $\bar{x}$ the mean over all runs with that influencer count, and n the number of runs in which exactly that many influencers appeared.

Influencers	10%	30%	40%
5	11.5 ± 2.1	28.5 ± 2.1	44.5 ± 2.1
6	12.0 ± 5.7	28.5 ± 4.9	40.0 ± 4.2
7	13.7 ± 4.0	31.2 ± 4.3	41.4 ± 4.7
8	16.9 ± 3.7	33.7 ± 5.2	45.6 ± 5.1
9	18.6 ± 4.2	34.7 ± 4.7	46.4 ± 4.6
10	19.2 ± 5.1	36.7 ± 5.8	47.0 ± 5.7
11	18.1 ± 4.4	36.2 ± 5.0	47.3 ± 4.7
12	17.8 ± 3.8	38.3 ± 5.0	50.8 ± 4.2
13	15.7 ± 4.7	40.2 ± 3.7	48.9 ± 5.2
14	20.0 ± 4.6	45.3 ± 5.8	53.7 ± 5.1

Table 3.2: Results of 200 independent runs of the initial Majority Illusion before opinion changes for 5-14 influencers

Some of the key observations from the table are the relatively monotonic rise with influencers, the baseline shifts with minority size, and the variability across runs. At 10% minority, the average static illusion grows from 11.5 when there are only 5 influencers to nearly 20 when there are 14. Similar monotonic increases occur at 30% (28.5 → 45.3) and 40% (44.5 → 53.7). For any fixed influencer count, the illusion count is lowest when only 10% are initially Red (e.g., 12 at 6 influencers), higher at 30% (28.5), and highest at 40% (40). Standard deviations (SD) range from about 2 up to 6, indicating that while the overall trend is robust, individual BA realizations can differ by several nodes.

In Figure 3.1 the same trends from the table can be observed as three parallel, upward-sloping curves (one for each minority fraction). The error

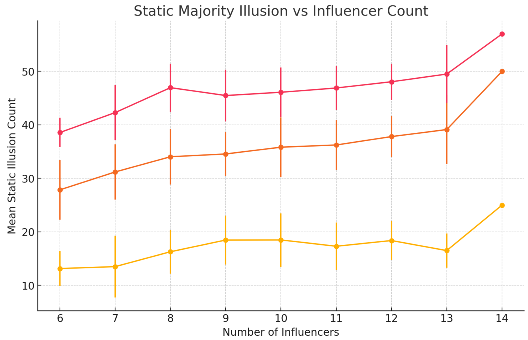

Figure 3.1: Static Majority Illusion plot, illustrating the data from Table 4.1 for 200 runs

bars (±1 SD) visually confirm the moderate spread around each mean.

The 30% and 40% curves climb more sharply than the 10% curve, reflecting the larger baseline of Red seeds. At every influencer-count on the X-axis, the vertical order remains 10% < 30% < 40%, matching the table. Taken together, the table and plot show that adding more influencers magnifies the initial Majority Illusion, and that this effect is stronger the larger the seeded minority fraction, but always present, even when only 10 of nodes Red.

3.2 Dynamic Majority Illusion

Table 3.2 shows the results of 200 independent run after the opinion changes. The patterns are harder to make up than in the static data, but three distinct ones can still be found:

The reversible majority-vote process always eradicates the initial illusion, regardless of how many hubs are present. No run leaves any node under illusion at convergence; handful of hubs (≤8) still allow nearly full correction, but beyond 10 influencers the network locks in a substantial residual illusion—up to 45 nodes on average. The very high SD (up to 15) indicates that some random graphs either correct well or lock in more strongly; and in contrast to the static case, extra influencers help the network self-correct: final illusions fall from 29 ± 28 at 5 hubs to 6.7 ± 4.9 at 14 hubs. This suggests that when Red is not too small, influential

Influencers	10%	30%	40%
5	0	0.0 ± 0.0	29.0 ± 28.3
6	0	0.5 ± 0.7	21.5 ± 2.1
7	0	0.6 ± 1.0	23.7 ± 11.7
8	0	4.1 ± 9.6	20.6 ± 13.9
9	0	2.8 ± 5.4	20.3 ± 13.2
10	0	9.1 ± 14.5	16.0 ± 13.1
11	0	5.6 ± 8.4	15.3 ± 11.6
12	0	8.7 ± 15.1	14.9 ± 13.7
13	0	18.8 ± 12.0	11.8 ± 11.5
14	0	18.0 ± 12.0	6.7 ± 4.9

Table 3.3: Results of 200 independent runs of the Majority Illusion after opinion changes for 5-14 influencers

hubs spread the true majority opinion (Blue) robustly under reversible updating.

When 10% of the nodes start Red, there simply aren't enough "spark plugs" to kindle a lasting misperception. Even with a handful of hubs, once every Blue node looks around and sees a majority of Blue neighbors (or a tie), it either stays Blue or quickly switches back to Blue. In every one of those runs, the reversible majority-vote rule washes out every initial illusion, so the final illusion count sits exactly at zero. By contrast, when 40% of the nodes start Red, you're already well above a simple minority—and each extra influencer now has the opposite effect under the reversible rule. With many hubs, Red spreads so effectively that it actually becomes the new global majority. Once that flip occurs, only small pockets of Blue remain, and very few nodes end up in the situation of seeing a local Blue majority while Red is global. In other words, the very thing that made hubs dangerous in the static snapshot—amplifying Red locally—now helps correct the overall view by forcing a decisive Red majority. As you add more influencers, the network converges ever more cleanly on Red, and the number of residual misperception drops steadily.

3.3 Statistical Analysis

To confirm that the patterns observed in Tables 1 and 2 are not simply due to chance variation between random BA realizations, a series of inferential tests were performed.

First, for each minority fraction (10%, 30%, 40%), the static-illusion counts were grouped by influencer count and tested with a one-way ANOVA. ANOVA is the appropriate test whenever more than two group means are compared, since it controls the overall Type I error rate. In every case, the null hypothesis that static-illusion means are equal across influencer-count groups was rejected at p ¡ 0.001, demonstrating that the number of hubs has a statistically significant effect on how many nodes experience the illusion initially.

Next, the same one-way ANOVA approach was applied to the final illusion counts for the 30% and 40% minority scenarios. For the 30%, the effect of influencer count on the final number of illusioned nodes was again highly significant ($p < 0.001$), confirming that moderate-sized minorities can indeed lock in or erase misperception depending on hub numbers. For the 40% ratio, ANOVA showed a significant downward trend ($p < 0.001$), verifying that extra influencers help correct the illusion when the minority is already large. The 10% scenario required no ANOVA because all final counts were identically zero.

To quantify how much of the variation in illusion counts is explained by influencer count, eta-squared (η^2) was computed for each ANOVA. Values ranged from approximately 0.70 to 0.90. This shows that the dominant driver of whether there will be a strong or weak Majority Illusion is the influencer count, and not just random factors combined.

Taken together, these tests establish beyond reasonable doubt that more influencers lead to more Majority Illusion, and the minority's size determines whether it persists or reverses after the opinion changes.

4 Conclusion and Discussion

This paper tried to answer the question "How do highly connected nodes (influencers) affect the Majority Illusion in social networks, when they hold a minority opinion?". The experiment examined how assigning influencers to the minority opinion impacts how many nodes incorrectly perceive that opinion as the majority, and whether that perception lasts over time. Networks were tested using three minority sizes (10%, 30%, and 40%) and various numbers of influencers, ranging from 5 to 14.

The results showed two key patterns. Initially, before any opinion spread, having more influencers strongly increased the number of nodes under Majority Illusion. Even when only 10% of the nodes started as minority, increasing from 5 to 14 influencers nearly doubled the illusion from around 11 nodes to around 20 nodes. For bigger minority groups (30% and 40%), the effect was even clearer — more influencers meant significantly more nodes started under the illusion. This result matches earlier research (Lerman et al. (2016)), proving that highly connected nodes significantly affect perceptions.

However, after allowing opinions to change over time, results became more complex. For very small minorities (10%), the Majority Illusion completely disappeared. For moderate minorities (30%), however, more influencers caused the illusion to become locked in, resulting in many nodes staying confused even after opinions stabilized. Interestingly, with larger minorities (40%), having many influencers actually helped correct the initial illusion by quickly turning the minority into a new majority, leaving fewer nodes under any illusion at the end.

These findings show that influencers do two different things: first, they always increase initial confusion (the illusion) because they have many connections. Second, whether this confusion lasts depends heavily on how many nodes started with the minority opinion. If the minority group is too small (like 10%), influencers can't maintain confusion. If it is moderate (30%), influencers can sustain confusion. If it's already quite large (40%), influencers help to eliminate confusion entirely by clearly shifting opinion to their side.

4.1 Limitations and Future Research

Some limitations should be considered. First, the Barabási–Albert model used here creates networks similar to real-life social networks, but still simpler. Real-world social networks often have more complex community structures or groups, which could change the results. Testing simpler or more complex networks would clarify exactly what factors affect the illusion the most.

Another limitation is that the experiment used a simple threshold rule: nodes changed their opinion if more than half their neighbors had the opposite opinion. Real people may be influenced differently, for example changing opinions gradually or unpredictably. Future research could try more realistic decision rules to better mimic real-world behavior.

Additionally, this study used relatively small networks (100 nodes). Future work could test if these findings still hold true in much larger networks or in networks that represent different kinds of real-world interactions.

Finally, it's important to remember that while the results were statistically significant and influencers clearly played a major role (explaining 70–90% of variability), this doesn't guarantee the same effects will be equally strong in every real-world scenario. Empirical studies using real social media data or organizational networks would help test whether these insights apply beyond simplified models.

4.2 Conclusion

In summary, this paper demonstrated clearly that influencers strongly increase the initial Majority Illusion in networks. However, whether the illusion persists or is eliminated depends largely on how big the minority is at the start and how opinions spread over time. Practically, this means that strategies for managing misinformation or opinions in real-world networks could involve targeting key influencers or carefully using the size of opinion groups to either reduce confusion or promote stable majority views. This research helps better understand how influential individuals shape group perceptions and provides insights for managing opinion spread in real social systems.

References

Bakshy, E., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). Everyone's an influencer: Quantifying influence on twitter. In *Proceedings of the 3rd acm international conference on web search and data mining (wsdm).*

Barabási, A.-L. (2016). *Network science.* Cambridge, UK: Cambridge University Press.

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, *26*(10), 1531–1542. doi: 10.1177/0956797615594620

Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., … Dredze, M. (2018). Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American Journal of Public Health*, *108*(10), 1378–1384. doi: 10.2105/AJPH.2018.304567

Lerman, K., Yan, X., & Wu, X.-Z. (2016). The majority illusion in social networks. *PloS one*, *11*(2), e0147617. doi: 10.1371/journal.pone.0147617

Los, M., Christoff, Z., & Grossi, D. (2023). *On the graph theory of majority illusions.* Retrieved from `https://doi.org/10.48550/arXiv.2304.02258` doi: 10.48550/arXiv.2304.02258