

An Improved ISCR Model for Rumor Propagation on Social Networks: Dynamics, Reproduction Number, and Intervention Strategies

Tianyi Xu*

Department of Mathematical Sciences, University of Nottingham, Ningbo, 315104, China

*Corresponding author: smyt4@nottingham.edu.cn

Abstract

Rumor diffusion on online social networks (OSNs) is a fast, non-linear and high-impact phenomenon that increasingly threatens public order. We extend the classical SEIR epidemic framework to a four-compartment Improved ISCR (Ignorant-Hesitant-Spreader-Stifler) model that explicitly distinguishes the cognitive hesitation stage and incorporates an official intervention intensity u . Specializing to constant control we derive, via the next-generation matrix method, a closed-form effective reproduction number $R_{\text{eff}} = \beta\sigma/[(\sigma + \delta)(\gamma + u)]$, which yields an interpretable threshold $u^* = \beta\sigma/\sigma + \delta - \gamma$ for rumor extinction. We complement the mean-field analysis with agent-based simulations on Barabási-Albert (BA) scale-free and Watts-Strogatz (WS) small-world networks, run on 5 independently generated graphs $\times$ 6 random seeds ($n=30$ per configuration) and extended until $S(t) \rightarrow 0$ to avoid endpoint censoring. We benchmark five governance policies including a novel combined early-intervention plus opinion-leader-immunization rule. Results show that (i) on BA networks, random immunization is essentially inefficient (76.7% $\rightarrow$ 74.7%); opinion-leader immunization alone reduces final reach to $17.5 \pm 14.3\%$; the combined policy further compresses it to 5.3 ± 0.3 under the same 5% immunization budget jointly with early control; (ii) topology-effect claims are robust: under random-seed initialization the BA reach drops to 45.7% but remains higher than WS's 14.5%, while hub-seeded BA still dominates; (iii) the joint policy is the only one that simultaneously achieves low mean and low variance, which is critical for governance reliability.

Keywords

Rumor propagation, ISCR model, basic reproduction number, complex networks, opinion leader, social-media governance.

1. Introduction

Online social networks (OSNs) have grown to billions of monthly active users; on platforms such as Weibo, WeChat, X (Twitter) and TikTok, an unverified message can reach millions in a matter of hours through a chain-like cascade combining one-to-many broadcast and many-to-many resharing. From the “shuanghuanglian rush” during the early COVID-19 outbreak, to repeated financial-market hoaxes that wiped billions of dollars from listed companies, to natural-disaster misinformation that hampered rescue operations, such episodes show that rumor diffusion is bursty, non-linear, and socially amplified [1], whereas ex-post debunking is often too late once the post-truth narrative has crystallised in users’ belief networks.

Two questions are central to platforms, regulators, and crisis-management agencies. (i) *When* should an authoritative refutation be released? Releasing too early wastes scarce attention budget; releasing too late lets the cascade self-organise into a hard-to-reverse meme. (ii) *Whom* should the limited immunization budget (e.g. verified labels, recommendation suppression,

account warnings) be applied to? Random sampling is fair but inefficient on heavy-tailed graphs; opinion-leader targeting is efficient but politically fragile. Despite a sizeable mathematical literature on rumor dynamics (Section 2), few studies offer an *integrated* quantitative answer that closes the loop from a continuous-time threshold condition to network-level governance policies and reports its statistical reliability.

The classical mathematical study of rumor spreading dates back to the Daley–Kendall (DK) model [2,3], an adaptation of the Kermack–McKendrick SIR framework [4], and the Maki–Thompson refinement [5]. Modern work has placed these models on complex networks, and added cognitive ingredients such as forgetting, hesitation, memory, and emotional contagion. Recent papers also explore game-theoretic interactions among official agencies, platforms, and rumor clarification strategies [6]. Most of these works, however, study a single dynamical property at a time (existence of equilibria, closed-form reproduction numbers, or one numerical scheme) and rarely cross the bridge from analytical thresholds to operational governance recommendations. The main contribution of this paper:

1) We construct an *Improved ISCR* (Ignorant–Hesitant–Spreader–Stifler) model that explicitly separates “contacted-but-undecided” hesitants from active spreaders and embeds a controllable intervention intensity u , capturing both an information-quality channel ($\sigma: \delta$) and an authority channel (u).

2) We derive, via the next-generation matrix (NGM) method, a closed-form effective reproduction number $R_{\text{eff}}(u) = \beta\sigma/[(\sigma + \delta)(\gamma + u)]$ and the explicit minimum-effort governance bound $u^* = \beta\sigma/(\sigma + \delta) - \gamma$, providing a one-line, interpretable threshold rule for resource allocation.

3) We design statistically rigorous agent-based experiments on BA and WS topologies with 5 independent graphs $\times$ 6 random seeds ($n = 30$ replicates per configuration), simulated until extinction ($S(t) \rightarrow 0$), and benchmark five policies including a *joint* “early intervention + opinion-leader immunization” rule never tested as a single package in prior rumor-modelling studies. We report mean and standard deviation throughout, and quantify the topology-seeding confound by running random-seed controls.

Section 2 reviews related rumor-propagation and immunization work. Section 3 presents the Improved ISCR model and derives R_{eff} . Section 4 reports the mean-field numerical experiments. Section 5 reports the agent-based network experiments and the sensitivity analysis. Section 6 discusses governance implications and limitations, and Section 7 concludes.

2. Related Work

Compartmental rumor models. The DK [3] and Maki–Thompson [5] stochastic rumor models established the SIR-style “ignorant / spreader / stifler” template still used today. Zhao et al. [7] added a forgetting–remembering mechanism that better matches empirical decay. Xia et al. [8] introduced a hesitating mechanism to capture the cognitive lag between hearing a rumor and deciding to spread it. Gao et al. [9] examined suppression strategies on online social networks. Our Improved ISCR model unifies the hesitation compartment of [8] with a forgetting term [7], and additionally exposes a controllable intervention intensity u that admits a closed-form threshold.

Spreading on complex networks. A long line of work on epidemic and information spreading shows that scale-free topology drastically lowers (or eliminates) the percolation threshold [10,11], while later work further demonstrated fast, hierarchical spreading through hubs in scale-free networks [12]. Watts–Strogatz small-world graphs [13] and Barabási–Albert scale-free graphs [14] are now the canonical synthetic substrates. Zanette [15] and Moreno et al. [16] studied rumor dynamics on these substrates. Our work uses both substrates and explicitly disentangles the topology effect from the seeding effect by including random-seed controls.

Targeted immunization. Cohen et al. [17] showed that on heavy-tailed networks, targeted immunization of high-degree nodes is dramatically more efficient than random immunization, sometimes by orders of magnitude. Kitsak et al. [18] further showed that influential spreaders are not always captured by degree alone, motivating robustness checks for opinion-leader targeting. We confirm this on rumor diffusion (Table 2) and, importantly, show that the variance of opinion-leader immunization alone is large; combining it with early authority intervention is what delivers a low-mean, low-variance outcome.

Position of this work. Compared with prior studies, the present work distinguishes itself in three respects. First, we close the loop from a closed-form mean-field threshold to an empirical comparison of five network policies. Second, we report mean and standard deviation across 30 replicates per configuration and disentangle topology from seeding—both rare in earlier rumor-network papers. Third, we identify and quantify a combined policy that achieves both lower mean and lower variance than any single-axis intervention.

3. The Improved ISCR Model

3.1. Compartments and Transitions

We partition users into four classes: Ignorant I (not yet exposed), Hesitant E (exposed but undecided), Spreader S (believes and forwards), and Stifler R (no longer participating). Total population $N = I + E + S + R$ is constant. The transition rules are: $I \rightarrow E$ at rate $\beta IS/N$; $E \rightarrow S$ at rate σE (belief); $E \rightarrow R$ at rate δE (skeptical exit); $S \rightarrow R$ at rate $(\gamma + u(t))S$, where γ is natural forgetting and $u(t) \geq 0$ is the official refutation intensity.

3.2. Dynamical Equations

Under the standard mean-field approximation, with homogeneous mixing in a closed population of constant size NNN , the transition rules of Section 3.1 yield the following system of ordinary differential equations (ODEs):

$$\begin{aligned}\dot{I} &= -\beta IS/N \\ \dot{E} &= \beta IS/N - (\sigma + \delta) E \\ \dot{S} &= \sigma E - (\gamma + u(t)) S \\ \dot{R} &= \delta E + (\gamma + u(t)) S\end{aligned}\tag{1}$$

The system is positively invariant on the simplex $I, E, S, R \geq 0$, $I + E + S + R = N$.

3.3. Effective Reproduction Number via the NGM Method

We now derive a closed-form expression for the effective reproduction number under a piecewise-constant intervention intensity u , using the next-generation matrix method of van den Driessche and Watmough [19]. The method decomposes the time-derivative of the “infectious” compartments—here, E and S , since these are the only states that can produce new spreaders—into a new-infection part $\mathcal{F}$ and a transition part $\mathcal{V}$, namely

$$\begin{aligned}\mathcal{F}(E, S) &= \begin{pmatrix} \beta IS/N \\ 0 \end{pmatrix}, \\ \mathcal{V}(E, S) &= \begin{pmatrix} (\sigma + \delta) E \\ -\sigma E + (\gamma + u) S \end{pmatrix}.\end{aligned}\tag{2}$$

Linearising at the rumor-free equilibrium $(I, E, S, R) = (N, 0, 0, 0)$ yields the constant Jacobian matrices

$$F = \begin{pmatrix} 0 & \beta \\ 0 & 0 \end{pmatrix}, \quad V = \begin{pmatrix} \sigma + \delta & 0 \\ -\sigma & \gamma + u \end{pmatrix}, \quad (3)$$

where V is invertible whenever $u > \gamma$. A direct computation gives

$$V^{-1} = \frac{1}{(\sigma + \delta)(\gamma + u)} \begin{pmatrix} \gamma + u & 0 \\ \sigma & \sigma + \delta \end{pmatrix}, \quad (4)$$

and therefore

$$FV^{-1} = \frac{1}{(\sigma + \delta)(\gamma + u)} \begin{pmatrix} \beta\sigma & \beta(\sigma + \delta) \\ 0 & 0 \end{pmatrix}. \quad (5)$$

The spectral radius of this rank-one matrix is its non-zero eigenvalue, hence

$$R_{\text{eff}}(u) = \rho(FV^{-1}) = \frac{\beta\sigma}{(\sigma + \delta)(\gamma + u)}. \quad (6)$$

Setting $u = 0$ recovers the basic reproduction number $R_0 = \beta\sigma/[(\sigma + \delta)\gamma]$. By the standard result of [19], the rumor-free equilibrium is locally asymptotically stable iff $R_{\text{eff}}(u) < 1$, and unstable iff $R_{\text{eff}}(u) > 1$. Solving $R_{\text{eff}}(u) < 1$ for u yields the explicit minimum-effort threshold

$$u > u^* := \frac{\beta\sigma}{\sigma + \delta} - \gamma. \quad (7)$$

Interpretation. (i) u^* is increasing in the contact rate β and in the believe-rate σ , and decreasing in the doubt-rate δ and the natural-forgetting rate γ . (ii) If natural skepticism is already strong enough that $\gamma \geq \beta\sigma/(\sigma + \delta)$, then $u^* \leq 0$ and no intervention is required: rumors die out by themselves. (iii) Equation(4) provides a one-line operational rule, since $\beta, \sigma, \delta, \gamma$ can in principle be estimated online from repost-velocity and decay-rate data. We use eq(3) as an instantaneous risk indicator in Section6 and tabulate the resulting u^* surface in Section5.5. Although derived under the constant-control linearisation, the threshold persists as a strong empirical predictor of the take-off / containment dichotomy in our piecewise-constant numerical experiments (Section4).

4. Mean-Field Numerical Experiments

Table 1. Mean-field outcomes under three intervention timings. “Final reach” is the cumulative population fraction in **R** after the dynamics terminate ($\mathbf{S}(\mathbf{T}) < 10^{-3}$).

Scenario	PeakS	Peak day	Final reach
No intervention	2156.4	57.3	95.76%
Delayed (t = 10)	4.2	10.0	0.31%
Early (t = 3)	1.3	3.0	0.09%

We set $N = 10,000$, $\beta = 0.45$, $\sigma = 0.55$, $\delta = 0.20$, $\gamma = 0.10$, with initial condition $(N - 1, 0, 1, 0)$ and a long horizon $T = 200$ days that is sufficient for $S(t) \rightarrow 0$ in every scenario; these values lie inside the empirical ranges reported by [7,8]. Equation[eq:ode] is solved with the LSODA adaptive solver. The baseline reproduction number is $R_0 = 3.30$, far above the extinction threshold, and from(4) the minimum required intervention intensity is $u^* \approx 0.23$.

Figures 1–3 and Table 1 show the non-linear premium of timeliness in the mean-field regime: each day delayed in the take-off region multiplies the final reach by orders of magnitude, while any timely $u > u^*$ collapses the outbreak.




Figure 1. ISCR compartmental dynamics under no intervention ($u = 0$). The classical S-curve drives most of the population through the spreader/stifler chain by $t = 200$.




Figure 2. Effect of intervention timing on the spreader trajectory $S(t)$. Once $u = 0.35 > u^*$, both early and delayed responses force $R_{\text{eff}} < 1$ and contain the outbreak; the early response prevents take-off entirely.




Figure 3. Basic reproduction number $R_0(\beta, \gamma)$ at $u = 0$. The black contour $R_0 = 1$ partitions the parameter plane; the white star marks the baseline operating point.

5. Agent-Based Network Experiments

5.1. Network Construction and Protocol

Algorithm 1 Agent-based ISCR simulation kernel (one repli-cate)

Require: Graph $G = (V, E)$, horizon T , immunized set $\mathcal{I} \subseteq V$, intervention onset t_u , intensity u , rates $(P_{\text{inf}}, P_{\text{bel}}, P_{\text{dbt}}, P_{\text{frg}})$, seed strategy

Ensure: Trajectory $\{(I_t, E_t, S_t, R_t)\}_{t=0}^T$

1. Initialise state $x_v \leftarrow I$ for all $v \in V$
 2. Set $x_v \leftarrow R$ for all $v \in \mathcal{I}$ // immunization (excluded from infection)
 3. Choose seed v^* : highest-degree node (hub) or uniform random
 4. $x_{v^*} \leftarrow S$
 5. for $t = 1$ to T do
 6. $u_t \leftarrow u$ if $t \geq t_u$ else 0
 7. for each spreader v with $x_v = S$ do
 8. for each neighbour w of v with $x_w = I$ do
 9. With probability P_{inf} , set $x_w \leftarrow E$
 10. end for
 11. end for
 12. for each hesitant v with $x_v = E$ do
 13. With probability P_{bel} , set $x_v \leftarrow S$
 14. Else with probability $P_{\text{dbt}}/(1 - P_{\text{bel}})$, set $x_v \leftarrow R$
 15. end for
 16. for each spreader v with $x_v = S$ do
 17. With probability $P_{\text{frg}} + u_t$, set $x_v \leftarrow R$
 18. end for
 19. Record (I_t, E_t, S_t, R_t)
 20. end for
-




Figure 4. Empirical degree distribution of one BA realization, exhibiting the expected power-law signature.

We simulate the discrete-time analogue of [eq:ode] on (i) a BA scale-free graph ($n = 2000, m = 3$) emulating Weibo-like “few-hubs/many-followers” topology, and (ii) a WS small-world graph ($n = 2000, k = 6, p = 0.1$) emulating WeChat-like acquaintance networks. Discretized probabilities are calibrated to reproduce the mean-field rates: $P_{\text{infect}} = 0.08$, $P_{\text{believe}} = 0.55$, $P_{\text{doubt}} = 0.20$, $P_{\text{forget}} = 0.10$. Each policy is run on 5 independently generated graphs $\times$ 6 random initialisations ($n = 30$ replicates), and the horizon is fixed at $T = 200$ days, beyond which $S(t) \equiv 0$ in all configurations. To guard against confounding “topology effect” with

“adversarial seeding”, we evaluate two seeding regimes: hub-seeded (highest-degree initial spreader, mimicking “celebrity origination”) and random-seeded.

We benchmark five policies on BA: NoInt, RandomImm5, OLIImm5 (top-5% degree immunised), EarlyU (constant u activated from $t = 3$), and Combined (EarlyU jointly with OLIImm5); two policies on WS for cross-topology check (NoInt, OLIImm5); and the RandomSeed-NoInt variant on each topology. Algorithm1 formalises the per-replicate simulation kernel.

5.2. Topology effects.

Figure5 contrasts the spreader curves under hub-seeded no-intervention. The BA topology produces a much larger and more sustained outbreak than WS. To check that this is not an artefact of always seeding from the highest-degree node, Table2 reports the RandomSeed variants: BA reach drops from 76.7 to 45.7 but still exceeds the WS reach (14.5), so the scale-free topology indeed amplifies diffusion in expectation, even though hub-seeding clearly compounds it.

Table 2. Network simulation results, mean $\pm$ s.d. over 30 replicates (5 graphs $\times$ 6 seeds), horizon $T = 200$.

Configuration	PeakS	Final reach
BA, no intervention	399.3 ± 32.6	76.72 ± 2.26
BA, random imm.(5%)	347.2 ± 34.3	74.66 ± 2.61
BA, early u ($t \geq 3$)	50.6 ± 17.3	22.31 ± 5.81
BA, opinion-leader imm.(5%)	26.3 ± 29.1	17.47 ± 14.32
BA, combined	2.3 ± 1.3	5.34 ± 0.26
BA, random seed (no int.)	228.6 ± 191.1	45.68 ± 37.91
WS, no intervention	35.4 ± 31.5	20.65 ± 20.68
WS, opinion-leader imm.(5%)	4.7 ± 5.3	6.03 ± 2.06
WS, random seed (no int.)	26.2 ± 26.1	14.45 ± 15.71




Figure 5. BA versus WS topologies under no intervention, hub-seeded (mean of 30 runs).

5.3. Intervention Strategies.

On BA we compare the four governance policies plus the no-intervention baseline; results are summarised in Fig.6 and Table2. Random immunization is essentially inefficient (76.7). Opinion-leader immunization alone reduces final reach to 17.5 but with very high variance (± 14.3), reflecting hub-fragility outcomes that depend on whether the hub-rich neighbourhood is fully sealed. Early intervention without targeting yields 22.3 ± 5.8 . The novel combined policy - early u jointly with opinion-leader immunization, compresses the reach to 5.34 ± 0.26 , a 93 relative reduction with the same 5 immunization budget combined with the early- u

control, and crucially the smallest variance of any policy, an essential property for governance reliability.




Figure 6. Five policies on BA, hub-seeded. The combined “early u + opinion-leader immunization” rule attains both the lowest mean and the lowest variance.




Figure 7. Final reach across all configurations (mean $\pm$ s.d., $n = 30$). The combined BA policy is the only one that simultaneously achieves low mean and low variance.

5.4. The WS small-world Result Deserves A Comment.

On WS, opinion-leader immunization yields a smaller mean reach (6.0) than no intervention (20.7), consistent with intuition. However the variances are large under no-intervention and modest under targeting, indicating that on small-world topologies a rumor either fizzles out (when the seed lands in a low-conductance neighbourhood) or finds a few long-range shortcut edges and percolates farther. Because WS lacks high-degree hubs, the immunized 5 are not dramatically more central than the average node, so the relative gain is more modest than on BA. This contrast reinforces our main finding that topology dictates which intervention is worth the budget.

5.5. Sensitivity Analysis of the Threshold u^*

Because u^* is closed-form, equation(4), we can quantify how robust our governance prescription is to misspecification of the underlying behavioural rates without re-running any simulation. Table3 shows u^* on a ± 50 grid around the baseline $(\beta, \sigma, \delta, \gamma) = (0.45, 0.55, 0.20, 0.10)$ used throughout this paper. Two patterns stand out. First, u^* is monotonically increasing in β and roughly proportional to it; doubling the contact rate roughly doubles the required intervention intensity. Second, u^* is strongly sensitive to γ in the small- γ regime, but less so for moderate γ , since natural forgetting and authority refutation enter symmetrically as $\gamma + u$. As long as the empirical contact rate β on a target platform is bounded

above (e.g. by 0.6), $u^* \leq 0.32$ suffices, so a moderate official refutation budget is enough to cross the extinction threshold.

Table 3. Sensitivity of the closed-form threshold $u^* = \beta\sigma/(\sigma + \delta) - \gamma$ to the four model parameters. All combinations except β are at the baseline values; the lowest row varies γ alone.

Varying	−50	−25	Baseline	+25	+50
β	0.025	0.130	0.230	0.330	0.430
σ	0.151	0.197	0.230	0.255	0.275
δ	0.230	0.230	0.230	0.230	0.230
γ	0.280	0.255	0.230	0.205	0.180

The sensitivity table delivers an actionable upper bound: if a regulator can keep $u \geq 0.35$ at all times when a rumor episode is detected, then for the entire ± 50 uncertainty box around our baseline parameters the system is below the extinction threshold. Combined with the agent-based finding that an early authority intervention plus opinion-leader immunization compresses both the mean and the variance of final reach, this gives a robust two-knob governance prescription.

6. Discussion: Implications for Social Governance

The combined evidence from the mean-field analysis (Sections 3,4), the network experiments (Section 5) and the parameter-sensitivity analysis (Section 5.5) supports three concrete principles for OSN platforms and regulators.

6.1. Principle 1: Real-time Monitoring Framed as R_{eff}

Real-time monitoring of an unfolding rumor episode should be re-cast in terms of the effective reproduction number R_{eff} . Observable proxies on a platform such as Weibo include the per-minute repost velocity (proxy for β), the half-life of the originating post’s engagement (related to γ), the share of replies that are “believing” versus “disputing” (proxy for $\sigma: \delta$), and the reach of competing refutation posts (related to u). A platform-side dashboard can then estimate R_{eff} in near-real-time and trigger a structured response playbook the moment it crosses 1. The closed-form threshold (4) also tells the dashboard how much intervention is needed, not only whether some intervention is needed.

6.2. Principle 2: Time-sensitive Marginal Value of Intervention

The marginal value of intervention is highly time-sensitive. In the mean-field regime, an early intervention at $t = 3$ contained the outbreak to 0.09 final reach, while delaying to $t = 10$ already raised it to 0.31 (Table 1)—a relative cost of more than $3 \times$ for one week of inaction in the take-off region. In the network regime, early intervention alone reduces final reach by a factor of 3 – 4 relative to no intervention. The closed-form threshold $u^* = \beta\sigma/(\sigma + \delta) - \gamma$ furnishes a minimal-cost rule for resource allocation, formalising the “earlier is exponentially cheaper” intuition that has long been folklore in crisis-communication practice.

6.3. Principle 3: Pre-cultivated Opinion Leaders Deliver Low-variance Outcomes

Opinion leaders should be profiled, verified, and on-boarded before crises occur, so that targeted immunization is executable in real time when an episode begins. Our network experiments show that opinion-leader immunization alone has a low mean (17.5) but a very high variance (± 14.3) on BA: governance reliability is poor, because outcomes depend on the precise hub-rich neighbourhood structure surrounding the rumor seed. The combined policy-

early authority intervention jointly with opinion-leader immunization-is the only policy that simultaneously achieves both low mean (5.3) and low variance (± 0.3). For risk-averse governance practices that must guarantee a containment level rather than merely meet it on average, this joint policy is the recommended default.

6.4. Limitations and Future Work

Several limitations should be acknowledged. (i) Model parameters are calibrated against the published literature ranges [7,8,9] rather than estimated from a specific platform’s traces; doing the latter requires de-identified API access. (ii) Algorithmic recommendation systems, paid amplification and bot networks introduce a non-stationary effective contact rate $\beta(t)$ that our piecewise-constant analysis only captures coarsely. (iii) Cross-platform spillover (e.g. from Weibo to WeChat to TikTok) is not modelled; a multilayer network extension is a natural next step. (iv) The hesitancy compartment E is treated as a single homogeneous pool, whereas in reality believers and skeptics are influenced by their personal network neighbourhoods—an opinion-dynamics coupling that could be added on top of the present epidemic backbone. (v) Strategic interaction between platform governance, rumor-mongers seeking attention, and a regulator setting compliance rules is not addressed; a game-theoretic extension along the lines of [6] could yield welfare-optimal governance equilibria. We leave these directions to future work.

7. Conclusion

This paper proposed an Improved ISCR (Ignorant–Hesitant–Spreader–Stifler) model that explicitly incorporates cognitive hesitation and a tuneable official intervention intensity u . By applying the next-generation matrix method to the linearised system at the rumor-free equilibrium, we derived a closed-form effective reproduction number $R_{eff} = \beta\sigma/[(\sigma + \delta)(\gamma + u)]$ and the corresponding minimum-effort governance threshold $u = \beta\sigma/(\sigma + \delta) - \gamma$, both of which admit direct policy interpretation: each parameter maps to an observable platform metric, enabling real-time risk monitoring via an R_{eff} dashboard. Statistically rigorous agent-based simulations on Barabási–Albert scale-free and Watts–Strogatz small-world networks (30 replicates per configuration, each run to extinction) were conducted to benchmark five governance strategies. The results demonstrate that random immunization is essentially ineffective on scale-free topologies, while opinion-leader immunization alone, despite achieving a low mean reach of 17.5%, exhibits high variance ($\pm 14.3\%$), making it unreliable for risk-averse governance. The novel combined “early u + opinion-leader immunization” policy compresses the final rumor reach to $5.34 \pm 0.26\%$ on BA networks, a 93% relative reduction under the same 5% immunization budget plus early control, and is the only strategy that simultaneously achieves both low mean and low variance across all tested configurations. The sensitivity analysis further confirms that the governance prescription is robust to $\pm 50\%$ parameter misspecification, with $u \geq 0.35$ sufficing to guarantee extinction across the entire uncertainty box. These findings provide a transferable analytic template and a practical two-knob governance prescription, real-time R_{eff} monitoring coupled with pre-cultivated opinion-leader networks, for evidence-based social-media rumor containment.

Acknowledgments

The authors thank the instructors and peers who provided guidance and constructive comments during the preparation of this manuscript. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- [2] Daley, D. J., & Kendall, D. G. (1964). Epidemics and rumours. *Nature*, 204, 1118. <https://doi.org/10.1038/2041118a0>
- [3] Daley, D. J., & Kendall, D. G. (1965). Stochastic rumours. *IMA Journal of Applied Mathematics*, 1(1), 42–55. <https://doi.org/10.1093/imamat/1.1.42>
- [4] Anderson, R. M., & May, R. M. (1991). *Infectious diseases of humans: Dynamics and control*. Oxford University Press.
- [5] Maki, D. P., & Thompson, M. (1973). *Mathematical models and applications: With emphasis on the social, life, and management sciences*. Prentice-Hall.
- [6] Agarwal, P., Aziz, R. A., & Zhuang, J. (2022). Interplay of rumor propagation and clarification on social media during crisis events-A game-theoretic approach. *European Journal of Operational Research*, 298(2), 714–733. <https://doi.org/10.1016/j.ejor.2021.06.060>
- [7] Zhao, L., Wang, Q., Cheng, J., Chen, Y., Wang, J., & Huang, W. (2011). Rumor spreading model with consideration of forgetting mechanism: A case of online blogging LiveJournal. *Physica A: Statistical Mechanics and its Applications*, 390(13), 2619–2625. <https://doi.org/10.1016/j.physa.2011.03.010>
- [8] Xia, L., Jiang, G., Song, B., & Song, Y. R. (2015). Rumor spreading model considering hesitating mechanism in complex social networks. *Physica A: Statistical Mechanics and its Applications*, 437, 295–303. <https://doi.org/10.1016/j.physa.2015.05.113>
- [9] Gao, F., Wang, X., He, Q., Qiu, L., & Huang, M. (2024). Multi-feature rumor suppression mechanism based on community division in social networks. *IEEE Transactions on Network Science and Engineering*, 11(2), 2047–2061. <https://doi.org/10.1109/TNSE.2023.3336649>
- [10] Pastor-Satorras, R., & Vespignani, A. (2001). Epidemic spreading in scale-free networks. *Physical Review Letters*, 86(14), 3200–3203. <https://doi.org/10.1103/PhysRevLett.86.3200>
- [11] Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167–256. <https://doi.org/10.1137/S003614450342480>
- [12] Barthélemy, M., Barrat, A., Pastor-Satorras, R., & Vespignani, A. (2004). Velocity and hierarchical spread of epidemic outbreaks in scale-free networks. *Physical Review Letters*, 92(17), 178701. <https://doi.org/10.1103/PhysRevLett.92.178701>
- [13] Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393(6684), 440–442. <https://doi.org/10.1038/30918>
- [14] Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. <https://doi.org/10.1126/science.286.5439.509>
- [15] Zanette, D. H. (2002). Dynamics of rumor propagation on small-world networks. *Physical Review E*, 65(4), 041908. <https://doi.org/10.1103/PhysRevE.65.041908>
- [16] Moreno, Y., Nekovee, M., & Pacheco, A. F. (2004). Dynamics of rumor spreading in complex networks. *Physical Review E*, 69(6), 066130. <https://doi.org/10.1103/PhysRevE.69.066130>
- [17] Cohen, R., Havlin, S., & ben-Avraham, D. (2003). Efficient immunization strategies for computer networks and populations. *Physical Review Letters*, 91(24), 247901. <https://doi.org/10.1103/PhysRevLett.91.247901>
- [18] Kitsak, M., Gallos, L. K., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. E., & Makse, H. A. (2010). Identification of influential spreaders in complex networks. *Nature Physics*, 6(11), 888–893. <https://doi.org/10.1038/nphys1746>
- [19] van den Driessche, P., & Watmough, J. (2002). Reproduction numbers and sub-threshold endemic equilibria for compartmental models of disease transmission. *Mathematical Biosciences*, 180(1–2), 29–48. [https://doi.org/10.1016/S0025-5564\(02\)00108-6](https://doi.org/10.1016/S0025-5564(02)00108-6)