

Disperse to ignite, concentrate to endure: a decay-driven crossover in modular complex contagions

Paul Bufort

Independent Researcher, ORCID 0009-0000-6080-1887
paulbufort@paulbufort.com

1 Introduction and related work

Getting the staff of an organization to adopt a new habit is a practical management question. When an organization wants to spread a costly practice (a tool such as an AI assistant, a work habit), should it concentrate the first adopters (seeds) in a few dense teams, or disperse them across teams? In our simulated organisations, dispersing seeds spreads the practice further, but concentrating them sustains it when people often drop practices unused by colleagues. A team seeded together can sustain adoption; a lone seed lacks this support. A complex contagion does not occur on a single exposure: threshold models need multiple exposures and reinforcements [4]. Cascades on random graphs obey a cascade condition [12]. Complex contagions require wide bridges of overlapping ties, which long-range shortcuts destroy [2, 1]. Intervention taxonomies emphasise matching interventions to context [10]. The expectation that clustered seeding should win comes from the wide-bridges results, obtained largely in populations not already embedded in cohesive groups. Other results challenge this idea: modularity only aids cascades up to a certain point [7]; dispersed neighbour seeding outperforms clustered seeding in threshold models [8]; and multiple spreaders perform best when placed apart [6]. Random ties outperform clustered ties unless interest wanes quickly [11, 9]. Where these designs vary the tie structure, we vary only seed placement and ask whether decay reverses the ordering. Real organisations are modular, with almost everyone belonging to a small, dense team, so what benefit does clustered seeding offer when every seed is already surrounded by a cluster?

2 Model and paired protocol

The generator produces two-tier organisations: a formal hierarchy (eight departments, eight people per line team on average; 249 teams, 2,000 people) and an informal influence network. It resembles a nested stochastic block model: nearly complete teams ($p = 0.9$), about six ties per person to other teams in the same department, mostly adjacent ‘sister’ teams, and about 0.45 to other departments (silo = 0.85), plus sparse random ties.




Fig. 1. Synthetic data: 50 paired organisations per cell. (a) Paired difference in terminal adoption, random minus cluster (pp), across mean threshold $\bar{\theta}$ and seed budget; Holm-corrected difference and TOST families, grey = equivalent within ± 2 pp. (b) Same contrast with decay (95% CIs; hollow marker = uncertain). r : reinforcement required for retention.

Agent i has a threshold of $\theta_i \sim \text{Beta}(\mu\kappa, (1-\mu)\kappa)$, where $\mu = 0.30$ and $\kappa = 20$. Agent i is an innovator ($\theta_i = 0$) with probability 0.025. Exposure is the credibility-weighted share of adoptions:

$$\text{share}_i(t) = \frac{\sum_{j \in N(i)} w_{i,j} a_j(t-1) + b(t)}{\sum_{j \in N(i)} w_{i,j} + b(t)},$$

where $N(i)$ is i 's contact set, $w_{i,j}$ the weight of j for i , $a_j(t-1)$ the adoption of j at $t-1$, and $b(t)$ a communication term used only in the broadcast baseline (0.3 in its first step, 0 thereafter).

Credibility is 1.0 for close cooperation partners, 0.7 for managers, 0.6 for distant peers and 0.3 for communication partners.

Adoption requires being ready ($\text{share}_i \geq \theta_i$ and $\text{share}_i > 0$), willing (Bernoulli, $p = 0.85$, drawn once at $t = 0$) and able (1.0 here) [3]. An adopter relapses with probability ρ per step as long as share_i is less than $r\theta_i$, where r sets how much reinforcement retention requires. Relapsed agents may re-adopt and relapse is off ($\rho = 0$) unless stated. Seeds adopt at $t = 0$ regardless of readiness and willingness, then follow the same rules as other agents. Our main measure is terminal adoption: adopters as a share of all agents when runs end at an absorbing state (no agent can adopt and no adopter is below $r\theta_i$) or after 100 steps. Cumulative reach (secondary) is the share of agents that have adopted by a given time. Four seeding strategies are compared: random, champions (top informal degree), cluster (whole teams) and one-per-team (round-robin coverage), each with a budget of 100 seeds. These are rules fixed in advance, with no claim of optimality [5, 8]. Broadcast uses no seeds and serves as the baseline.

Paired protocol. In each condition cell, each strategy is run on the same 50 new organisations using the same random numbers for agent selection, so contrasts are paired differences within organisations; pairing is used as a confounding control, not for precision. We pre-specified the two confirmatory families (regime map and decay grid) in the decision log after an exploratory panel ($n = 12$ organisations per cell, not used in Fig. 1). Holm correction was applied within each family to the difference and the TOST equivalence against a practical band of ± 2 percentage points (pp; ≈ 40 people). Each cell receives one of three verdicts: ‘win’, ‘equivalent’ (within ± 2 pp, for the tested cells and model), or ‘uncertain’. Seeds were frozen before execution.

3 Results

In the headline scenario, random seeding achieves a terminal adoption rate of $72.0\% \pm 9.2\%$ (mean $\pm$ SD), compared to $32.4\% \pm 11.0\%$ for cluster seeding, a paired difference of $+39.7$ pp (95% CI [35.4, 43.9]). Degree targeting (champions) adds about 3.9 pp (descriptive, CI [0.6, 7.1]). One-per-team coverage shows no clear difference from random seeding (-1.4 pp, CI [$-5.0, +2.2$]). On average, random already places seeds in 82 of the 249 line teams, compared to 13 for the cluster. Figure 1a shows that dispersion dominates in 14 out of the 40 cells, with the diagonal ignition band ranging from $+2.8$ to $+42.6$ pp (boxed cell: an independent paired sample). The strategies are equivalent to ± 2 pp (TOST) in 23 cells (saturation with low $\bar{\theta}$ and starvation with high $\bar{\theta}$ /low budget), and three cases are uncertain. No cell shows a detected cluster advantage of practical interest (≥ 2 pp); the largest observed cluster edge is 0.8 pp. Terminal adoption under broadcast was $5.5 \pm 3.4\%$ (0.0% in the 12 tests without innovators; independent panel).

As Figure 1b shows, under relapse the dispersion advantage decreases monotonically in ρ and changes sign between the tested values. When retention requires full reinforcement ($r = 1$), the paired difference falls from $+8.4$ percentage points [4.2, 12.7] at $\rho = 0.10$ to -5.3 pp [$-9.0, -1.5$] at $\rho = 0.25$ and -11.4 pp [$-14.8, -8.0$] at $\rho = 0.40$. With weaker coupling ($r = 0.5$), reversal is seen only at the highest tested rate (-7.7 pp [$-11.5, -3.9$] at $\rho = 0.40$). All runs at $\rho \geq 0.25$ ended in absorbing states, ruling out a horizon artefact. Clustered adoption barely moves (32.4% to 29.3% under the harshest setting), while dispersed adoption collapses (72.0% to 17.9%). For ($r = 1, \rho = 0.25$), the terminal contrast has been reversed while the contrast in cumulative reach (secondary, descriptive) remains inconclusive (-1.8 pp [$-5.4, +1.8$]; clearly negative only at $\rho = 0.40$: -7.4 pp [$-10.7, -4.1$]); the terminal reversal is therefore driven mainly by differential retention, the share of ever-adopters who remain adopters (0.99 for cluster vs. 0.86 for random).

The advantage persists across organisational variants, varying one parameter at a time ($N = \{500, 2000, 8000\}$, team size = $\{5, 12\}$, silo = $\{0.5, 0.7, 0.95\}$): the paired difference ranged from $+14.7$ to $+48.3$ pp, and no confidence interval crossed zero. These variants were tested without relapse; the crossover was measured in a single organisational setting. Applied to the real email-Eu-core

topology (986 nodes, 42 ground-truth departments, all behavioural attributes synthetic), the mean increase over 50 paired attribute draws is 13.5 points (95% CI [1.0, 26.0]). Adoption there is bimodal: a low mode (median 9%) and a high mode (median 85%). This replicates the sign without relapse on a real modular topology; it does not empirically validate the mechanism.

4 Discussion

A random seed, usually alone in its team, starts far below its threshold and may relapse immediately. A fully seeded team can keep its members well above their thresholds. Dispersed seeding must then activate teams before its seeds relapse, which becomes harder as ρ increases. The strategy most efficient for spreading a practice is thus not necessarily the most efficient for maintaining it. A local predictor of team ignition, pre-specified in the same decision log, failed its validation criterion (AUC 0.64 vs. 0.70) and does not sufficiently explain dispersed activation; team-level trajectories remain for future work. All of this holds within an uncalibrated synthetic model: thresholds unobserved, decay parameters set on a grid, time abstract. Reproducibility: code, decision log (D1–D23) and reproduction pipeline at <https://github.com/PaulBafort/adoption-sim> (AGPL-3.0); every reported estimate is machine-extracted from the versioned results and checked by an audit script.

References

1. Centola, D.: The spread of behavior in an online social network experiment. *Science* 329(5996), 1194–1197 (2010)
2. Centola, D., Macy, M.: Complex contagions and the weakness of long ties. *Am. J. Sociol.* 113(3), 702–734 (2007)
3. Coale, A.J.: The demographic transition reconsidered. In: *Proceedings of the International Population Conference, Liège*, vol. 1, pp. 53–72. IUSSP, Liège (1973)
4. Granovetter, M.: Threshold models of collective behavior. *Am. J. Sociol.* 83(6), 1420–1443 (1978)
5. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: *Proc. 9th ACM SIGKDD (KDD '03)*. pp. 137–146 (2003)
6. Kitsak, M., et al.: Identification of influential spreaders in complex networks. *Nat. Phys.* 6, 888–893 (2010)
7. Nematzadeh, A., Ferrara, E., Flammini, A., Ahn, Y.Y.: Optimal network modularity for information diffusion. *Phys. Rev. Lett.* 113(8), 088701 (2014)
8. Ou, J., et al.: Influence maximization under limited network information: Seeding high-degree neighbors. *J. Phys. Complex.* 3(4), 045004 (2022)
9. Sassine, J., Rahmandad, H.: How does network structure impact socially reinforced diffusion? *Organ. Sci.* 35(1), 52–70 (2024)
10. Valente, T.W.: Network interventions. *Science* 337(6090), 49–53 (2012)
11. Wan, A., Riedl, C., Lazer, D.: Diffusion of complex contagions is shaped by a trade-off between reach and reinforcement. *Proc. Natl. Acad. Sci. USA* 122(28), e2422892122 (2025)
12. Watts, D.J.: A simple model of global cascades on random networks. *Proc. Natl. Acad. Sci. USA* 99(9), 5766–5771 (2002)