



Recovering from Viral Narratives
Delayed Interventions and Strategic Seeding in Polarised Networks

Alexandru Galoiu¹
Supervisor(s): Dr. Anna L. D. Latour¹

¹**EEMCS, Delft University of Technology, The Netherlands**

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 23, 2026

Name of the student: Alexandru Galoiu
Final project course: CSE3000 Research Project
Thesis committee: Dr. Anna L. D. Latour, Dr. Megha Khosla

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

Online social networks enable the rapid spread of competing narratives, including misinformation and corrective information campaigns. While previous work has studied competitive diffusion, community structure and influence maximisation separately, less attention has been given to how these factors interact when competing narratives spread according to different diffusion models.

We study the Competing Narrative Diffusion Problem, in which a viral narrative competes against a reinforcement-based intervention. We model the viral narrative using the Independent Cascade (IC) model and the corrective intervention using the Linear Threshold (LT) model. Using Stochastic Block Model and Lancichinetti-Fortunato-Radicchi networks, we investigate how community segregation, intervention timing and strategic seed placement affect diffusion outcomes.

Our results show that community structure plays a central role in determining narrative dominance. Strongly segregated networks favour LT diffusion, while increasing cross-community connectivity consistently shifts outcomes towards IC diffusion. Delayed interventions become less effective as delay increases, but remain competitive in highly segregated networks. Finally, strategic seed placement substantially improves intervention performance and allows LT diffusion to outperform IC diffusion across a much wider range of network structures.

These findings demonstrate that the effectiveness of counter-narratives depends not only on the diffusion model itself, but also on the interaction between community structure, intervention timing and seed selection. Understanding these interactions can help inform the design of more effective interventions against harmful information propagation in online social networks.

1 Introduction

The rapid spread of information through online social networks has become a major societal concern, particularly in politically polarised environments. Recent events, such as the cancelled Romanian presidential election in 2024, highlighted the visible impact of online influence campaigns and coordinated narrative propagation. Some narratives can spread quickly, after limited exposure, especially when they are sensational, emotionally engaging and easy to share. Other narratives, such as policy explanations or evidence-based messages, may require repeated exposure or social reinforcement before they are accepted. Understanding how these different types of diffusion processes interact is important for designing interventions that could potentially limit the spread of harmful narratives and support more reliable information.

A common way to study information propagation in networks is through diffusion models, which provide a framework for understanding how behaviours, opinions and narratives spread between individuals. Different diffusion processes capture different mechanisms of social influence. Two widely used models are the *Independent Cascade* (IC) and *Linear Threshold* (LT) models [Kempe et al., 2003]. These models capture different mechanisms of influence: IC represents viral spreading through repeated exposure opportunities, while LT represents reinforcement-based adoption that requires social confirmation from multiple neighbours. Diffusion models therefore represent a useful tool for analysing how different types of narratives propagate and compete within social networks.

Most work on competitive diffusion studies competition between campaigns that follow the same or closely related diffusion models. However, competing narratives in online environments may not spread in the same way. For example, a viral post may spread quickly through single exposures, while a fact-checking campaign may depend on repeated exposure, credibility and social reinforcement. This raises the question of how competition changes when the two narratives follow different diffusion dynamics.

Network structure further complicates this problem. Online social networks often contain *communities*, groups of users that are more densely connected internally than with other groups. Such community structure can influence whether information remains concentrated within one group or spreads across the wider network. In *highly segregated* networks, communities are strongly separated and only weakly connected to one another. This can affect which type of narrative gains an advantage. Therefore, understanding community structure is essential for explaining the outcomes of competing diffusion processes.

In this paper, we study the *Competing Narrative Diffusion Problem*, in which a viral narrative competes with a reinforcement-based intervention in a network structured in communities. In particular, we investigate the following research question:

“How do network structure, intervention timing and intervention strategy influence competition between viral and reinforcement-based diffusion processes?”

In order to answer this question, we study three complementary aspects of competitive diffusion. Firstly, we investigate how community structure affects competition between IC and LT diffusion. Secondly, we examine whether a reinforcement-based intervention can remain effective when introduced after a viral narrative has already begun spreading. Finally, we evaluate whether strategic seed placement can improve the effectiveness of the delayed interventions and alter their dependence on community structure.

The rest of this paper is structured as follows. In Section 2, we formally define the competing narrative diffusion problem and introduce the notation used throughout the paper. Then, in Section 3, we review the most relevant literature on competing diffusion processes and community structure. In Section 4, we describe the diffusion models and methodology.

In Section 5, we present the experimental setup, parameter settings and research questions. In Section 6, we present and discuss the experimental results. In Section 7, we reflect on the ethical considerations and reproducibility of our work. Finally, in Section 8, we summarise our findings and discuss directions for future work.

2 Problem Description

We study the *Competing Narrative Diffusion Problem* (CNDP). The problem considers two narratives that spread through the same social network while competing for their adoption by the same set of users.

2.1 Network Model

We represent a social network as an undirected graph $G = (V, E)$, where V is the set of users and E is the set of social connections between them. If $(u, v) \in E$, users u and v can influence each other in the adoption of one of the two narratives (A and B).

Each user can be in one of three states: $s(v) \in \{U, A, B\}$, where U denotes an uncommitted user, A denotes the adoption of Narrative A and B denotes the adoption of Narrative B. Once a user adopts a narrative, they cannot switch to the other narrative or become uncommitted again. This *first-adoption-wins* rule is a simplification that allows us to isolate the effects of network structure, intervention timing and seed placement.

2.2 Competing Narrative Diffusion

The two narratives originate from seed sets $S_A \subseteq V$ and $S_B \subseteq V$, which represent the initial adopters of Narrative A and Narrative B. More generally, we denote the diffusion mechanisms of the two narratives by D_A and D_B .

Different narratives may spread through different social mechanisms. Some forms of information are adopted after a single exposure, while others require repeated reinforcement from multiple social contacts before users become convinced.

In this study, we are interested in understanding how an intervention could limit the spread of harmful or misleading information in online social networks. As a result, we distinguish between two competing narratives. Narrative A represents the information that is already spreading throughout the network, such as sensational claims or misinformation. Narrative B represents a counter-narrative that aims to reduce the influence of Narrative A, for example a fact-checking campaign.

2.3 Interventions

We use the term *intervention* to refer to the deliberate introduction of Narrative B into the network as a counter-narrative to Narrative A [Budak et al., 2011]. The intervention begins from the seed set S_B and may be introduced either simultaneously with Narrative A or after a delay.

We denote the intervention delay by d , representing the number of timesteps between the introduction of Narrative A and the introduction of Narrative B. When $d = 0$, both narratives begin spreading simultaneously. Larger values of d allow Narrative A to spread for longer before the intervention begins, giving it an initial advantage.

2.4 Community Structure

We assume that the network is partitioned into a set of disjoint communities $C = C_1, C_2, \dots, C_K$, where each community $C_i \subseteq V$ contains users that are more densely connected to one another than to users outside the community.

A central property of the networks we study is *community segregation*, which describes the extent to which edges remain within communities rather than between them. Highly segregated networks contain dense communities with relatively few connections between them, while weakly segregated networks contain a larger proportion of cross-community edges.

Community structure influences how easily information can spread both within and between groups. Information may diffuse rapidly within a community but encounter difficulty crossing into other communities when few bridging connections exist. We therefore investigate whether different levels of segregation favour particular diffusion dynamics and whether community structure affects the effectiveness of delayed interventions.

2.5 Illustrative Example

In Figure 1, we provide an example of the setting studied in this paper. The figure shows a snapshot taken during the diffusion process rather than the final outcome, illustrating how competing narratives may occupy different parts of the network as they spread depending on seed placement, community structure and diffusion models.

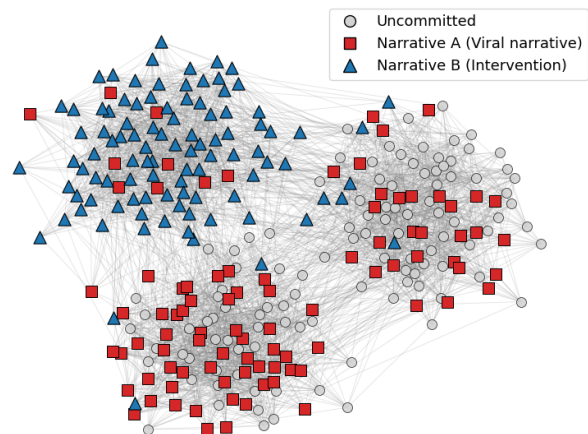


Figure 1: Example of competing narrative diffusion on a network structured in three communities. Nodes represent users and edges represent relationships. Red nodes have adopted Narrative A, the viral narrative, while blue nodes have adopted Narrative B, the intervention. Grey nodes remain uncommitted.

2.6 Problem Statement

Given a social network $G = (V, E)$, two competing narratives with seed sets S_A and S_B , diffusion mechanisms D_A and D_B and an intervention delay d , we investigate how network structure, intervention timing and seed selection influence the final adoption of the competing narratives.

3 Related Work

This section reviews prior work on competitive diffusion, community structure and influence maximisation in social networks.

Simple and Complex Contagions. An important distinction for this work is the difference between *simple* and *complex contagions* [Centola and Macy, 2007]. Simple contagions may spread after a single successful exposure, whereas complex contagions require reinforcement from multiple social contacts before adoption occurs [Centola and Macy, 2007]. This distinction suggests that different types of narratives may spread according to fundamentally different diffusion mechanisms and motivates our focus on competition between viral and reinforcement-based diffusion processes.

Competitive Diffusion. The study of competitive diffusion is an important area within the influence maximisation literature, especially between actors that follow the same diffusion mechanisms. In *Competitive Influence Maximisation* [Bharathi et al., 2007], competing actors attempt to maximise their influence under limited seed budgets. Competitive single-exposure diffusion has been examined in settings where multiple actors seek to maximise influence simultaneously [Carnes et al., 2007], while competitive reinforcement-based diffusion has been investigated in settings where adoption depends on influence accumulated from neighbouring nodes [Borodin et al., 2010].

A related problem, and one that is foundational to our work, is the containment of misinformation through competing interventions. The *Misinformation Containment Problem* [Budak et al., 2011] was introduced as the task of selecting a limited set of seed nodes for a corrective campaign in order to minimise the eventual spread of misinformation. This work shows that strategic seed placement can substantially reduce misinformation spread.

However, these studies assume that competing narratives follow the same or similar diffusion processes. In contrast, we study competition between narratives that may spread according to different diffusion models. We examine how a viral diffusion process competes with a reinforcement-based intervention and how this affects the effectiveness of the intervention.

Community Structure. Community structure is also known to strongly influence diffusion outcomes. Research on segregation, *homophily* (the tendency of individuals to form connections with others who are similar to themselves) and misinformation propagation has shown that community segregation can either help or hinder information spread. Network segregation can substantially affect the propagation of misinformation [Stein et al., 2023]. Studies of polarisation dynamics indicate that strongly segregated communities can amplify differences in beliefs and exposure to information [Azzimonti and Fernandes, 2023]. These findings suggest that community structure should play an important role in determining which competing narrative becomes dominant.

Seeding strategies. Other studies have investigated how community structure can be exploited for influence maxi-

misation. Seeding strategies that take communities into account have been shown to outperform approaches that ignore network structure in competitive LT settings [Bozorgi et al., 2017]. Similarly, leveraging community structure can produce influential seed sets while reducing computational cost [Angell and Schoenebeck, 2017]. Other work has proposed targeting influential bridge nodes in order to maximise diffusion across otherwise disconnected communities [Sadaf et al., 2022]. These studies motivate the strategic seeding strategies we examine.

Contribution. Our paper works on bridging the gap between these lines of research. We investigate how community structure, intervention timing and strategic seed placement influence the diffusion of competing narratives within the same framework. In particular, we examine whether a delayed reinforcement-based intervention can remain effective when competing against an already spreading viral narrative.

4 Approach

In this section we describe how we instantiate the Competing Narrative Diffusion Problem introduced in Section 2. We first describe the diffusion dynamics used by the two narratives, then present the network models, seed-selection strategies and evaluation metrics used throughout the experiments.

4.1 Diffusion Mechanism

In Section 2 we define the diffusion mechanisms of the two narratives abstractly as D_A and D_B . For this study, we model D_A using the Independent Cascade (IC) model and D_B using the Linear Threshold (LT) model [Kempe et al., 2003].

We chose IC for Narrative A because it represents the manner in which viral content often spreads online. IC captures this behaviour as a simple contagion process in which there is a probabilistic chance that a single successful exposure may be sufficient for the adoption of the view. This makes it a suitable model for information that can spread rapidly through a network and require little reinforcement. IC or models based on it have been widely used to model the spread of misinformation or misinformation in online social networks [Budak et al., 2011; Zhu et al., 2020; Zhang et al., 2026].

In contrast, we model Narrative B using the LT model because corrective information campaigns may depend on social reinforcement, where individuals seek confirmation from multiple trusted sources before adopting a new message. LT captures this process by allowing adoption only when the influence from neighbouring adopters exceeds a specific threshold. As a result, LT has often been used to study influence propagation and competitive diffusion in settings where adoption depends on the repeated exposures or the combined influence of multiple neighbours [Kempe et al., 2003; Borodin et al., 2010; Bozorgi et al., 2017].

Each simulation begins with Narrative A seeded in one community. Narrative A spreads using IC diffusion. Narrative B is introduced either at the same time or after a delay d , and spreads using LT diffusion [Kempe et al., 2003].

At each timestep, active IC nodes attempt to activate inactive neighbours with probability p_{IC} . LT adoption is then

evaluated for inactive nodes using the fractional threshold rule $\frac{|N(v) \cap B(t)|}{|N(v)|} \geq \theta$, where v is the node currently being evaluated, $N(v)$ is the set of neighbours of node v , $B(t)$ is the set of nodes that have adopted Narrative B by time t .

4.2 Simulation Procedure

Given a network $G = (V, E)$, seed sets S_A and S_B , diffusion dynamics D_A and D_B and intervention delay d , we simulate competing narrative diffusion in discrete timesteps.

Initially, all nodes are uncommitted except the nodes in S_A , which adopt Narrative A. Narrative A then begins spreading according to D_A . After d timesteps, the intervention is introduced by activating the nodes in S_B , which adopt Narrative B and begin spreading according to D_B .

At each timestep, both diffusion processes attempt to activate additional nodes according to their respective diffusion rules. If both narratives try to activate the same node in the same timestep, the winning narrative is chosen randomly. Once a node has already adopted a narrative, it remains in that state for the rest of the simulation.

The simulation terminates when neither narrative produces any new adopters. We then compute the evaluation metrics described in Section 4.4.

4.3 Seed-Selection Strategies

The seed-selection strategy determines how the intervention seed set S_B is chosen. More formally, a seed-selection strategy is a function $f : G \rightarrow S_B$, which maps a network G to a set of intervention seed nodes.

In all experiments, we select S_A uniformly at random. This helps us ensure that any observed differences in diffusion outcomes can be attributed to the intervention strategy rather than the initial placement of Narrative A.

To isolate the effect of seed placement, all strategies select nodes from the same intervention community. The strategies differ only in how nodes are ranked within that community.

To investigate whether interventions can benefit from strategic placement, we compare four methods for selecting the intervention seed set S_B :

- *Random*: seed nodes are selected uniformly at random from the intervention community. This acts as the baseline.
- *Degree*: seed nodes with the highest degree within the intervention community are selected. This represents a standard influence heuristic based on local reach.
- *Betweenness*: seed nodes with the highest *betweenness centrality* within the intervention community are selected. Betweenness centrality measures how often a node lies on the shortest paths between other pairs of nodes, identifying nodes that act as bridges through which information can spread between different parts of the network [Bozorgi et al., 2017; Angell and Schoenebeck, 2017]. Since IC diffusion often benefits from easily spreading through such bridge nodes, strategically placing LT seeds at these locations may reduce IC’s structural advantage and improve LT recovery.

- *Community core*: seed nodes with the highest internal degree within the intervention community are selected. This strategy is intended to favour local reinforcement, which is especially relevant for LT diffusion [Bozorgi et al., 2017; Sadaf et al., 2022].

4.4 Evaluation Metrics

The first metric we study is final adoption, defined as the number of nodes adopting each narrative at the end of the simulation.

The second metric is narrative dominance: $D = \frac{|A| - |B|}{|V|}$, where A is the set of nodes that adopted Narrative A (IC), B is the set of nodes that adopted Narrative B (LT) and V is the set of all nodes in the network. Positive values indicate IC dominance, while negative values indicate LT dominance.

5 Experimental Setup

In this section, we describe how we evaluate the problem introduced in Sections 2 and 4. We first summarise the software and hardware used for the experiments, then present the research questions that guide our evaluation. We also present the models used to generate the networks used in the experiments. Finally, we describe the experimental configurations we used to investigate the effects of community structure, intervention timing and strategic seed placement on competing diffusion processes.

5.1 Software and Hardware

Software. We implemented all experiments in Python 3.13.14 using the NumPy, NetworkX, Pandas, SciPy and Matplotlib libraries. We used NetworkX to generate SBM and LFR networks, manage graph structures, compute node centrality measures and implement the competing diffusion simulations. We used Pandas to store experimental results in data frames and perform grouping and aggregation operations on simulation outcomes across repeated runs. We used SciPy for statistical analysis, such as t-tests and linear regression modelling. We generated all figures presented in our paper with tools from Matplotlib. The source code we used for all experiments is available in the project repository ¹.

Hardware. We executed all experiments on a machine running Windows 11, equipped with an Intel i7-10750H CPU processor and 16 GB RAM.

5.2 Research Questions

Our experiments intend to answer the following research questions:

- **Q1:** How does community segregation influence competition between viral narratives and reinforcement-based interventions?
- **Q2:** Does the effect of community segregation remain consistent across different community configurations?
- **Q3:** Can a reinforcement-based intervention still remain effective when introduced after a viral narrative has already begun spreading?

¹<https://github.com/alexandrugaloiu/RPpolarisation.git>

- **Q4:** How does community structure affect the ability of delayed interventions to remain competitive?
- **Q5:** Can strategic seed placement improve the performance of delayed interventions?
- **Q6:** How does strategic seed placement alter the dependence of intervention success on community structure?

These questions guide the three experiments we present in this paper. In Experiment 1, we investigate how different aspects of community structure influence competing diffusion. In Experiment 2, we examine the effectiveness of delayed interventions and the role of community segregation in intervention recovery. In Experiment 3, we evaluate whether strategic seed placement can improve intervention performance and reduce its dependence on network structure.

5.3 Network Generation

We evaluate the competing diffusion process on synthetic networks with community structure. We made this choice because synthetic networks allow us to systematically vary the strength of community segregation, while controlling other structural properties. This makes it easier to isolate the effect of community segregation on the diffusion of competing narratives.

For Experiments 1 and 2, we generate the graphs using the Stochastic Block Model (SBM) [Holland et al.,]. An SBM assigns nodes to communities and connects pairs of nodes with probabilities determined by their community memberships. Within-community edges are created with probability p_{in} , while between-community edges are created with probability p_{out} . Increasing p_{out} reduces segregation in the network.

For Experiment 3, we use the Lancichinetti-Fortunato-Radicchi benchmark to generate the networks [Lancichinetti et al., 2008]. These networks preserve the idea of community structure but add heterogeneous community sizes and node degrees. The uneven degree distribution allows us to compare the effectiveness of different structural seed-selection strategies, unlike SBM networks, which are uniform and neutralised the effect of the strategies in preliminary tests. A mixing parameter μ controls the fraction of edges that connect outside a node’s community. Similarly to p_{out} in SBM networks, higher values of μ correspond to lower levels of segregation.

5.4 Setup of Experiment 1

We first investigate how community structure affects competition between viral narratives and reinforcement-based interventions. Specifically, we examine two aspects of community structure: the level of segregation between communities (**Q1**) and the number and size of communities (**Q2**). Our goal is to determine whether these structural properties favour one diffusion mechanism over the other.

Community Segregation Analysis

We generate SBM networks with 3000 nodes, divided equally among three communities, with fixed p_{in} and varied p_{out} . In each simulation, Narrative A is seeded in one community and

spread through IC, while Narrative B is seeded in a different community and spread through LT; both narratives are introduced simultaneously. We record final adoption counts, dominance scores and diffusion duration.

To evaluate whether the observed effects of segregation are robust to parameter choices, we repeat the experiment across a parameter grid (see Table 1). For each combination, we perform 30 independent simulation runs, resulting in a total of 12960 simulations.

Parameter	Values
Nodes	3000
p_{out}	0.005, 0.01, 0.02, 0.03, 0.05, 0.06, 0.07, 0.08, 0.09, 0.10
p_{in}	0.10, 0.15, 0.20
p_{IC}	0.03, 0.05, 0.07, 0.10
θ	0.15, 0.20, 0.25, 0.30

Table 1: Parameter settings used in the community segregation analysis.

Community Configuration Analysis

We next investigate whether the relationship between segregation and narrative dominance depends on the number and size of communities. We again use SBM networks with 3000 nodes, but vary the community configuration while keeping the diffusion parameters fixed (see Table 2).

The evaluated configurations include networks with 3, 4, 5 and 6 equal-sized communities, as well as networks with unequal community sizes (such as [500, 1000, 1500] and [250, 350, 450, 550, 650, 750]). For each configuration, we perform 30 independent simulation runs, resulting in an additional 3300 simulations.

Parameter	Values
Nodes	3000
p_{out}	0.005, 0.01, 0.02, 0.03, 0.05, 0.06, 0.07, 0.08, 0.09, 0.10
p_{in}	0.15
p_{IC}	0.05
θ	0.20

Table 2: Parameter settings used in the community configuration analysis.

5.5 Setup of Experiment 2

In this experiment, we address **Q3** and **Q4** by investigating how intervention timing influences the outcome of competing diffusion processes. Specifically, we examine whether a delayed intervention can still remain competitive after Narrative A has already begun spreading.

Narrative A is introduced at the beginning of the simulation and spreads according to IC diffusion. Narrative B acts as the intervention and is introduced after a delay d , after which it spreads according to LT diffusion. We evaluate delay values

of $d \in \{0, 2, 4, 6, 8, 10\}$ (for reference, runs in the first experiment generally ended in 10-20 timesteps). For the remaining parameters, we use the same ranges as in Experiment 1 (see Table 1). For each parameter combination, we perform 30 independent simulation runs, resulting in a total of 77760 simulations.

To determine whether the effect of intervention timing depends on community structure, we additionally repeat the experiment using the alternative community configurations introduced in Experiment 1. This results in an additional 19800 simulations.

5.6 Setup of Experiment 3

In this experiment, we address **Q5** and **Q6** by investigating whether strategic seed placement can improve the effectiveness of reinforcement-based interventions. Rather than selecting intervention seeds uniformly at random, we compare several seed-selection strategies to determine whether certain placements provide an advantage for Narrative B.

For this experiment, we use LFR networks rather than SBM networks. As discussed earlier, in LFR networks node degrees are not uniform, allowing meaningful comparisons between the randomised seeds and the different seed-selection strategies. The network contains 1000 nodes with an average degree of 20.

We used a fixed parameter configuration for p_{IC} , θ and d (see Table 3 for exact parameter values). We vary the community mixing parameter μ and the intervention delay d , allowing us to evaluate intervention strategies across different levels of community segregation. For each value of μ , we perform 30 independent simulation runs, resulting in a total of 5760 simulations.

Parameter	Value
Nodes	1000
μ	0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.35, 0.40
p_{IC}	0.05
θ	0.20
d	0, 2, 4, 6, 8, 10

Table 3: Parameter settings used in Experiment 3.

We assess statistical significance using paired t-tests comparing each seeding strategy against the random baseline.

6 Results

In this section, we present the results of the three experiments and answer the research questions introduced in the previous section. To keep the discussion focused, we only include the figures and analyses necessary to answer the research questions. Additional results and visualisations can be found in Appendix A.

6.1 Experiment 1: Community Structure

Figure 2 shows a clear transition in narrative dominance as community segregation decreases. For low values of p_{out} ,

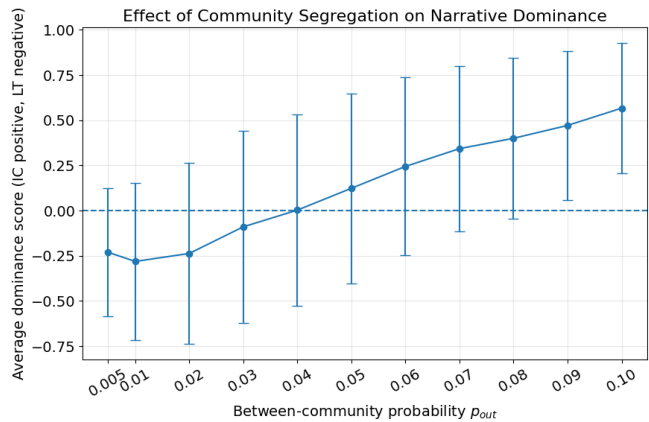


Figure 2: Average narrative dominance as a function of between-community connectivity p_{out} . Error bars indicate standard deviation across the aggregated simulation outcomes for each value of p_{out} . Increasing p_{out} shifts outcomes from LT dominance ($D < 0$) to IC dominance ($D > 0$), with the transition occurring around $p_{out} \approx 0.04$.

LT diffusion consistently outperforms IC diffusion, producing negative dominance values. As p_{out} increases, dominance gradually shifts towards IC diffusion. The transition occurs around $p_{out} = 0.04$, after which IC becomes the dominant process. The average dominance score increases from approximately -0.25 in highly segregated networks to approximately 0.60 in weakly segregated networks.

The regression analysis confirms that the observed relationship is not exclusive to this particular parameter configuration. The model explains approximately 79% of the variance in dominance outcomes ($R^2 = 0.79$). Among all predictors, p_{out} remains the strongest positive predictor of dominance ($\beta = 9.68$, $p < 0.001$). This indicates that decreasing the segregation between communities consistently shifts outcomes towards IC diffusion, even after accounting for network density and diffusion strength.

Q1: Segregation. Community segregation strongly influences competition between viral narratives and reinforcement-based interventions. Highly segregated networks favour the latter, whereas increasing cross-community connectivity increasingly favours the former.

Figure 3 presents the results obtained using alternative community configurations. Although the exact transition point varies slightly between networks, the same pattern appears in every configuration. Increasing p_{out} consistently shifts outcomes from LT dominance towards IC dominance. Networks containing more or smaller communities generally require slightly higher values of p_{out} before IC became dominant. However, no tested configuration reversed the overall trend.

Overall, the experiment shows that community structure plays a central role in determining which diffusion mechanism will dominate. Increasing cross-community connectivity consistently shifted outcomes from LT to IC winning, and this phenomenon is not tied to a particular parameter setting or community configuration.

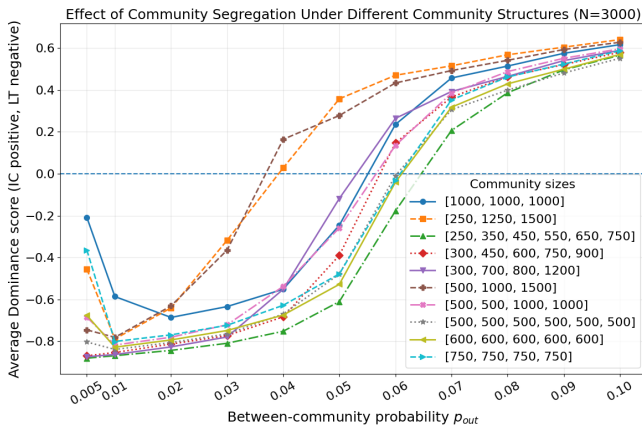


Figure 3: Effect of community segregation on narrative dominance across different community configurations. The overall relationship between p_{out} and dominance remains consistent despite variations in community size and number.

Q2: Number and size of communities. The relationship between community segregation and narrative dominance remains consistent across all tested community configurations. Community size and number influence the exact transition point, but do not alter the overall effect of segregation on diffusion outcomes.

6.2 Experiment 2: Delayed Introduction of a Counter-Narrative

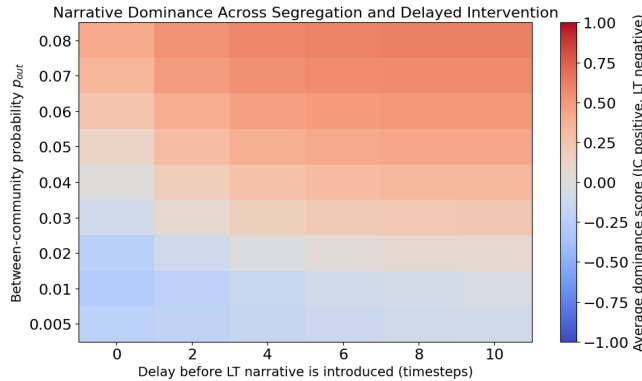


Figure 4: Average dominance score across combinations of intervention delay and community segregation. Delayed LT interventions remain competitive in highly segregated networks but become increasingly ineffective as community mixing and delay increase.

Figure 4 shows the combined effect of intervention delay and community segregation on narrative dominance. While increasing the delay consistently reduces the effectiveness of LT diffusion, the magnitude of this effect depends strongly on the level of segregation.

In highly segregated networks, LT diffusion remains competitive even when introduced several timesteps after IC diffusion had already begun spreading. Although the intervention loses some influence compared to the case where $d = 0$,

the community structure limits the advantage gained by IC diffusion.

In contrast, weakly segregated networks increasingly favour IC diffusion as the delay increases. Once IC diffusion gains an initial advantage, it became extremely difficult for the LT diffusion to recover.

Q3: Delayed Intervention Effectiveness. Delayed interventions can remain competitive against viral diffusion, but their effectiveness decreases as the delay increases. Earlier interventions consistently produce better outcomes.

Q4: Delay and Community Structure. The impact of intervention delay depends strongly on community segregation. While highly segregated networks reduce the negative effect of delayed intervention, weakly segregated networks amplify it.

6.3 Experiment 3: Seed Selection Strategies

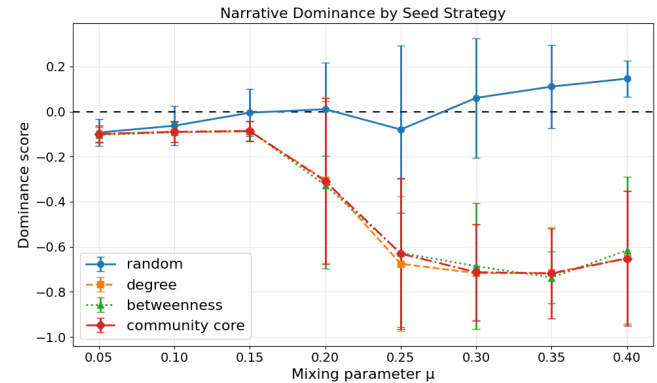


Figure 5: Narrative dominance under different LT seed selection strategies for an intervention delay of $d = 4$. Random seeding increasingly favours IC diffusion as community mixing rises, whereas strategic seeding reverses this pattern and produces consistent LT dominance $D < 0$ for $\mu \geq 0.20$

Results: Figure 5 shows that strategic seed placement substantially improves the ability of LT diffusion to recover from an initial disadvantage. At low values of μ , the difference between strategies is relatively small. However, as community mixing increases, the advantage of strategic seeding becomes increasingly visible. Placing seeds on influential nodes allowed LT to outperform IC despite the delayed introduction of the intervention.

Interestingly, strategic seeding reverses the trend observed under random seeding. While randomly seeded LT diffusion benefits from strong segregation, strategically seeded LT diffusion performs best at higher levels of community mixing. This is because influential seed nodes can use their cross-community edges (which there would be less of in strongly segregated networks) to sustain reinforcement-based diffusion outside of the starting community.

Another unexpected behaviour is that the three new strategies produce very similar outcomes. They generate nearly identical dominance scores across all values of μ . The t-tests find no statistically significant differences between the three

strategic methods ($p > 0.92$ for all three comparisons). Upon further inspection of the sets of nodes getting chosen by the strategies, there appears to be a substantial overlap between them, likely because of the way LFR graphs are generated.

However, all strategic methods significantly outperform random seeding. Across all simulations, strategic seeding increased average LT adoption by more than 400 nodes compared to the random baseline. These improvements were highly significant ($p < 10^{-40}$ for all three comparisons).

Q5: Strategic Seed Placement. Strategic seeding substantially improves the effectiveness of delayed interventions. Selecting structurally important nodes consistently outperforms random seeding across all levels of community mixing.

Q6: Strategic Seeding and Community Structure. Strategic seeding substantially alters the relationship between community structure and diffusion outcomes. In weakly segregated networks, the intervention achieves greater success and even outperforms its spread in strongly segregated networks. This suggests that targeting structurally influential nodes can fundamentally reshape diffusion dynamics.

7 Responsible Research

Privacy and Data. In this project, we investigate competing information diffusion processes using synthetic networks generated through the SBM and LFR models. Since no human participants or real social media data were involved, the research does not raise any direct concerns regarding privacy, informed consent or personal data protection.

Reproducibility. We took several measures to support research integrity and reproducibility. We conducted all experiments using publicly available algorithms and clearly documented parameter settings. We described the diffusion models, network-generation methods, evaluation metrics and experimental procedures in sufficient detail to allow the experiments to be reproduced. We used fixed random seeds to ensure consistent experimental conditions. We performed multiple independent simulation runs for each parameter configuration, in order to reduce the influence of randomness on the results.

Limitations. The fact that we used simplified diffusion models and synthetic networks is a potential source of bias. Real social networks contain unpredictable user behaviour, dynamic relationships, varying influence strengths and external information sources that are not captured by the IC and LT models. Similarly, the SBM and LFR models provide useful abstractions of community structure, but cannot fully represent the complexity of real online platforms. As a consequence, the results should be interpreted as insights into the behaviour of diffusion mechanisms rather than direct predictions of information spread in the real world.

Use of Generative AI. We used generative AI tools during the project to support software development and academic writing. We treated these tools as assistants rather than sources of the absolute truth and independently verified all claims before incorporating them into any part of the project. At no point did LLMs replace our own judgement in the design, implementation or analysis of the research. We take full

responsibility for the content and quality of both the project and this paper, including any errors or omissions that may remain. A detailed overview of this process can be found in Appendix B.

Societal Impact. Finally, the topic itself is socially relevant because the spread of misinformation and political narratives can have significant societal consequences. Our findings could potentially be used both for beneficial and harmful purposes. On the one hand, understanding how community structure, intervention timing and strategic seed placement influence diffusion may help governments, public organisations and fact-checking initiatives design more effective communication campaigns. This could help the spread of accurate information and help counter harmful narratives before they become widespread. On the other hand, the same insights could be used to amplify misleading information. An actor intending to maximise the spread of misinformation could use similar techniques to identify influential users, optimise the timing of a campaign or exploit structural properties of a social network. We believe that the benefits of studying this subject outweigh the risks. It is a necessary step towards building a safer online environment for everyone. Nevertheless, researchers and developers should consider the societal implications of these techniques before applying them in the real world.

8 Conclusions and Future Work

In this paper, we investigated the *Competing Narrative Diffusion Problem*, in which a viral narrative competes with a reinforcement-based intervention for adoption in a social network separated in communities. Our goal was to understand how *community structure*, *intervention timing* and *seed placement* influence the success of competing diffusion processes.

Our results show that community structure has a substantial impact on diffusion outcomes. Highly segregated networks generally favour reinforcement-based diffusion, while increasing connectivity between communities shifts the advantage towards viral diffusion. This pattern remained consistent across a wide range of diffusion parameters and community configurations, assuming the intervention is randomly seeded.

We also found that delayed interventions can remain effective even after a viral narrative has already begun spreading. Highly segregated networks reduce the negative impact of the intervention delay, whereas weakly segregated networks allow viral diffusion to exploit its early advantage more effectively.

Finally, we have found strategic seed placement to be highly effective. Strategies that select nodes by *degree*, *betweenness centrality* or the their position within the *community core* consistently outperformed random seeding and enabled the intervention to remain competitive across a wider range of network conditions.

Taken together, these findings suggest that the success of an intervention depends not only on the containing message, but also on where it is introduced and how the surrounding network is structured. Even when viral diffusion benefits from an early advantage, reinforcement-based interventions

can achieve similar or superior outcomes when supported by favourable structural conditions and strategically chosen seed nodes.

This work also opens multiple opportunities for future research. Firstly, the experiments could be extended to real social network data to evaluate whether the observed patterns persist in realistic settings. Secondly, more sophisticated diffusion models could be considered, such as *heterogeneous thresholds* and *weighted influence* [Kempe et al., 2003]. Thirdly, our study assumes that intervention seed nodes are selected before the viral diffusion begins. Future work could explore intervention strategies that dynamically adjust seed placement in response to the evolving state of the diffusion process.

Ultimately, the long-term goal of this line of research to support the development of more effective public communication tools. By identifying how interventions can be timed, targeted and adapted to different network structures, future systems may help public institutions and other educators to communicate accurate information more effectively, counter harmful narratives and build healthier online environments.

References

- [Angell and Schoenebeck, 2017] Angell, R. and Schoenebeck, G. (2017). Don't be greedy: Leveraging community structure to find high quality seed sets for influence maximization. In Devanur, N. R. and Lu, P., editors, *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings*, Lecture Notes in Computer Science, pages 16–29. Springer.
- [Azzimonti and Fernandes, 2023] Azzimonti, M. and Fernandes, M. (2023). Social media networks, fake news, and polarization. *European journal of political economy*, 76:102256.
- [Bharathi et al., 2007] Bharathi, S., Kempe, D., and Salek, M. (2007). Competitive influence maximization in social networks. In Deng, X. and Graham, F. C., editors, *Internet and Network Economics, Third International Workshop, WINE 2007, San Diego, CA, USA, December 12-14, 2007, Proceedings*, Lecture Notes in Computer Science, pages 306–311. Springer.
- [Borodin et al., 2010] Borodin, A., Filmus, Y., and Oren, J. (2010). Threshold models for competitive influence in social networks. In Saberi, A., editor, *Internet and Network Economics - 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, Lecture Notes in Computer Science, pages 539–550. Springer.
- [Bozorgi et al., 2017] Bozorgi, A., Samet, S., Kwisthout, J., and Wareham, T. (2017). Community-based influence maximization in social networks under a competitive linear threshold model. *Knowl. Based Syst.*, 134:149–158.
- [Budak et al., 2011] Budak, C., Agrawal, D., and Abbadi, A. E. (2011). Limiting the spread of misinformation in social networks. In Srinivasan, S., Ramamritham, K., Kumar, A., Ravindra, M. P., Bertino, E., and Kumar, R., editors, *Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011*, pages 665–674. ACM.
- [Carnes et al., 2007] Carnes, T., Nagarajan, C., Wild, S. M., and van Zuylen, A. (2007). Maximizing influence in a competitive social network: a follower's perspective. In Gini, M. L., Kauffman, R. J., Sarppo, D., Dellarocas, C., and Dignum, F., editors, *Proceedings of the 9th International Conference on Electronic Commerce: The Wireless World of Electronic Commerce, 2007, University of Minnesota, Minneapolis, MN, USA, August 19-22, 2007*, ACM International Conference Proceeding Series, pages 351–360. ACM.
- [Centola and Macy, 2007] Centola, D. and Macy, M. (2007). Complex contagions and the weakness of long ties. *American journal of Sociology*, 113(3):702–734.
- [Holland et al.,] Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. 5(2):109–137.
- [Kempe et al., 2003] Kempe, D., Kleinberg, J. M., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In Getoor, L., Senator, T. E., Domingos, P. M., and Faloutsos, C., editors, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 24 - 27, 2003*, pages 137–146. ACM.
- [Lancichinetti et al., 2008] Lancichinetti, A., Fortunato, S., and Radicchi, F. (2008). Benchmark graphs for testing community detection algorithms. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 78(4):046110.
- [Sadaf et al., 2022] Sadaf, A., Mathieson, L., Bródka, P., and Musial, K. (2022). Maximising influence spread in complex networks by utilising community-based driver nodes as seeds. In Lossio-Ventura, J. A., Valverde-Rebaza, J., Díaz, E., and Alatrística-Salas, H., editors, *Information Management and Big Data - 9th Annual International Conference, SIMBig 2022, Lima, Peru, November 16-18, 2022, Proceedings*, Communications in Computer and Information Science, pages 126–141. Springer.
- [Stein et al., 2023] Stein, J., Keuschnigg, M., and van de Rijt, A. (2023). Network segregation and the propagation of misinformation. *Scientific Reports*, 13(1):917.
- [Zhang et al., 2026] Zhang, X., Xu, Z., and Li, X. (2026). Minimizing the spread of misinformation using random influence subgraph sampling. *Expert Syst. Appl.*, 299:130017.
- [Zhu et al., 2020] Zhu, J., Ni, P., and Wang, G. (2020). Activity minimization of misinformation influence in online social networks. *IEEE Trans. Comput. Soc. Syst.*, 7(4):897–906.

A Supplementary Results

In this appendix, we provide additional figures supporting the results discussed in Section 6. In the main text, we only show the figures required to answer the research questions directly.

A.1 Experiment 1: Community Structure

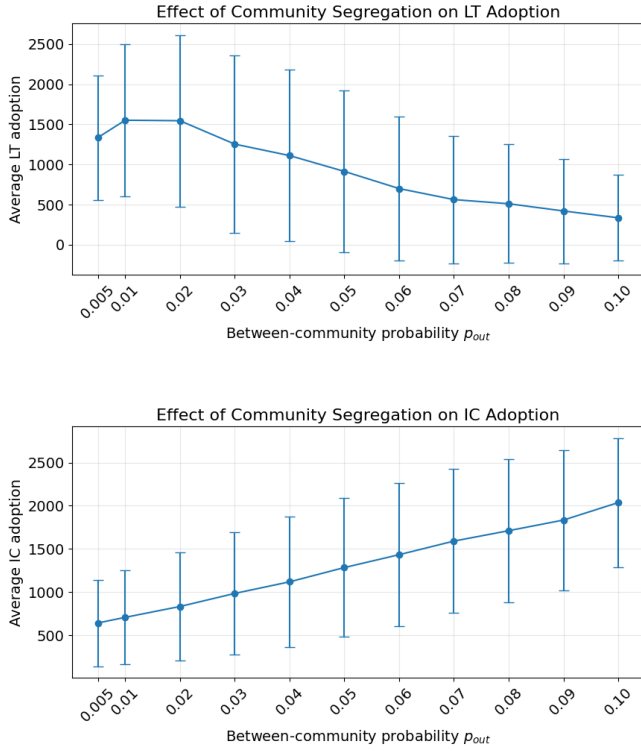


Figure 6: Final adoption of the LT and IC narratives as a function of between-community connectivity p_{out} . Error bars indicate one standard deviation across the aggregated simulation outcomes for each value of p_{out} . As p_{out} increases, LT adoption decreases while IC adoption increases, supporting the main finding that stronger cross-community connectivity shifts the advantage from reinforcement-based diffusion towards viral diffusion.

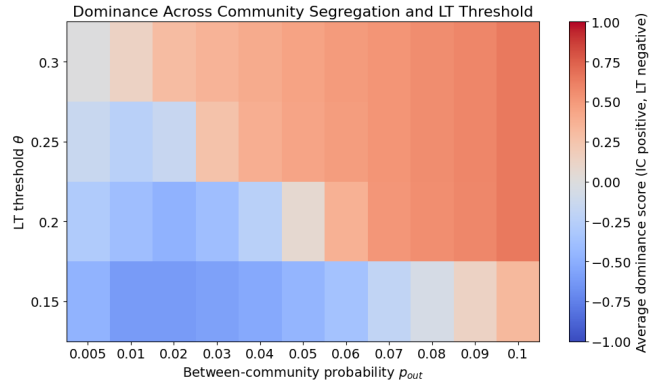


Figure 7: Dominance scores across between-community connectivity p_{out} and LT threshold θ . The transition from LT dominance to IC dominance remains visible across the tested threshold values, although higher thresholds make LT adoption more difficult and therefore shift outcomes towards IC.

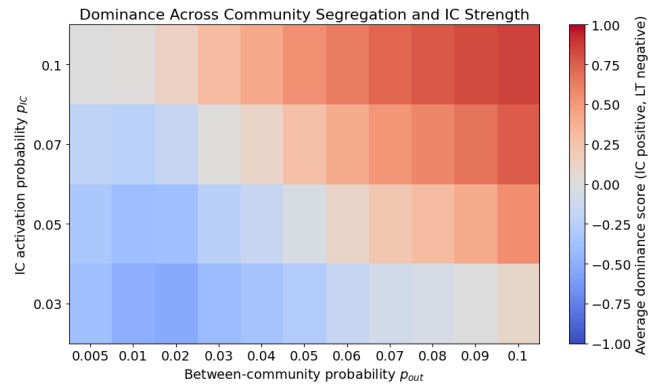


Figure 8: Dominance scores across between-community connectivity p_{out} and IC activation probability p_{IC} . Higher IC activation probabilities strengthen IC dominance, but the overall effect of increasing p_{out} remains consistent across the tested values of p_{IC} .

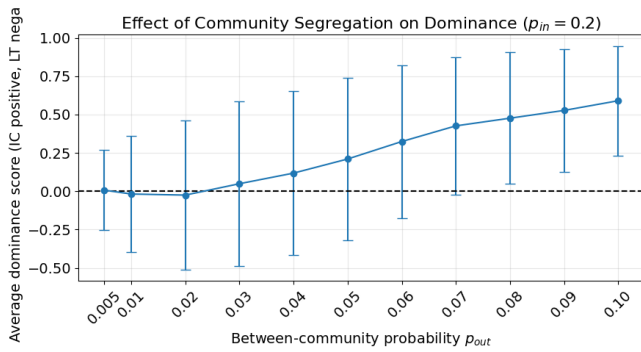
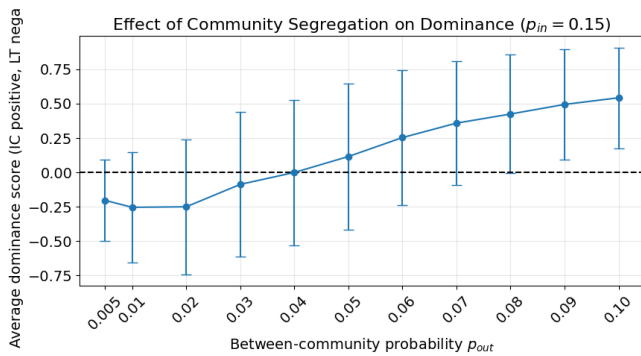
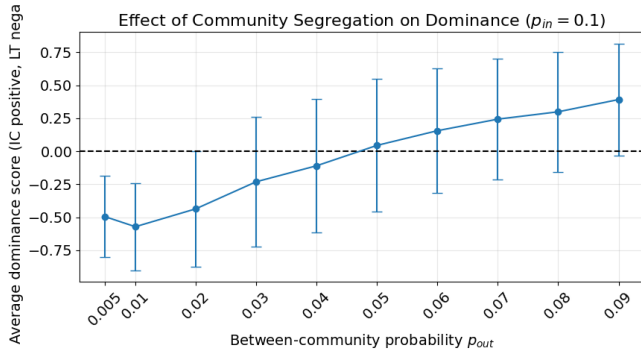


Figure 9: Effect of between-community connectivity p_{out} on narrative dominance for different values of within-community connectivity p_{in} . Positive dominance values indicate IC dominance, while negative values indicate LT dominance. Error bars indicate one standard deviation across simulation runs. Across all tested values of p_{in} , increasing p_{out} shifts outcomes from LT dominance towards IC dominance, supporting the robustness of the community-segregation effect.

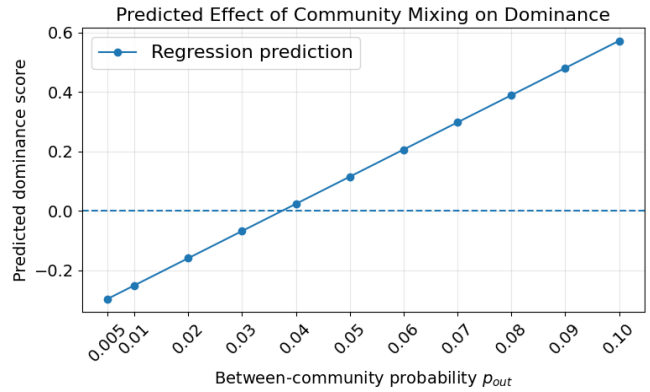


Figure 10: Predicted effect of between-community connectivity p_{out} on narrative dominance based on the regression model. The positive slope supports the main result that increasing cross-community connectivity shifts outcomes towards IC dominance.

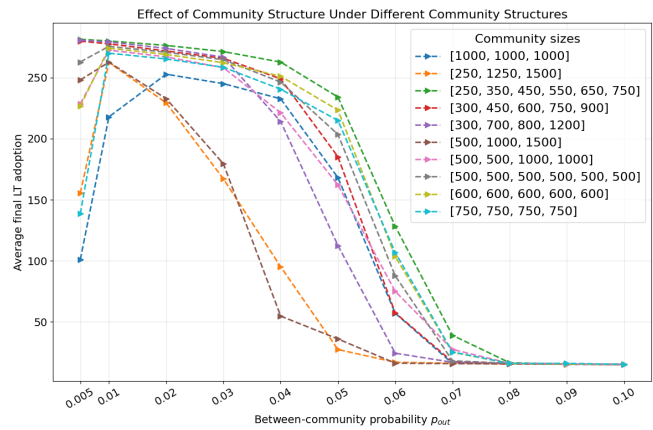


Figure 11: Final LT adoption across different community configurations. This figure complements the dominance results in the main text by showing that LT adoption generally decreases as p_{out} increases, although the exact transition depends on the number and size of communities.

A.2 Experiment 2: Delayed Interventions

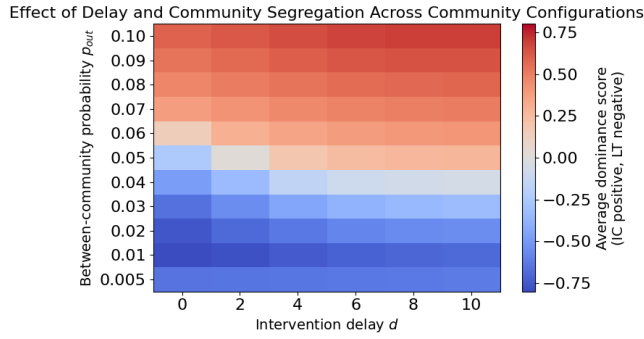


Figure 12: Average dominance score across intervention delays and between-community connectivity values for the alternative community configurations. Positive values indicate IC dominance, while negative values indicate LT dominance. The figure shows that delayed interventions generally become less effective as delay increases, but that this effect depends on community segregation. Strongly segregated networks remain more favourable to LT diffusion, whereas higher cross-community connectivity shifts outcomes towards IC dominance.

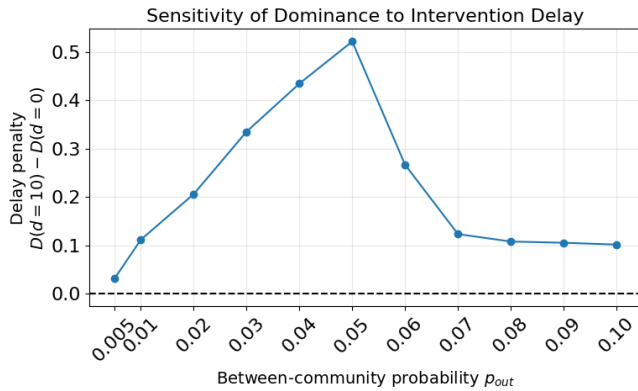


Figure 13: Delay penalty as a function of between-community connectivity p_{out} , computed as the difference between dominance at $d = 10$ and dominance at $d = 0$. Positive values indicate that delaying the intervention shifts outcomes towards IC dominance. This plot summarises how sensitive each segregation level is to intervention timing.

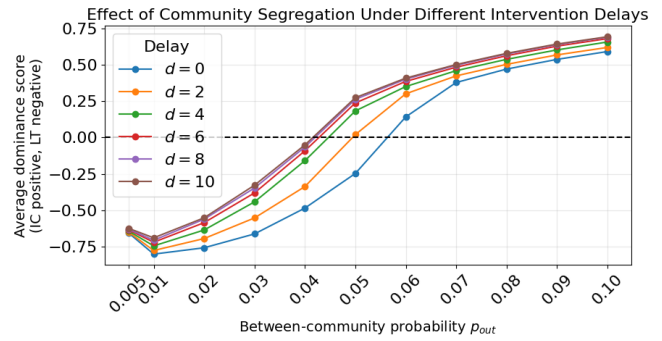


Figure 14: Average dominance as a function of between-community connectivity p_{out} for different intervention delays. Each line represents one delay value, averaged across the alternative community configurations. Increasing p_{out} shifts outcomes towards IC dominance for all delays, while larger delays generally strengthen this shift.

A.3 Experiment 3: Strategic Seed Selection

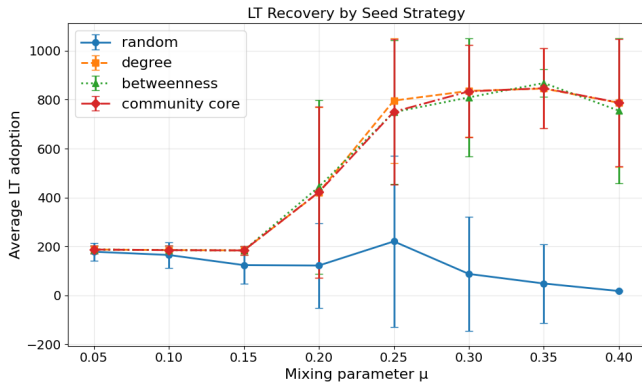


Figure 15: Average LT adoption under different seed-selection strategies for intervention delay $d = 4$. This figure complements the dominance plot in the main text by showing that strategic seed placement substantially increases the absolute reach of the LT intervention compared to random seeding.

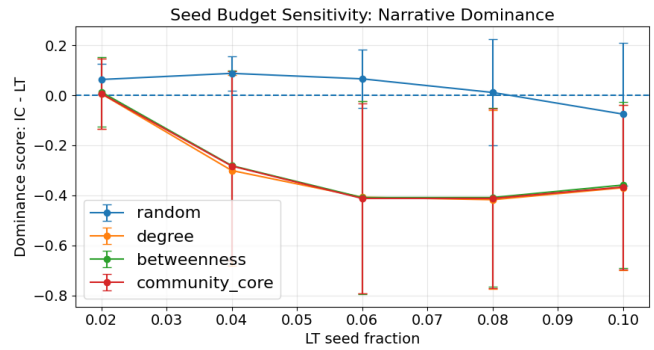


Figure 17: Sensitivity of narrative dominance to the intervention seed budget. Larger strategic seed sets shift outcomes further towards LT dominance, showing that seed budget and seed placement jointly influence intervention effectiveness.

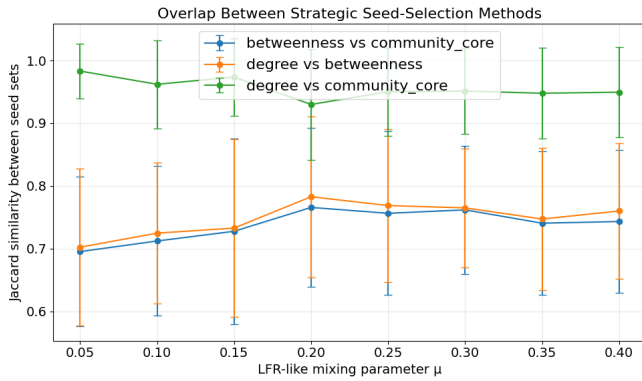


Figure 16: Jaccard similarity between seed sets selected by the strategic seed-selection methods. The high overlap helps explain why degree, betweenness and community-core strategies produce very similar diffusion outcomes in the LFR experiments.

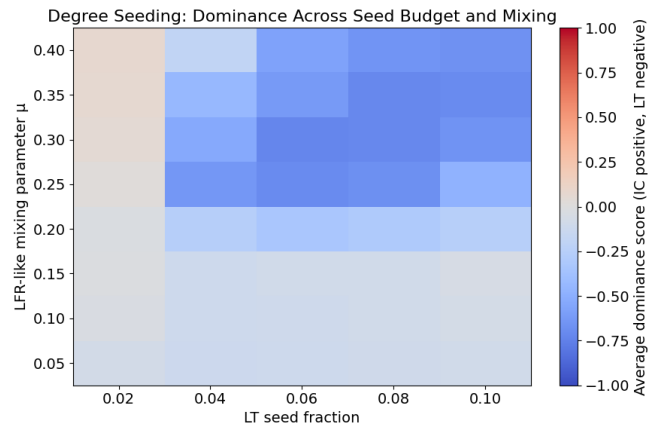
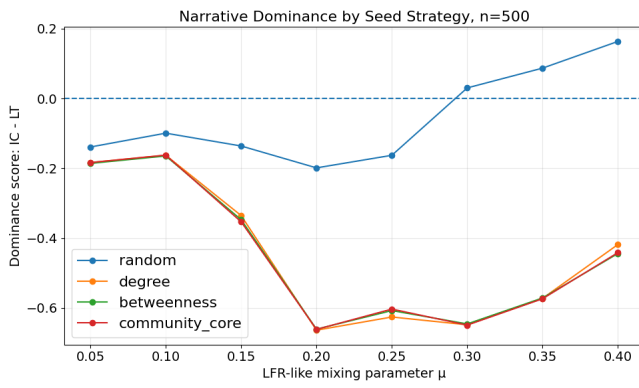
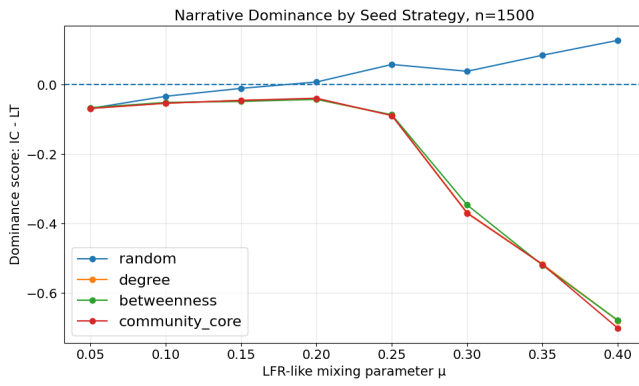


Figure 18: Average dominance across intervention delays and LFR mixing values for each seed-selection strategy. Strategic seed placement remains substantially more favourable to LT than random seeding across the tested delay values.



(a) $n = 500$



(b) $n = 1500$

Figure 19: Narrative dominance under different seed-selection strategies for alternative LFR network sizes. Positive dominance values indicate IC dominance, while negative values indicate LT dominance. Across both network sizes, strategic seed-selection methods remain more favourable to LT diffusion than random seeding, supporting the robustness of the strategic-seeding result beyond the main $n = 1000$ setting.

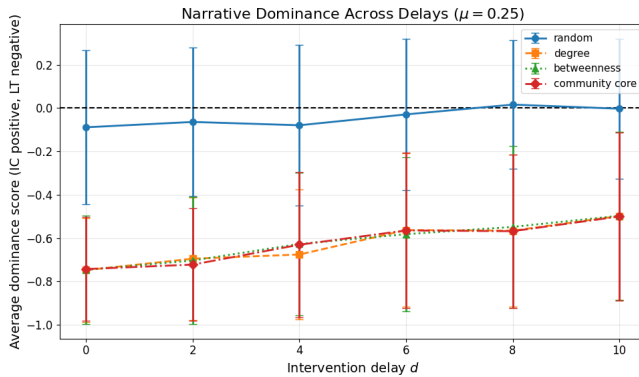
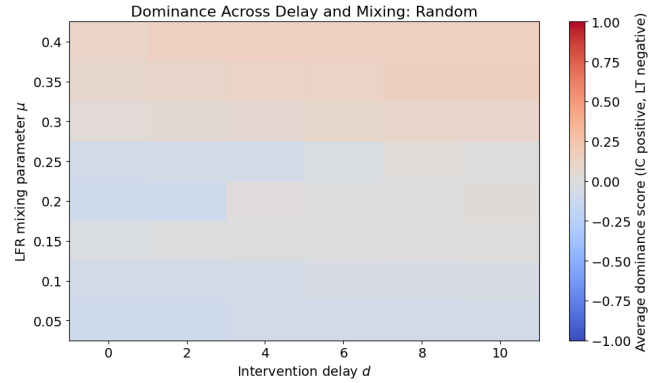
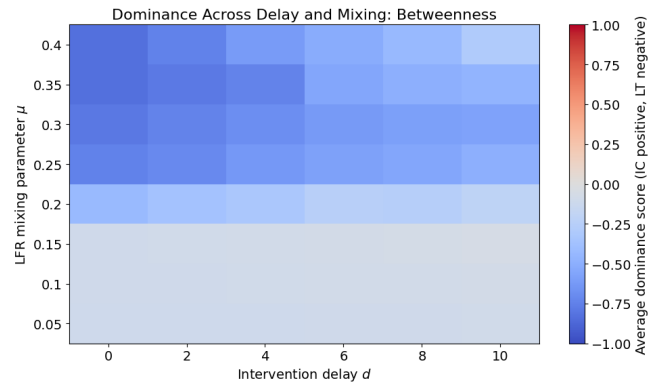


Figure 20: Narrative dominance across intervention delays for different seed-selection strategies at $\mu = 0.25$. Positive dominance values indicate IC dominance, while negative values indicate LT dominance. Error bars indicate one standard deviation across simulation runs. Strategic seed-selection methods remain substantially more favourable to LT diffusion than random seeding across all tested delay values, although the LT advantage weakens as the intervention delay increases.



(a) Random seeding



(b) Betweenness-based seeding

Figure 21: Narrative dominance across intervention delay d and LFR mixing parameter μ for random and betweenness-based seed selection. Positive values indicate IC dominance, while negative values indicate LT dominance. Random seeding remains close to the IC/LT boundary and becomes less favourable to LT as community mixing increases. In contrast, betweenness-based seeding produces stronger LT dominance across a wider range of delay and mixing values, showing that strategic placement can reduce the intervention's dependence on favourable community structure.

B GenAI Use

In this appendix, we provide an overview of how we used generative AI tools during this project. We first provide an overview of the tools that were used. Next, we present representative examples of the prompts, summaries of the outputs and how we used them.

B.1 Overview of tools used

ChatGPT 5.4. We used ChatGPT to debug code, explain algorithms and statistical methods (only where available documentation was not enough for us to understand the concepts), suggest experimental checks, improve structure and readability. We chose ChatGPT because it supports an interactive discussion and code review. It helped us explore alternative approaches and identify issues that needed further investigation. We used one chat for all miscellaneous questions that required no knowledge of the project and one chat in which we built the necessary context.

We never accepted ChatGPT outputs uncritically. We tested every code suggestion ourselves, inspected the outputs and determined whether the proposed changes aligned with our needs. We edited all writing suggestions ourselves and only incorporated ideas that accurately reflected our intended meaning.

Overleaf Built-in Rephrasing Tool. We used the Overleaf rephrasing tool for rephrasing sentences and grammar correction. We chose this tool because it is integrated directly into the writing environment and provides very convenient and accessible support during drafting. We reviewed every suggestion ourselves and decided whether it improved clarity without changing meaning.

Google Scholar Labs. We used Google Scholar Labs to identify relevant papers that were too specific to find through traditional keyword searches. We chose it because it was a feature of Google Scholar, which was the tool used in the rest of our searches, and it complemented the regular search. Google Scholar Labs was useful because it also provided quick comments explaining how the recommended papers fit our prompt. Before referencing any of the recommended papers, we read them ourselves, just like we would do with papers found through a normal search. We never cited a paper solely because the tool recommended it.

B.2 Prompts

In the following section, we provide prompts for each use case, the output received (or a summary of it, depending on the length), the validation process we applied and the effect on the project. For the sake of brevity and readability, this is not meant to be an exhaustive list, but it is representative of all the ways in which we used AI tools for this project.

1. Programming Support

Identifying why strategic seeding had little effect in SBM networks

- **Prompt:**

I am running an experiment where I compare different seed-selection strategies [insert code]. However, in there seems to be no difference in results between the

random strategy and the other strategies [insert output]. Can you help me find the issue?

- **Output:**

“ [irrelevant remarks about the code]

It is possible that the code is working correctly and that the SBM graph structure is the reason the strategies behave similarly.

In a basic SBM with equal-sized communities and uniform connection probabilities, nodes inside the same community are often structurally similar. Their expected degrees are close to each other, and there may not be strong hubs or clearly distinct bridge nodes. As a result, degree-based, betweenness-based and community-core selection may choose different nodes, but those nodes can still play very similar structural roles.

This is especially important if the LT process depends mainly on local reinforcement inside a dense community. In that case, many nodes inside the intervention community may be almost equally useful as seeds, so strategic placement will not produce a large advantage over random placement.

To check whether this is a modelling issue rather than a coding bug, you could:

- 1. compare the degree distribution of the SBM graph; inspect the variance of betweenness centrality within the intervention community;*
- 2. measure overlap between the selected seed sets;*
- 3. plot the selected seeds on a small generated graph;*
- 4. repeat the experiment on a network model with more heterogeneous degrees or more varied community sizes.*

[irrelevant remarks about other things and conclusion]”

- **Validation process:**

We applied the advice and checked that the strategies indeed selected the right nodes. While there was not necessarily a complete overlap between the selected nodes, their degrees were very close, meaning that this was indeed more of a structural problem with the network.

- **Impact on the project:**

This exchange resulted in us switching to the LFR model for the generation of the data for the last experiment. Other design decisions that resulted in a similar manner include:

- Using regression instead of t-tests in Experiment 1*
- Adding seed-overlap analysis in Experiment 3 (this was after the change to LFR)*
- Refactoring the IC/LT simulation update order (initially there was a bias for IC because it updated first at each timestep)*

- Adding robustness checks across alternative community configurations (this initially started because we were debugging Experiment 2)

Fixing a wrong function call inherited from an earlier experiment

- **Prompt:**

I am trying to run the delayed experiment over different community configurations. The experiment should use the delayed diffusion function, where Narrative A starts first and Narrative B is introduced after delay d. However, the results look almost identical to the last experiment. Here is the code: [correct code but with one function call accidentally from the previous experiment]. Can you help me find what is wrong?

- **Output:**

“ One possible issue is that the delay is not being applied inside the simulation loop. If the results are very similar for different delay values, then Narrative B may still be activated at the beginning of the simulation. I would first inspect where the B seed set is introduced and check whether the condition for activating B depends on the delay parameter.

You should also check whether the same random seed is being reused across delay values in a way that makes the networks and seed sets too similar. Similar networks are useful for controlled comparison, but you still need to make sure that the delay parameter is actually passed to the simulation function.

In the code you pasted, you are calling the function from the simultaneous experiment:

```
run_sbm_experiment_variable_communities(...)
```

but this function does not appear to use the intervention delay. For the delayed-intervention experiment, you should call the delayed version instead, for example:

```
run_sbm_delayed_experiment_variable_communities(..., delay_B=delay_B, ...)
```

[The response included an updated version of the loop in which the old function call was replaced by the delayed-intervention function and the delay parameter was passed explicitly.]

So the main thing to verify is whether the function you call actually contains the delayed introduction logic. If it does not, then changing the delay value in the outer loop will not affect the simulation.”

- **Validation process:**

The explanation given by the model was only partly correct. The issue was not caused by random seeds or by a problem in the delay logic itself. However, the response correctly pointed out that the code was still calling the function from the earlier experiment. We checked the function definitions manually and confirmed that the

function being called did not use the intervention delay parameter. After making the changes, we reran the experiment and verified that it now worked as intended.

- **Impact on the project:**

This helped us fix Experiment 2 across alternative community configurations. The overwhelming majority of bugs solved with ChatGPT were similar to this one or had similar levels of difficulty.

2. Experimental Design and Methodological Support

Using regression instead of t-tests in Experiment 1

- **Prompt:**

I want to test whether pout influences the dominance score. I was thinking of splitting the networks into strongly segregated and weakly segregated networks and then doing t-tests between the two groups. Is this a good idea?

- **Output:**

“ A t-test would work if you only had two groups and one thing changing. But in your case, you are not only changing segregation. You are also varying p_{in} , p_{IC} and θ , so a simple t-test between “strong” and “weak” networks would mix several effects together.

A better option is to use a regression model with dominance as the dependent variable and the parameters as predictors. For example: $D = \beta_0 + \beta_1 p_{out} + \beta_2 p_{in} + \beta_3 p_{IC} + \beta_4 \theta + \epsilon$.

This lets you ask whether p_{out} still has an effect after controlling for the other parameters. It also avoids having to choose an arbitrary threshold for what counts as “strong” or “weak” segregation.

You can still use t-tests later for specific pairwise comparisons, but for the main parameter sweep, regression is more appropriate. ”

- **Validation process:**

We checked whether this made sense for our data. Indeed, using t-tests would have required us to artificially split the data into groups and ignore the fact that p_{out} was continuous, which could be problematic.

- **Impact on the project:**

This exchange resulted in us adding a regression analysis to Experiment 1.

3. Literature Exploration

We used both ChatGPT and Google Scholar Labs to identify potentially relevant literature and discover papers related to specific aspects of the project.

- **Representative prompt:**

(In Google Scholar Labs) Are there papers that model Misinformation as Independent Cascade?

- **Output summary:**

List of relevant papers with short comments describing how they fit the prompt. For example: "Temporal cascade model for analyzing spread in evolving networks

A Haldar, S Wang, GV Demirci, J Oakley... - ACM Transactions on ... , 2023 - dl.acm.org

Introduces the Temporal Independent Cascade (T-IC) model, an extension of the Independent Cascade model, specifically designed to simulate and analyze the spread of various phenomena, including rumors and misinformation, through dynamically evolving networks.

Addresses Model Limitations: Explains that the T-IC model addresses the limitations of traditional static cascade models by incorporating temporal properties, such as evolving connections, dynamic propagation probabilities, and the limited active period of information like rumors.

Applications in Misinformation: Details the T-IC model's potential applications, which include identifying sentinel nodes and susceptible nodes for use in combating rumor spread and misinformation campaigns in social networks."

- **Validation process:**

Based on the comments, we chose a few papers to do a first pass on.

- **Impact on the project:**

Based on the first pass, the only paper we used in our work from that list was [Zhang et al., 2026]. Other papers we referenced and found this way are [Zhu et al., 2020; Sadaf et al., 2022; Azzimonti and Fernandes, 2023].

4. Writing and Structure Support

Rewriting the introduction around the new framing

- **Prompt:**

Can you help me rewrite this part of my introduction to be less convoluted: Narratives diffuse in different ways depending on their content and social context; Sensational, emotional and easily shareable narratives can spread quickly even with limited exposure. In contrast, narratives such as policy explanations or evidence-based messages often require repeated exposure, trust, or social reinforcement before they are accepted. Understanding how these diffusion processes interact is important for designing interventions that can limit the spread of harmful narratives while supporting the circulation of reliable information."

- **Output:**

"Rather than starting with the diffusion models themselves, consider starting from the broader observation that different types of information may spread through different social processes. Some information can spread after a small number of exposures, while other information may require reinforcement from multiple contacts before adoption occurs.

You could then motivate the use of diffusion models as abstractions of these processes. One possible framing is to distinguish between simple contagions and complex contagions, and then explain that your study examines competition between narratives that follow different diffusion mechanisms.

The introduction should focus on the motivation, research gap and research questions. Details such as SBM and LFR network generation are probably better left for the methodology section.

You may also want to emphasise three factors that are central to your study: community structure, intervention timing and seed placement. These provide a natural way to introduce the experiments without discussing implementation details too early.

Certain narrative structures may exhibit accelerated propagation dynamics following comparatively limited exposure, particularly when such narratives possess sensational characteristics, elicit pronounced emotional engagement, and demonstrate high communicability across social environments. Conversely, alternative narrative forms, including policy explanations and evidence-based messages, may necessitate repeated exposure, cumulative social reinforcement, or sustained interpretive processing prior to their broader acceptance. A systematic understanding of the interaction between these heterogeneous diffusion processes is therefore important for the design of interventions that could potentially constrain the dissemination of harmful narratives while facilitating the circulation of more reliable information."

- **Validation process:**

We used this response more as structural feedback rather than as text to include directly. The suggested wording was too rigid and complicated, it would have probably alienated readers especially in the introduction. We then rewrote the introduction ourselves based on our initial version, repurposing words or the structure of certain phrases from the AI generated version.

- **Impact on the project:**

Although we did not use the generated text directly, this exchange introduced several terms that ended up shaping the final text. Such terms include reinforcement-based interventions, *diffusion mechanisms*, *community core*, *community-structured network*, *first-adoption-wins rule*, *cross-community connectivity*, *controlled variation in community segregation*, *supported local reinforcement*. We also checked whether those wordings were taken from other papers we mention, in which case a reference was necessary (*between-ness centrality*, *simple contagion*).

B.3 Trivial or Low-Level Uses

In addition to the non-trivial examples described above, we also used generative AI tools for small technical and formatting questions. These interactions did not affect the research design, experimental results or interpretation of the findings.

Examples of such trivial prompts include:

- *How do I make LaTeX references appear as author-year citations instead of numbered citations?*
- *How do I write a footnote in LaTeX?*
- *How do I write a LaTeX figure with two subfigures stacked vertically?*
- *How do I make the legend text bigger in Matplotlib?*
- *Some parameters in this picture are not showing. How do I force Matplotlib to show all p_{out} values on the x-axis?*
- *How do I rotate x-axis labels in Matplotlib so they do not overlap?*
- *How do I assign a triangle marker to one group of nodes and a circle marker to another group in a graph visualisation?*