Controlling Complex Contagions

Alastair Langtry DJ Thornton

March 31, 2026

Abstract

Many social and economic behaviors, from participating in a protest to adopting a new technology, require reinforcement from multiple peers to be worthwhile. The spread of these behaviors is often called *complex contagion*. We develop a tractable model of complex contagion on random networks, where agents take an action only when enough peers do too. We identify a tipping point in behavior: a level of network density where participation (the fraction taking the action) jumps discontinuously. And we characterize how the location and size of this tipping point depends on heterogeneity in connectivity among agents. In a setting where a principal chooses network density and dislikes participation (such as a government who dislikes protest), the principal is pushed into "extreme" choices. She either chooses no participation, or participation significantly above the tipping point: to accept any participation is to accept a lot of it. The principal is "in for a penny, in for a pound".

JEL Codes: D85, Z13. *Keywords:* contagion, network design, random networks, threshold games, social norms.

1 Introduction

In many economic settings, an individual will take an action only when enough of their peers do as well. A citizen joins a protest only when a critical mass of friends will provide "safety in numbers"; a person adopts a new fashion, workplace practice, or social norm only when enough peers have already done so. These behaviors spread through social networks via a process of mutual reinforcement, and are known as *complex contagions*.[1] They play a central role in settings ranging from corporate culture and technology adoption to political mobilization and the evolution of social norms.[2]

An important feature emerging from empirical work is that behavior in these settings can change suddenly in response to modest changes in the environment. Controlled network experiments show that a committed minority of almost 25% can flip a prevailing social convention. Just below this cutoff, change rarely occurs, while just above it, adoption of the new behavior spreads rapidly (Centola et al., 2018). This "tipping point" dynamic (Schelling, 1978)—where modest changes in connectivity or early participation produce out-sized changes in outcomes—has also been observed in political mobilization (Madestam et al., 2013; Enikolopov et al., 2020).

A more limited empirical literature investigates the network design choices of principals—such as governments or firms—who have some control over the network. Such principals sometimes adopt "extreme" strategies. King et al. (2013) document that Chinese government censorship targets the coordination channels through which collective action can spread, disabling mass coordination against the regime. Internet shutdowns during protests and contested elections represent an even starker choice. They dramatically reduce connectivity: forgoing the economic benefits in order to suppress the *possibility* that coordinated and self-sustaining collective action can arise (Access Now, 2022; V-Dem Institute, 2024).

[1] This language was introduced by Centola and Macy (2007) to describe contagion processes requiring reinforcement from multiple sources, in contrast to "simple" contagions that spread through a single contact. Our model captures this through an absolute participation threshold $k \geq 3$; see Section 6 for a discussion of alternative approaches.

[2] A substantial body of empirical work has documented the importance of social reinforcement in settings as varied as the adoption of new agricultural technologies (Bandiera and Rasul, 2006; Beaman et al., 2021), communication technologies (Björkegren, 2019), birth control use (Munshi and Myaux, 2006), health behaviors (Christakis and Fowler, 2007, 2008), protest participation (Larson et al., 2019; González, 2020; Bursztyn et al., 2021), and public health interventions (Airoldi and Christakis, 2024).

Despite the accumulating evidence on tipping points and complex contagion, existing theories do not give us a useful map from network structure to behavior. Fixed-network analyses deliver rich microfoundations but struggle to deliver clean comparative statics with respect to network structure, while more tractable approaches push the network into the background. As a result, we still lack a simple benchmark that links explicit network structure to equilibrium behavior in complex contagion environments. This precludes clean analysis of a principal's optimal design problem in this environment—in particular, one that can examine whether the extreme policies highlighted by some empirical work are always necessary, or dependent on the environment.

We provide such a benchmark. Our model clarifies how complex contagion environments generate tipping points in agents' behavior, and how this in turn can lead to extreme design choices by a principal. We highlight the crucial role of the broader network structure in determining the sharpness of the tipping points and hence the extremeness of the principal's optimal policy. More concretely, we consider a setting where the principal can control the overall connectivity of a network—the average number of links—formed among a large number of agents, but cannot control its exact structure. Agents form links at random but have heterogeneous propensities to form links. This captures the idea that relationships form organically as a result of chance meetings: they are not fully chosen by agents, nor can they be micromanaged by a principal. The principal has concave benefits from network connectivity directly. But connectivity also opens up the possibility that agents participate in some collective action that the principal dislikes. The action in question is a complex contagion. Agents will only participate if at least some number, k (a participation threshold), of their friends in the network do so too. So the principal faces a fundamental tension: interaction among agents is productive, but it also creates the capacity for collective action.

Our first set of results characterizes the relationship between network connectivity and equilibrium participation—the fraction of agents who participate in the collective action. We show that, under minimal regularity conditions on the distribution of propensities to form links and for any participation threshold $k \geq 3$, there is a critical cutoff for the level of connectivity: below this cutoff, participation is zero; above it, a positive fraction of agents participate. Participation emerges *discontinuously*—it jumps from zero to a strictly positive level. As connectivity increases beyond that critical cutoff, participation grows monotonically. Separately, the critical cutoff is increasing, and the level of participation decreasing in the participation threshold k. The more friends needed to "support" participation, the fewer agents will participate.

A key finding is that the *size* of the discontinuous jump depends on heterogeneity of agents' propensities to form links in an economically interpretable way. In a homogeneous benchmark where all agents form links at the same rate, the jump in participation when $k = 3$ is approximately 0.27: whenever a self-sustaining group exists, it must contain at least a quarter of all agents. We show that this large jump is robust to small amounts of heterogeneity in agents' propensities to form links: sufficiently homogeneous environments always produce large jumps. In contrast, in environments with strong heterogeneity—where a vanishingly small "core" of agents form links at a much higher rate than the rest—the jump can be made arbitrarily small. We provide an explicit example in which heterogeneity confines participation to a small "core".

Figure 1 provides a graphical illustration of these results. The left panel plots equilibrium participation as a function of network connectivity when the participation threshold is $k = 3$ for two environments: a homogeneous benchmark where all agents have the same propensity to form links (blue), and a highly heterogeneous core-periphery environment where one fifth of agents have a propensity of 5 and the remaining four fifths do not form links at all (red). The right panel shows the effect of varying the participation threshold k in the homogeneous benchmark.

From a technical standpoint, these results are derived by analyzing the giant k-core of the random network—the largest subgraph in which every agent has at least k connections to other members. The giant k-core characterizes the set of agents whose participation is endogenously sustained by mutual peer support.[3] To our knowledge, our comparative statics showing that, above the critical cutoff, the size of the giant k-core—and hence equilibrium participation—are strictly monotonic in both network connectivity and the participation threshold are novel.[4]

This clean characterization of equilibrium participation then opens up our second set of results.

[3]The k-core is known to characterize equilibrium behavior in threshold games on finite, deterministic networks (Gagnon and Goyal, 2017; Langtry et al., 2024). We work with its large-network analogue, the giant k-core.

[4]Although some of these comparative statics may have been known in the mathematical literature for the homogeneous benchmark (an Erdös-Rényi random graph), to our knowledge the results we present have not been recorded elsewhere. Moreover, our results apply to a much broader class of graphs than the homogeneous benchmark.

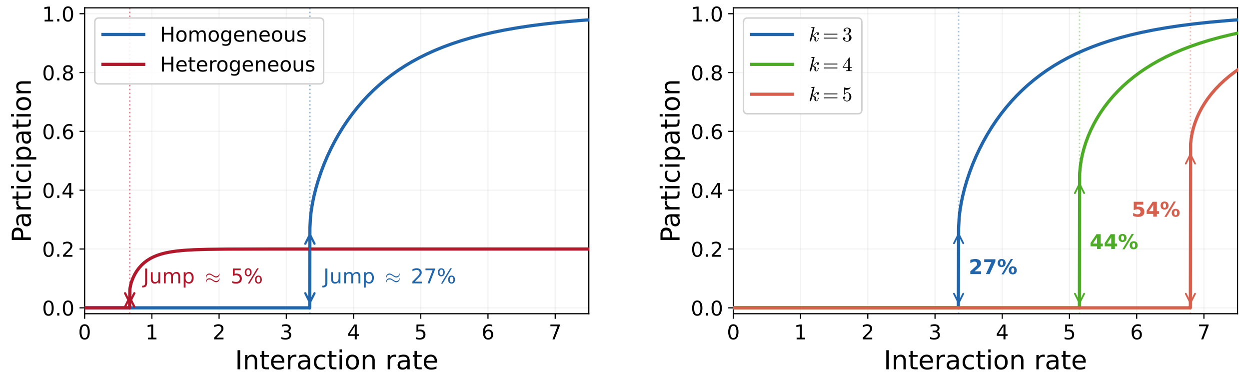

Figure 1: *Left:* Equilibrium participation as a function of network connectivity in two environments with the same mean degree: a homogeneous environment (blue) and a core–periphery environment (red). The arrows indicate the size of the jump—the minimum positive participation level. In the homogeneous environment, participation jumps by roughly 27%; in the core–periphery environment, by roughly 5%. *Right:* Equilibrium participation in the homogeneous benchmark for participation thresholds $k = 3, 4, 5$. Higher thresholds require more connectivity and produce larger jumps.

We characterize the principal's optimal choice of network connectivity and show how this depends on the heterogeneity of agents' propensities to form links. The presence—and more importantly, size—of the discontinuous jump in participation above the critical cutoff for network connectivity proves essential here. When the jump is large—as in homogeneous environments—the principal faces a significant "missing middle" in her strategy. To accept any participation is to accept a large amount of it: the discontinuous jump means that the principal either sets connectivity precisely at the critical cutoff (ensuring zero participation) or pushes well above it, where the benefits from connectivity are large enough to outweigh the costs of widespread participation. Intermediate levels of connectivity are never optimal. This is the stark, all-or-nothing logic that rationalizes extreme policy responses such as internet shutdowns.

But this extreme logic is not universal. When agents have highly heterogeneous propensities to form links, the jump above the cutoff can shrink, and with it, the missing middle. In a core–periphery environment, for instance, the missing middle vanishes as the "core" of highly connected agents becomes more concentrated. This is a key message of our paper: whether a principal facing a complex contagion is forced into extreme all-or-nothing policies depends on the heterogeneity of agents' link formation behavior. Homogeneous environments produce sharp tipping points and strong pressure toward extreme policies; heterogeneity tends to weaken that pressure by shrinking the jump above the cutoff.

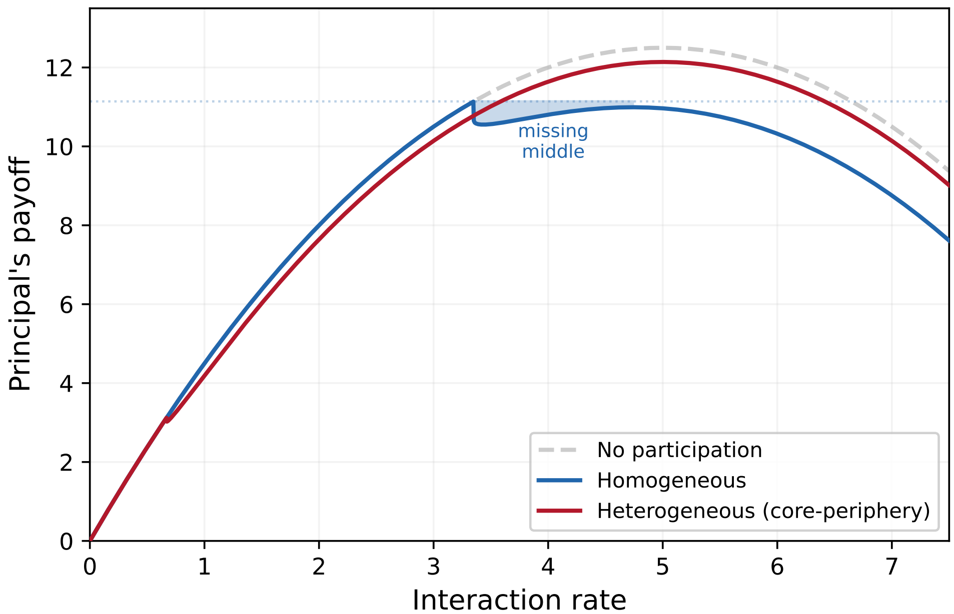

Figure 2: The principal's payoff in the homogeneous environment (blue) and core–periphery environment (red). The dashed gray curve is the no-participation benchmark. In the homogeneous case, the utility drop from the jump is large and generates a sizable missing middle interval (shaded). In the core–periphery case, the utility drop from the jump and implied missing middle are much smaller, so the optimum lies close to the benchmark.

Figure 2 plots the principal's payoffs in the same two environments, and hence shows our second set of results graphically. The blue and red lines correspond to the homogeneous and core–periphery settings in Figure 1, respectively. The dashed gray curve shows the no-participation benchmark. When agents are homogeneous (blue line), there is a large drop in the principal's payoff at the critical cutoff. This is precisely because the jump is large when interactions are homogeneous. In turn, this leads to a large "missing middle", as the principal must choose a significantly higher level of connectivity—entailing a significantly higher level of participation—to recoup her losses from allowing any participation at all.

The rest of the paper is organized as follows. Section 2 discusses related literature. Section 3 sets out the model. Section 4 characterizes agents' equilibrium behavior and the discontinuous emergence of participation. Section 5 solves the principal's control problem, derives the missing-middle result, and links its severity to heterogeneity. Section 6 discusses directions for future work and concludes.

2 Related Literature

Our paper builds on and contributes to four main literatures.

First, we develop a tractable model of complex contagions on a random network, and derive sharp comparative statics for equilibrium behavior, which have been elusive in fixed-network models. There is a large literature that studies models of complex contagion, with applications to a very wide range of economic behaviors.[5] These models primarily use a fixed network (see, e.g. Gagnon and Goyal, 2017; Reich, 2023; Chandrasekhar et al., 2025), and often focus on the diffusion of the action from some starting group of agents (Morris, 2000; Acemoglu et al., 2011).[6]

Second, by developing a model of collective behavior on a random network, we connect to a growing literature on simple contagion—where agents need only *one* neighbor to act—on random networks (Campbell, 2013; Akbarpour et al., 2020; Campbell et al., 2024a,b; Langtry, 2025). Sadler (2020) studies a general class of diffusion games on configuration model random graphs that are closely related to ours. A key distinction is that simple contagions produce continuous emergence of viral behavior: the giant component grows smoothly as connectivity increases (Centola and Macy, 2007; Centola, 2010). Complex contagions, by contrast, produce *discontinuous* transitions—participation jumps from zero to a positive fraction at the critical cutoff. Our model makes this contrast precise and tractable on the same class of random networks. Campbell et al. (2024b) also study a mixed Poisson random graph in a model of entry deterrence and pricing, focusing on the role of the giant component.

Third, we contribute to a literature on discontinuous phase transitions in networks—settings where small changes in network structure induce large, discontinuous changes in outcomes. Watts

[5]These include technology adoption (Reich, 2023), protest (Chwe, 2000), pricing (Zhang, 2025), financial contagion (Rogers and Veraart, 2013; Elliott et al., 2014), persuasion (Candogan, 2022), and choices of whether to participate in formal markets (Gagnon and Goyal, 2017). Granovetter (1978) and Schelling (1978) also suggest many other applications.

[6]Another approach has been to push the network into the background, by abstracting from it altogether (Granovetter, 1978; Schelling, 1978) or imposing a "mean field" assumption (Jackson and Yariv, 2006; López-Pintado, 2006, 2008), or to assume that agents only have partial information (Galeotti et al., 2010; Leister et al., 2022).

(2002) provides the seminal analysis of fractional-threshold contagion on random networks, showing that cascades can fail entirely or succeed on a large scale, with little in between. Elliott et al. (2022) show how a need for multiple types of input in supply networks can produce analogous discontinuous transitions, and provide comparative statics for the fragility of such networks. In our model, the discontinuity arises directly from the nature of threshold-based complex contagion itself: agents require a fixed number of participating neighbors, and this creates a sharp tipping point in aggregate behavior. We provide a number of new comparative statics and characterize how the distribution over agents' propensities to interact shapes the size and location of the discontinuity—including a comparison principle identifying when greater heterogeneity lowers the critical cutoff. Jackson and Storms (2025) also study complex contagion on random networks using stochastic block models, but focus on identifying which groups must take the same action and on optimal seeding, rather than on comparative statics for aggregate participation.[7]

Finally, we embed our analysis within a principal-agent framework to study optimal network design when the principal chooses aggregate connectivity, and show that homogeneity-versus-heterogeneity in the network structure moderates whether the optimal policy is extreme or interior. This connects us to the literature on network design and formation. This literature has increasingly focused on settings where a network forms and then agents play a game on the resulting network (Galeotti and Goyal, 2010; Sadler and Golub, 2021; Kinateder and Merlino, 2017, 2022). In part due to technical challenges, there is little work on network formation in a random-network setting. Elliott et al. (2022) and Langtry (2025) are exceptions, but consider endogenous formation by agents rather than control of overall connectivity by a principal. Galeotti et al. (2020) study the design of interventions in networks where a planner targets agents based on their degree. In a significantly different setting, Acemoglu et al. (2024) study platform design in a misinformation context and also find that the optimal policy can take a bang-bang form: the platform either allows wide sharing or restricts it sharply. Our analysis complements theirs by showing that this extreme logic is not universal—it depends on the heterogeneity in agents' propensities to form links.

[7]Their goal is to understand (a) which groups of agents must take the same action in any equilibrium, and (b) how to use this knowledge to best pick a set of initial adopters to maximize the spread of a new behavior. Our focus is instead on comparative statics for overall equilibrium behavior with respect to features of the network, as well as how a principal chooses overall network connectivity.

3 Model

We consider a sequence of games $\{\Gamma^{(n)}\}_{n\in\mathbb{N}}$ indexed by the number of agents n. We now describe the structure of a specific game $\Gamma^{(n)}$.

Agents, Actions & Timing. There is a single principal ("she"), P, and n agents ("he"), indexed $i \in N \equiv \{1, \dots, n\}$. In the first period ($t = 1$), the principal chooses an *interaction rate* $r \geq 0$, taking an *interaction profile* F as given.[8] Then Nature forms a large random network among agents (which we describe shortly). In the second period ($t = 2$), each agent simultaneously chooses whether or not to take a binary action, $a_i \in \{0, 1\}$. For clarity, we will say that an agent who takes action 1 *participates*, and an agent who takes action 0 *abstains*.

Random network formation and interaction profile. The principal controls only the interaction rate r (i.e. the network connectivity). But agents differ in their propensity to interact and form links. This heterogeneity is captured by a distribution F over nonnegative interaction weights ($W \geq 0$). An agent's expected degree is then $r \times W_i$.[9] The interaction weights W_i capture this variation. We impose the following condition on the interaction profile.

Assumption 1. *Weights* $(W_i)_{i=1}^n$ *are i.i.d. with* $W_i \sim F$*, satisfying* $\mathbb{E}[W] = 1$*, and* $\mathbb{E}[W^2] < \infty$.[10]

The assumption that $\mathbb{E}[W] = 1$ is simply a normalization. Conditional on W_i and r, each agent draws a number of stubs (their degree) $D_i \mid W_i \sim \text{Poisson}(rW_i)$ independently across i. Then conditional on the realized degree sequence $(D_i)_{i=1}^n$, Nature forms a (simple) random graph via the *configuration model*: stubs are paired uniformly at random, and we condition on the resulting graph being simple[11] (i.e. no self-loops or multiple edges). We let $G^{(n)}$ denote the adjacency matrix of a realization of this random graph: $G_{ij}^{(n)} \in \{0, 1\}$.

[8] We view F as a slow-moving, difficult-to-change feature of the environment, while r can be manipulated. Exploring what happens when the principal can also shape F is an interesting direction for future work; see Section 6.

[9] Such heterogeneity is well-documented empirically: degree distributions in real-world networks typically have a small fraction of agents maintaining many more connections than the median (Jackson, 2008).

[10] While mild, this can be violated by sufficiently heavy-tailed profiles. We discuss this in Online Appendix B.3.

[11] Under Assumption 1, the configuration model produces a simple graph with asymptotically positive probability, so this conditioning is without loss.

Preferences: principal. The principal dislikes *participation* among agents, but receives (diminishing) benefits from their interaction. Let $\alpha > 0$ be the net marginal benefit of interaction and let $\beta > 0$ be the per-capita cost she bears when agents participate. With interaction rate r and *participation rate* $\bar{a}$, her per-capita payoff is

$$\pi(r, F, \bar{a}) \; = \; \alpha\, r \; - \; \frac{1}{2} r^2 \; - \; \beta\, \bar{a}, \tag{1}$$

where $\bar{a} = \frac{1}{n}\sum_{i=1}^{n} a_i$. Note that the interaction profile F only affects the principal's objective through equilibrium participation.[12]

The two cost terms have distinct economic interpretations. *Infrastructure cost* ($\frac{1}{2}r^2$): the cost of providing interaction opportunities at rate r, regardless of how connections are distributed. This is an *extensive-margin* cost: it depends on the total volume of interaction, not its interaction profile. *Participation cost* ($\beta\bar{a}$): the cost the principal bears when agents participate.

Preferences: agents. Agents' actions are strategic complements: participating is more attractive to agent i when more of his neighbors participate. Specifically, i prefers to participate (choose $a_i = 1$) if and only if at least $k \geq 3$ of his neighbors also participate, where k is the *participation threshold*. Agents' preferences can be represented by a utility function $u_i(a_i, M_i)$ such that:

$$u_i(1, M_i) \geq u_i(0, M_i) \quad \Longleftrightarrow \quad M_i \geq k, \tag{2}$$

where $M_i = \sum_{j \neq i} G_{ij}^{(n)} a_j$ is the number of participating neighbors of i.

Information. For expositional simplicity, we assume that the realization of the network is common knowledge to agents. This means that actions a_i are functions $a_i = a_i(G^{(n)})$ of the realized network $G^{(n)}$ observed by agents. In the Online Appendix we provide a microfoundation which demonstrates that our model here can be viewed as the reduced-form of a model in which agents participate in an explicit diffusion process and only observe the actions of their neighbors.

[12] Online Appendix B.5 provides a congestion-based microfoundation for Equation (1).

Equilibrium. A strategy profile $\left(r^*, a_1^*, a_2^*, \ldots, a_n^*\right)$ is a *subgame perfect Nash equilibrium* of the game $\Gamma^{(n)}$ if, for any realization of the graph $G^{(n)}$, we have:

$$\text{for all } i \in N, \quad a_i^* = 1 \iff u_i(1, M_i^*) \geq u_i(0, M_i^*),$$

and the principal chooses r^* such that

$$r^* \in \arg\max_{r \geq 0} \ \mathbb{E}_{\mathbb{P}_{r,F}}\left[\pi(r, F, \bar{a}^*)\right],$$

where $\bar{a}^* = \frac{1}{n}\sum_{i=1}^n a_i^*$ is the equilibrium level of participation in the second stage and $\mathbb{P}_{r,F}$ is the distribution over networks generated by the principal's choice of r and the interaction profile F.

This definition requires that (i) agents best-respond to other agents given the realized network, and (ii) the principal maximizes her expected payoff, anticipating both the network structure and the resulting actions. In this paper we focus on the largest (i.e. highest participation) equilibrium, because it captures the worst feasible self-sustaining outcome for the principal and is also selected by natural local best-response dynamics. Online Appendix B.1 formalizes the latter interpretation.[13]

4 How Agents Behave

We begin with the second stage of the game, and characterize how agents' behavior depends on the interaction rate r and the interaction profile F. The second stage is a threshold game played on the realized network G. So before going further, it is important to discuss the role of the k-*core*—a notion of "cohesive" parts of a network—in characterizing equilibrium behavior.

4.1 Equilibrium and the k-core

In threshold games on networks, the largest equilibrium is determined by the network's "core structure". The k-core of a graph G is the largest induced subgraph H of G such that every agent

[13] Agents' actions are strategic complements, so the second stage of the game on any realized network $G^{(n)}$ has a *complete lattice* of Nash equilibria under the component-wise order on $\{0,1\}^n$ (Milgrom and Roberts, 1990; Topkis, 1998). In particular, there exists a *largest* equilibrium $a^{\max}(G^{(n)})$—the one with the most participating agents.

in H has degree at least k within H (Seidman, 1983). In our model the set of agents who participate in the largest equilibrium is precisely the k-core of the network G.

Remark 1. *In the largest equilibrium, agent i participates ($a_i = 1$) if and only if he belongs to the k-core of the realized network $G^{(n)}$.*[14]

We can see why this is the case by considering how agents might reason about a stable outcome. Any agent with fewer than k neighbors in the entire network knows he can never meet the participation threshold. So he will choose to abstain. Knowing this, other agents can revise what they expect their neighbors to do. An agent who initially had k neighbors might now expect fewer than k to participate, causing him to also abstain. This iterated removal of agents who lack sufficient support continues until only a stable group remains. The agents left are exactly those in the k-core; each has at least k connections to others who are also participating.

A key implication of Remark 1 is that characterizing equilibrium behavior in our model amounts to characterizing the fraction of agents belonging to the k-core of the random network induced by (r, F). The challenge in fixed-network models has been that it is very difficult to make strong predictions about how equilibrium behavior changes as a function of the primitives of the network. Using random networks allows us to obtain clean results in the limit of a large population. Therefore the remainder of our analysis focuses on the asymptotic properties of the model as the number of agents $n \to \infty$. As such, the equilibria we characterize are limits of subgame perfect equilibria of the finite games $\Gamma^{(n)}$.[15] We will often omit the limit in our statement of results, but any such results should be understood as applying with high probability as $n \to \infty$.

4.2 Equilibrium behavior

Our characterization of equilibrium behavior in the second stage builds on established results from random graph theory. The existence of a *giant k-core* (a k-core containing a constant fraction of agents with high probability as $n \to \infty$) is known to depend on a sharp cutoff condition.[16] Above

[14]The direct link between the largest equilibrium and the k-core has been established in prior studies of threshold games on fixed (finite) networks (Gagnon and Goyal, 2017; Langtry et al., 2024).

[15]All references to "equilibrium" from here on should be taken to mean "limits of equilibria in the sequence of finite games".

[16]For Erdős–Rényi graphs—which, in our model, corresponds to the homogeneous-interaction case—the canonical reference is Pittel et al. (1996); for configuration models with general degree distributions, see Janson and Luczak

this cutoff, the size of the giant k-core is strictly increasing in network connectivity. This provides the crucial link between the interaction rate r and equilibrium participation.

To fix ideas, let $a_k(r;F) \equiv \lim_{n\to\infty} \frac{1}{n}\sum_i a_i^*$ denote the limiting fraction of participating agents—the (level of) participation—given an interaction rate r and profile F. It is convenient to define $\psi_m(x) \equiv \mathbb{P}(\text{Poisson}(x) \geq m)$, and the functions

$$\Phi_k(x) \equiv \mathbb{E}[W\,\psi_{k-1}(Wx)]\,, \qquad A_k(x) \equiv \mathbb{E}[\psi_k(Wx)]\,. \tag{3}$$

When $W \equiv 1$ (the Erdős–Rényi benchmark), these reduce to $\Phi_k(x) = \psi_{k-1}(x)$ and $A_k(x) = \psi_k(x)$. Then define the critical cutoff

$$c_k(F) \equiv \inf_{x>0} \frac{x}{\Phi_k(x)}. \tag{4}$$

Under Assumption 1 and $k \geq 3$, $\Phi_k(x) = o(x)$ as $x \downarrow 0$ and $\Phi_k(x) \leq \mathbb{E}[W] = 1$ for all x, so $x/\Phi_k(x) \to \infty$ as $x \downarrow 0$ and as $x \to \infty$; hence the infimum in (4) is attained at some finite $x_k^\star(F) > 0$.[17] For expositional simplicity, we assume that the minimizer of $x/\Phi_k(x)$ is unique.

Assumption 2. *$g(x) \equiv x/\Phi_k(x)$ has a unique minimizer.*

This assumption ensures that the k-core emerges through a single discontinuous jump. It always holds when the interaction profile is sufficiently homogeneous (such as the Erdős–Rényi benchmark), and generally holds when heterogeneity is sufficiently "nice". In Appendix A.10 we provide a single-crossing condition on F that is sufficient for Assumption 2 to hold. We also discuss the more general analogues to our main results. Unless explicitly stated otherwise, all results throughout the rest of the paper operate under Assumptions 1 and 2.

Next, let $\rho = \rho_k(r;F)$ be the largest solution in $[0,1]$ to

$$\rho = \Phi_k(r\rho) = \mathbb{E}[W\,\psi_{k-1}(rW\rho)]\,. \tag{5}$$

Equivalently, letting $x \equiv r\rho$, this is the fixed-point equation $x = r\,\Phi_k(x)$. With this in place, we

(2007); Janson and Łuczak (2008). The critical cutoff condition is usually referred to as the critical *threshold* in this mathematics literature. But to avoid confusion, we reserve the term "threshold" for the participation threshold k.

[17] For completeness, Lemma 2 in the Appendix A.9 shows $\Phi_k(x) = o(x)$ as $x \downarrow 0$.

are now ready to characterize equilibrium behavior.

Theorem 1 (Participation)**.** *Fix $k \geq 3$ and an interaction profile F. Equilibrium participation:*

(i) is positive if and only if the interaction rate is sufficiently high ($r > c_k(F)$);

(ii) for every $r > c_k(F)$, is given by

$$a_k(r; F) = A_k(r\rho) = \mathbb{E}[\psi_k(rW\rho)]\,, \tag{6}$$

where ρ is as in (5).[18]

Theorem 1 establishes a sharp tipping point in equilibrium behavior. Unless the interaction rate r exceeds the critical cutoff $c_k(F)$, the network is too sparse to support a giant k-core, and equilibrium participation is zero. Above this cutoff, a strictly positive fraction of agents participate, characterized by part (ii). As we show in Section 4.4, this transition is not gradual—participation *jumps* discontinuously from zero to a strictly positive level at the cutoff.

The intuition for this lies in the nature of social reinforcement required for participation. For any group of agents to form a stable, self-sustaining equilibrium (i.e., to be a k-core), each member must have their participation supported by at least k friends who are also part of that same group.

This stands in sharp contrast to the *continuous* emergence of the giant component—a related feature of large random networks that has been shown to play a key role in driving behavior in models where the network spreads information (see, for example, Dasaratha (2023); Campbell et al. (2024b)).[19]

A sketch of the proof. Theorem 1 specializes Proposition 4.1 in Janson (2009) to the mixed-Poisson configuration model. In our setting the random degree sequence satisfies Condition 2.1 of Janson (2009) in probability, and Assumption 2 guarantees the non-local-maximum condition in

[18] Exactly at the cutoff $r = c_k(F)$, the fixed-point equation admits tangency solution, but the k-core law of large numbers (Proposition 4.1 of Janson and Luczak (2007)) does not apply there. Accordingly, we define $a_k(c_k(F); F) = 0$ by convention and study the right-limit $J_k(F) = \lim_{r \downarrow c_k(F)} a_k(r; F)$.

[19] The giant component is closely related to the k = 2-core. Taking the giant component and removing all agents with degree 1 yields the 2-core. The critical cutoff for the emergence of the giant 2-core is identical to the critical cutoff for the emergence of the giant component ($r = 1$).

Proposition 4.1(ii) for every $r > c_k(F)$. Writing D_p for the p-thinning of the limiting mixed-Poisson degree law, one has $h_1(p) = A_k(rp)$ and $h(p) = rp\,\Phi_k(rp)$, so the root equation $h(p) = \lambda p^2$ with $\lambda = r$ reduces to $x = r\Phi_k(x)$ after the change of variables $x = rp$. This yields (6). At the exact cutoff $r = c_k(F)$, Proposition 4.1 does not identify the finite-n k-core, so we make no probabilistic claim there and use the convention $a_k(c_k(F); F) = 0$.

4.3 Comparative statics: how agents respond to the environment

We now consider how equilibrium participation varies with the interaction rate r and the participation threshold k. The level of participation is increasing in the interaction rate and decreasing in the participation threshold.

Theorem 2 (Comparative statics: monotonicity)**.** *Fix an interaction profile F.*

(i) (Monotonicity in r.) *For every fixed $k \geq 3$, equilibrium participation, $a_k(r; F)$, is weakly increasing in r, and strictly increasing on $(c_k(F), \infty)$.*

(ii) (Monotonicity in k.) *For every fixed $r \geq 0$, equilibrium participation, $a_k(r; F)$, is weakly decreasing in $k \geq 3$, and strictly so whenever $a_k(r; F) > 0$.*

Intuitively, as the interaction rate increases, the degree distribution shifts to the right—agents have more neighbors on average. This increases the likelihood that the number of an agent's neighbors who participate will exceed k, reinforcing participation. Conversely, a higher threshold k means fewer agents will have sufficient participating neighbors, causing them to abstain. This initial abstention can then induce other agents to do the same, ultimately reducing participation.[20] Formally, these results are unlocked by our characterization in Theorem 1. To our knowledge, these comparative statics were not previously known.

A particular point of interest is the jump in participation due to the discontinuous emergence of the giant k-core. We now examine this jump in more detail. It is both of intrinsic interest—little is currently known about the jump beyond its existence—and will prove important for understanding the principal's optimal behavior.

[20] In contrast, when there is zero participation in equilibrium, changes in the network connectivity and/or the participation threshold have no impact. This follows immediately from Theorem 1(i).

4.4 The critical cutoff and the size of the jump

We now focus on the critical cutoff $c_k(F)$ and the *size* of the jump at the cutoff, as together these determine when participation first becomes possible and how much participation must arise when it does. We show that both are economically important, and characterize how they depend on the interaction profile F. But before stating any results, we need to define the size of the jump formally. It is the right-limit of participation at the critical cutoff:

$$J_k(F) \equiv \lim_{r \downarrow c_k(F)} a_k(r; F). \tag{7}$$

With this in place, we can now show that the jump is in fact a jump. That is, that above the critical cutoff, the participation rate is bounded above zero—no matter how close we get to the critical cutoff.

Proposition 1. *Fix any $k \geq 3$ and any interaction profile F.*

$$J_k(F) > 0.$$

As the interaction rate increases continuously, there is always a discontinuous jump from zero participation to some positive amount. Intuitively, this is because agents need *several* friends to take the action alongside them. When the interaction rate is low (below the critical cutoff), some set of agents may have at least k friends. And some subset of them may have at least k friends who in turn have at least k friends. But not enough of those friends will themselves have at least k friends who have k friends, and so on. There will be some "weak links" preventing mutually sustaining actions. So nobody participates. As the interaction rate rises just past the critical cutoff, those "weak links" will gain enough friends. This will be the "final piece of the puzzle" for a significant group of agents—leading to this whole group now being able to sustain participation. Hence there is a jump.

But is this group of agents large relative to the overall population? It turns out that when interactions are homogeneous (i.e. when $W \equiv 1$), this group makes up at least one quarter of the whole population. So the jump is *economically* significant. Further, this is *not* a knife-edge case reliant on there being no heterogeneity. The jump remains large when the interaction profile F has some heterogeneity.

Theorem 3. *Fix any $k \geq 3$. Let δ_1 denote the degenerate interaction profile with $W \equiv 1$.*
(i) (Homogeneous case.) *If $W \equiv 1$, then*[21]

$$J_k(\delta_1) > 0.26.$$

(ii) (Limited heterogeneity.) *If the variance of the interaction profile F is small enough, then both the critical cutoff and the size of the jump are close to those of the homogeneous case. Formally, for every $\eta > 0$ there exists $\varepsilon > 0$ such that for any F with $\sqrt{\mathrm{Var}(W)} \leq \varepsilon$, we have:*

$$|c_k(F) - c_k(\delta_1)| < \eta \quad \textit{and} \quad |J_k(F) - J_k(\delta_1)| < \eta.$$

Recent empirical work has identified a critical mass of approximately 25% as a "tipping point" required for a committed minority to overturn an established social norm (Centola et al., 2018). This highlights the necessary conditions for *initiating* a large-scale behavioral cascade. Theorem 3 provides a complementary theoretical perspective by characterizing the conditions for the resulting behavior to be *self-sustaining*: under a moderate participation threshold ($k = 3$), at least a quarter of the population must participate in the homogeneous benchmark. Our result suggests a possible explanation for this empirical tipping point: individuals may be willing to adopt a new norm only when they perceive the movement has enough momentum to become self-sustaining.

However, this large jump is not universal. It relies on interactions between agents being somewhat homogeneous. Sufficient heterogeneity can make the jump arbitrarily small.

Remark 2. *For every $\varepsilon > 0$ there exists a profile F with $J_k(F) < \varepsilon$.*

To help build an intuition of what such an interaction profile could look like, we show a simple example of pushing *all* of the interactions into an increasingly-smaller group of agents. For any

[21] This is for the $k = 3$ case. As k increases, so does the minimum participation: for example, when $k = 10$, $J_{10}(\delta_1) \approx 0.74$, and when $k = 100$ it is around 0.95. The derivation of the closed-form expression from the critical-point equation is given in the proof.

interaction profile F with weights W, we define a *diluted* interaction profile, $F^{(\nu)}$ for any $\nu \in (0, 1]$ with weights

$$W^{(\nu)} \equiv \frac{B}{\nu} W,$$

where $W \sim F$ and $B \sim \text{Bernoulli}(\nu)$ are independent. This dilution is a mean-preserving spread of F. In expectation, a fraction $1 - \nu$ of agents will not interact at all, and the remaining fraction ν will have their interaction rate scaled up by a factor $1/\nu$. In this special case, everything scales linearly.

Remark 3. *If F satisfies Assumption 2, then so does $F^{(\nu)}$, and for all $k \geq 3$*

$$c_k(F^{(\nu)}) = \nu \, c_k(F) \qquad \textit{and} \qquad J_k(F^{(\nu)}) = \nu \, J_k(F).$$

This shows how heterogeneity can shrink the size of the jump not by changing the within-core reinforcement technology, but by restricting the set of agents who can plausibly form a mutually reinforcing group. The network may be extremely dense *within* a small active set (mean degree r/ν), while the overall participation fraction scales with the size ν of that set.

The location of the jump. Under Bernoulli dilution, a more heterogeneous environment also has a lower cutoff: $c_k(F^{(\nu)}) = \nu \, c_k(F)$. As the active fraction shrinks, participation becomes possible at lower interaction rates—the network needs less overall connectivity to trigger collective action among the remaining active agents. This creates a tension for the principal. A more heterogeneous environment lowers the cutoff (participation is possible sooner) but also shrinks the jump (the discrete cost from crossing the cutoff is smaller). Whether this tension resolves in the principal's favor depends on parameters; Section 5 characterizes the tradeoff precisely through the missing-middle gap $d_k(F)$, which depends on both $c_k(F)$ and $J_k(F)$.

But does more heterogeneity always lower the cutoff? It turns out that in general the convex order alone is not enough to determine the effect of a change in the interaction profile. What matters is how a mean-preserving spread shifts mass around the margin of self-sustaining participation, which is determined by the threshold k. We now give a general result identifying when a mean-preserving spread lowers c_k.

A comparison principle for the critical cutoff. It is convenient to define, for $t \geq 0$,

$$C_F(t) \equiv \mathbb{E}[(W - t)_+] \,. \tag{8}$$

This measures the expected "excess interaction weight" that remains after stripping away the first t units of each agent's weight. In the finance literature, C_F is known as the *call-price curve* of the random variable W. For equal-mean profiles, ordering profiles by convex order is equivalent to ordering these curves pointwise: $F_1 \leq_{\text{cx}} F_2$ if and only if $C_{F_1}(t) \leq C_{F_2}(t)$ for all $t \geq 0$.

We can write Φ_k as a weighted integral against the call-price curve. So for every $x > 0$,

$$\Phi_{k,F}(x) = \frac{x^{k-1}}{(k-2)!} \int_0^\infty C_F(t)\, t^{k-2} e^{-xt} (k - xt)\, dt. \tag{9}$$

An analogous expression can be given for A_k. Together, Φ_k and A_k (from eq. (5) and eq. (6), respectively) pin down the size of the giant k-core. We derive both of these expressions formally in Section A.6.

Importantly, the term $t^{k-2} e^{-xt}(k - xt)$ in eq. (9) changes sign exactly once—at $t = k/x$. This sign change is the key to understanding which forms of heterogeneity move the cutoff, and which do not. Recall that $x = r\rho$, so an agent with weight w has wx participating neighbors in expectation. So an agent with weight $w = k/x$ expects exactly k participating neighbors—right at the participation threshold. Agents below this weight fall just short of the participation threshold; agents above it are safely inside the giant k-core.

With this representation in hand, we can answer the question posed above: when does a mean-preserving spread in F lower the critical cutoff? For some initial interaction profile F_1, there is an associated value of x, denoted $x^\star_{1,k}$, that is the level of x at the critical cutoff $c_k(F_1)$. In turn this implies a threshold weight $w^\star_{1,k} := k/x^\star_{1,k}$ such that agents with higher weights are able to sustain participation when connectivity is at the cutoff, while those just below are marginal—close to, but below the participation threshold. Intuitively, agents with weights at $w^\star_{1,k}$ are those who just make it into the giant k-core at the point where it emerges. Since the kernel in (9) changes sign at the threshold weight and decays exponentially away from it, the balance of the call-price curve on either side of $w^\star_{1,k}$ determines which direction the cutoff moves. Theorem 4 makes this precise: whether a mean-preserving spread raises or lowers the critical cutoff depends on how it shifts weight in the call price curve around $w^\star_{1,k}$. Increases in "excess interaction weight" among agents who fall just short of the participation threshold enter the integral with positive sign and tend to lower the cutoff, while increases among safely participating agents enter with negative sign and tend to raise it.

Theorem 4 (Critical cutoff comparison)**.** *Fix $k \geq 3$. Let $F_1 \leq_{\text{cx}} F_2$, so that $D(t) \equiv C_{F_2}(t) - C_{F_1}(t) \geq 0$ for all t. Suppose F_1 satisfies Assumption 2, and let $x^\star_1 \equiv x^\star_k(F_1)$ denote the unique minimizer of $g_{F_1}(x) = x/\Phi_{k,F_1}(x)$. If*

$$\int_0^\infty D(t)\, t^{k-2} e^{-x^\star_1 t}(k - x^\star_1 t)\, dt \geq 0, \tag{10}$$

then $c_k(F_2) \leq c_k(F_1)$. In particular, a sufficient condition for (10) *is $D(t) = 0$ for all $t > k/x^\star_1$.*

Extra dispersion among agents just below the margin of self-sustaining participation—those whose expected participating neighbors fall short of k—raises Φ_k and lowers the cutoff. This is because spreading mass among underconnected agents has increasing returns, creating some who become newly capable of sustaining participation. Formally, this is driven by the fact that $w \mapsto w\,\psi_{k-1}(xw)$ is convex in this region. In contrast, extra dispersion among agents already safely above the threshold has little effect: these agents participate regardless, so redistributing weight among them does not create new participants.

Theorem 4 clarifies why the critical cutoff cannot be ranked by the convex order alone. The integral condition (10) asks whether, evaluated at the minimizer $x^\star_1$ under F_1, the positive contribution to Φ_k from agents below the threshold weight outweighs the negative contribution from those above it. Our sharper sufficient condition—that $D(t) = 0$ for all $t > k/x^\star_1$—is the special case where the negative side vanishes entirely: all the extra dispersion under F_2 is concentrated among agents who fall short of the participation threshold. Note that neither condition requires solving for the equilibrium under F_2. [22]

Theorem 4 also makes it immediately clear why Bernoulli dilution behaves so cleanly. The identity $C_{F^{(\nu)}}(t) = C_F(\nu t)$ yields $\Phi_{k,F^{(\nu)}}(x) = \Phi_{k,F}(x/\nu)$ and hence $c_k(F^{(\nu)}) = \nu\, c_k(F)$ immediately. Bernoulli dilution is special because everything rescales together—the general result covers c_k; dilution is the tractable case where J_k and a_k come along for free.

A numerical illustration. Theorem 4 shows that the impact of a mean-preserving spread can, in principle, go either way—it depends on the full distributional object $D(t)$. However, for the gamma and lognormal distributions—parametric families commonly used in applied network analysis—the pattern is unambiguous. Figure 3 plots the critical cutoff and the size of the jump as a function of the dispersion for both families. For these two distributions, the fact that c_k is decreasing in dispersion can be shown from Theorem 4, by evaluating Equation (10). The figures show that for both families, the jump size J_k also decreases with dispersion. This is suggestive of a broader pattern: for regular parameterized families, greater heterogeneity tends to lower the cutoff *and* shrink the jump.

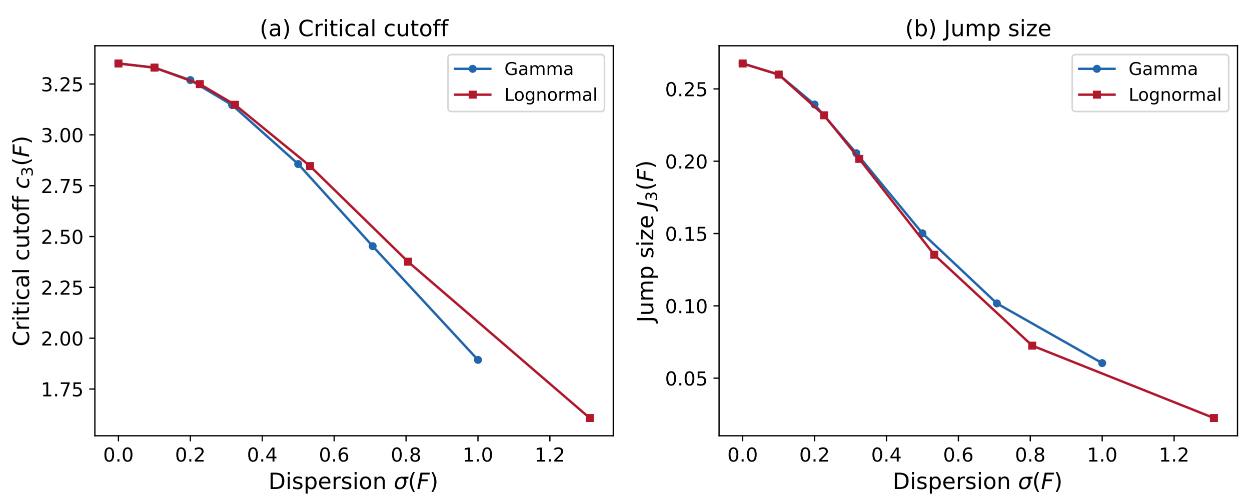

Figure 3: Critical cutoff $c_3(F)$ (left) and jump $J_3(F)$ (right) against dispersion $\sigma(F)$ for the gamma and lognormal families. Both families show the same qualitative pattern: increasing heterogeneity lowers the cutoff and shrinks the jump.

[22] The analogous forces act on J_k through a parallel signed-kernel formula, but a ranking of J_k requires additional conditions because $J_k(F) = A_k(x_k^\star(F))$ also depends on the movement of the minimizer (see Section A.6). This is consistent with the counterexample in Section B.7, where J_k increases despite the more dispersed profile having greater convex order.

With this understanding of second-stage behavior and how it depends on key features of the environment in place, we now turn to the principal's problem.

5 How the Principal Designs the Network

The principal's problem is to choose the interaction rate r to maximize her payoff (1), taking as given how agents will behave in the second stage—which is in turn a function of the participation threshold, k, and the interaction profile, F. Substituting the equilibrium participation $\bar{a} = a_k(r; F)$ from Section 4, she maximizes

$$\pi(r) = \alpha r - \tfrac{1}{2}\, r^2 - \beta\, a_k(r; F). \tag{11}$$

The headline result is that if the interaction profile is homogeneous, the principal will never choose an interaction rate that is close to but above the critical cutoff, $c_k(F)$. This is precisely due to the large jump in the level of participation (i.e. the discontinuous emergence of the giant k-core) at the critical cutoff. Setting the interaction rate exactly at the cutoff delivers no participation at all. But going even a very small amount above the cutoff suddenly yields significant participation. This creates significant costs for the principal, but only small incremental benefits. So she is better off by either "returning" to the critical cutoff, or "pushing through" to a much higher interaction rate—generating large enough direct benefits from interactions to offset the participation costs.

We call this push to extreme choices the "missing middle". But note that the missing middle is only economically meaningful when the jump is large. In particular, highly heterogeneous interaction profiles can generate small jumps at the critical cutoff, and hence small losses from choosing an interaction rate just above that cutoff. The rest of the section formalizes this intuition.

5.1 Choosing the optimal interaction rate

Preliminary observations. To start, let $r^{\text{naive}} = \alpha$. This is what the principal would choose absent any concerns about participation (i.e. $\beta = 0$). And because the level of participation, $a_k(r; F)$, is weakly increasing in the interaction rate, the principal never chooses an interaction rate above this level (i.e. $r > r^{\text{naive}}$). Next, note that the optimum exists by upper semicontinuity and compactness,[23] and that the optimum is generically unique. This is because the value function $V(\beta) \equiv \sup_{r \geq 0} \pi(r; F)$ is convex in β (as the pointwise supremum of affine functions), hence differentiable except at a countable set. At any β where V is differentiable, the global maximizer is unique (we prove this in Appendix A.14). So if there are two or more equilibria then a small perturbation of β restores uniqueness.

Finally, notice that the problem is trivial if $\alpha \leq c_k(F)$. In this case, the principal's unconstrained optimum, $r^{\text{naive}} = \alpha$, lies *below* the critical cutoff and triggers no participation. So throughout this section we focus on the case where $\alpha > c_k(F)$. This is where the principal faces a genuine tension between the benefits of interaction and the costs of the participation it induces.

The missing middle. The discontinuity in equilibrium participation creates a distinctive feature of the principal's problem. Because participation jumps discretely from zero to at least $J_k(F)$ when r crosses the critical cutoff $c_k(F)$, the principal will never choose an interaction rate "just above" the cutoff. If she moves even slightly beyond $c_k(F)$, she incurs a participation cost of at least $\beta J_k(F)$—a discrete penalty—while the direct benefit $\alpha r - \frac{1}{2}r^2$ improves only marginally relative to its value at $c_k(F)$. For each $\beta \geq 0$, define

$$\pi_\beta(r; F) \equiv \alpha r - \frac{1}{2}r^2 - \beta\, a_k(r; F), \qquad f(\beta; \alpha, F) \equiv \sup_{r > c_k(F)} \pi_\beta(r; F).$$

Then

$$\bar{\beta}(\alpha; F) = \sup\big\{\beta > 0 : f(\beta; \alpha, F) \geq \pi_\beta\big(c_k(F); F\big)\big\}\,. \tag{12}$$

With this notation in place, we can now formalize the "missing middle" intuition.

Proposition 2. *Suppose $\alpha > c_k(F)$, and let $\bar{\beta}(\alpha; F)$ be defined in* (12)*.*

(i) (High participation costs.) *if $\beta > \bar{\beta}$, the optimal interaction rate is exactly at the critical cutoff (and the level of participation is zero): $r^* = c_k(F)$.*

[23] Since $a_k(r; F) \in [0, 1]$ for all r, the term $-\frac{1}{2}r^2$ eventually dominates, so there is a compact interval $[0, b]$ on which all optima lie. On this interval, $\pi(r; F) = \alpha r - \frac{1}{2}r^2 - \beta a_k(r; F)$ is upper semicontinuous: participation jumps (if any) are upward in a_k, so the induced jumps in π are downward. An upper semicontinuous function on a compact set attains its maximum.

(ii) (Low participation costs.) *if $\beta < \bar{\beta}$, the optimal interaction rate is at least a distance $d_k(F) > 0$ above the critical cutoff: $r^* \geq c_k(F) + d_k(F)$, where*

$$d_k(F) = \frac{\beta\, J_k(F)}{\alpha - c_k(F)} > 0. \tag{13}$$

The result splits optimal behavior into two regimes, with a sharp boundary at $\bar{\beta}$. When participation is sufficiently costly for the principal, she will choose the highest interaction rate that yields no participation. And when participation is not too costly, she will choose an interaction rate significantly above the critical cutoff, accepting a high level of participation in the process.[24]

The term $d_k(F)$ provides a lower bound on how far the principal must move up from the critical cutoff to recover her losses. It of course depends on the parameters α and β. But its dependence on the rest of the environment is entirely through the critical cutoff $c_k(F)$ and the jump $J_k(F)$—the two key summary statistics from Section 4. The comparison theorem (Theorem 4) clarifies which forms of heterogeneity move these objects: heterogeneity weakens the missing middle when it creates more agents near the margin of self-sustaining participation, rather than concentrating interaction among agents already safely inside the core. A lower c_k, combined with a smaller J_k, reduces $d_k(F)$ and narrows the missing middle.

Heterogeneity in interactions therefore determines the size of the missing middle. In homogeneous environments, $J_k(F)$ is large, so the missing middle is large. In contrast, more heterogeneous environments have a small jump, which weaken the discrete tradeoff near the cutoff and makes an interaction rate closer to the critical cutoff less costly.[25]

The missing middle also creates a sensitivity in the principal's optimal choice of r. Although any equilibrium with positive participation has *a lot* of participation, small changes in the incentives (α or β) may induce the principal to "flip" from an equilibrium with substantial participation to one with zero (or vice versa). To explore this sensitivity, we now describe how the principal's

[24]The cutoff $\bar{\beta}(\alpha; F)$ is a knife-edge in the middle, where there is both an optimal choice at the critical cutoff and one significantly above it. The principal is of course indifferent between these choices.

[25]Of course, the global optimum may lie well above $c_k(F_\nu)$. This will depend on the parameters (α, β). Proposition 2 characterizes a region above the cutoff $c_k(F)$ that can *never* be optimal. Additionally, note that our analysis has focused on the case where there is only one jump in the size of the giant k-core (this is the case where the function g has a unique minimizer), and hence in the level of participation. If participation has additional jumps at higher r, an analogous missing middle can arise there; we leave that extension for future work.

equilibrium choice of r depends on the participation cost β.

5.2 Comparative statics and tipping points

The next result gives both a robust (set-valued) comparative static and a stronger derivative statement at regular interior optima. Let $R(\beta) \equiv \arg\max_{r\geq 0} \pi_\beta(r;F)$ denote the set of optimal interaction rates, and let $\bar{a}^* \equiv a_k(r^*;F)$ denote the equilibrium participation at the optimal policy.

Proposition 3 (Comparative statics and tipping points)**.** *Fix $\alpha > c_k(F)$ and let $\bar{\beta} = \bar{\beta}(\alpha;F)$.*

(i) (Monotonicity.) *If $0 \leq \beta_1 < \beta_2$ and $r_i \in R(\beta_i)$, then $r_2 \leq r_1$.*

(ii) (Strict monotonicity.) *Moreover, if $0 \leq \beta_1 < \beta_2 < \bar{\beta}$ and each $R(\beta_i)$ is a singleton $\{r_i\}$, then $r_2 < r_1$ and $a_k(r_2;F) < a_k(r_1;F)$.*

(iii) (Local derivative.) *At any $\beta < \bar{\beta}$ where the maximizer is unique, interior, and nondegenerate,*

$$\frac{\mathrm{d}r^*}{\mathrm{d}\beta} < 0 \quad \textit{and} \quad \frac{\mathrm{d}\bar{a}^*}{\mathrm{d}\beta} < 0.$$

Proposition 3 shows the difference between the two regimes. When the principal tolerates the behavior in equilibrium ($\beta < \overline{\beta}$), adjustments are smooth: as β increases, the principal reduces the interaction rate. This, in turn, leads to a reduction in participation. In contrast, when the optimal strategy is to completely suppress participation ($\beta > \overline{\beta}$), the principal's choice of r^* is "sticky". Small changes in her incentives are insufficient to move the interaction rate away from the critical cutoff; she continues to choose the highest possible interaction rate that guarantees zero participation.

The key takeaway from this result is that when β is close to $\overline{\beta}$, a small change in the cost of participation can move her from choosing an equilibrium with positive participation to one with zero participation (or the reverse). This highlights a tipping-point phenomenon in the principal's decision-making: her strategy, and more importantly the resulting participation, can shift dramatically in response to relatively mild changes in her incentives.

6 Concluding Remarks

We finish by briefly commenting on some of our modeling choices and how they relate to directions for future research.

Strategic complementarities. The model's core mechanism is a participation threshold driven by strategic complementarities. We interpret this broadly as a process of social reinforcement, where an agent's incentive to act is fundamentally linked to the actions of his direct network connections. The parameter k represents the critical mass of local support required to make an action worth taking. In our setting, this threshold is an *absolute number* of neighbors, rather than a fraction of all neighbors. This is technically useful, as it allows us to build on existing work in random graph theory. And it corresponds to settings where agents need a certain "critical mass" of local support for an action—such as participating in a protest, where safety in numbers requires enough fellow participants regardless of total network size, or adopting a new norm that requires a minimum number of compatible peers. Settings where agents instead care about the *share* of neighbors taking an action—such as language choice, where speakers switch when most of their contacts have adopted the new language—would benefit from a different modeling approach.

The random network. Formally, our random network is a mixed-Poisson model.[26] In the configuration model, agents form links uniformly at random. This is a deliberate assumption to isolate the interplay between complex contagion and network heterogeneity in expected degrees (i.e. agents' propensities to form links). However, our model abstracts from several important features of real-world networks. These include homophily—the tendency for agents to form links with similar others—and multiplexity, the idea that different types of links can play different roles. Additionally, agents in our framework do not make endogenous decisions about forming or severing links; the network is determined by the principal's choice of interaction rate and the realization of the random graph. In reality, agents may have some control over both how many connections to form and with whom to form them.

[26]This nests the canonical Erdős–Rényi network as a special case when $W \equiv 1$. When $W \equiv 1$, agents' degrees are i.i.d. Poisson(r), which is asymptotically equivalent to $G(n, r/n)$.

Relaxing these assumptions provides a clear agenda for future work. Incorporating homophily into the configuration model—for instance through a stochastic block model with heterogeneous degrees—would allow the analysis to capture group-level variation in both connectivity and behavior. Endogenous link formation, where agents can respond to the principal's choice by rewiring their connections, represents a more significant challenge but is important for applications where agents strategically manage their exposure to contagion.

The principal's levers. In our setting, the principal's policy levers are summarized by a single parameter, the interaction rate r. In practice, this may reflect a bundle of decisions—such as the frequency of meetings, the layout of shared spaces, or the length of breaks—that collectively shape how much agents interact. But we are considering a setting where the principal lacks more fine-grained tools required to target specific groups of agents, or induce agents to interact with specific people. Hence she treats the interaction profile F as fixed. A natural extension is a joint dynamic problem in which policy choices can also shape F; we leave that extension for future work.

Another natural tool for the principal is targeted interventions—such as degree-dependent subsidies or selective seeding of agents to shift their behavior. In our framework, random seeding decomposes cleanly via Poisson thinning, but the interaction between seeding and endogenous connectivity choice is richer: when the principal can respond to seeding by adjusting r, the net effect on participation may depend on the interaction profile in ways that our analysis of the missing middle suggests but does not fully resolve. We view this as another promising direction for future work.

An advantage for empirical work. Despite the abstractions, our framework's tractability offers a practical advantage: the key tradeoffs are driven by the interaction rate r and the interaction profile F, rather than by the full network adjacency matrix. This makes the model's predictions testable with aggregate data on connectivity patterns, and should serve as a useful benchmark for richer models that incorporate more detailed network features.

Moreover, the heterogeneity in our interaction profile F is directly tied to degree moments. The variance in realized degrees depends only on the interaction rate r and the variance in the weights

W_i. Formally: $\text{Var}(D_i) = \mathbb{E}[\text{Var}(D_i \mid W_i)] + \text{Var}(\mathbb{E}[D_i \mid W_i]) = r + r^2\text{Var}(W_i)$. So the variance in our (unobservable) interaction profile F is identified by the mean and variance of the realized degree distribution.

Wrapping up. This paper sets out a tractable framework for analyzing complex contagions on heterogeneous networks. By using a random networks approach we are able to provide clean comparative statics for how aspects of network structure affect equilibrium behavior. Notably, we highlight that a sudden jump in behavior around a critical level of network connectivity arises directly from the threshold-based nature of complex contagion, and how heterogeneity in the network structure influences the size of this jump.

These results then unlock an analysis of a principal's optimal design problem, where she can choose network connectivity but not the finer structure of the network. When the network structure is homogeneous, the principal is forced into extreme all-or-nothing policies. She either suppresses connectivity to prevent any participation or sets connectivity well above the critical cutoff, allowing significant participation. There is no middle ground. This is driven by the jump. If the principal accepts any participation, she must accept significant participation, so she is "in for a penny, in for a pound."

But in highly heterogeneous environments—such as those with a core–periphery structure—more moderate policies can arise. This is precisely because the jump in participation at the cutoff is smaller, which weakens the force driving the extreme policies. Together, these insights provide a new link between network heterogeneity and the qualitative character of optimal policy, and help explain why principals in some environments resort to extreme measures (such as internet shutdowns) while in others, the pressure toward extremes is weaker and moderate policies can emerge.

References

Access Now (2022). Weapons of control, shields of impunity: Internet shutdowns in 2022. Technical report, Access Now.

Acemoglu, D., Ozdaglar, A., and Siderius, J. (2024). A model of online misinformation. *Review of Economic Studies*, 91(6):3117–3150.

Acemoglu, D., Ozdaglar, A., and Yildiz, E. (2011). Diffusion of innovations in social networks. In *2011 50th IEEE conference on decision and control and European control conference*, pages 2329–2334. IEEE.

Airoldi, E. M. and Christakis, N. A. (2024). Induction of social contagion for diverse outcomes in structured experiments in isolated villages. *Science*, 384(6695):eadi5147.

Akbarpour, M., Malladi, S., and Saberi, A. (2020). Just a few seeds more: value of network information for diffusion. *Available at SSRN 3062830*.

Bandiera, O. and Rasul, I. (2006). Social networks and technology adoption in northern mozambique. *The Economic Journal (London)*, 116(514):869–902.

Beaman, L., BenYishay, A., Magruder, J., and Mobarak, A. M. (2021). Can network theory-based targeting increase technology adoption? *American Economic Review*, 111(6):1918–43.

Björkegren, D. (2019). The adoption of network goods: Evidence from the spread of mobile phones in rwanda. *The Review of Economic Studies*, 86(3):1033–1060.

Bursztyn, L., Cantoni, D., Yang, D., Yuchtman, N., and Zhang, Y. (2021). Persistent political engagement: Social interactions and the dynamics of protest movements. *American Economic Review: Insights*, 3(2):233–250.

Campbell, A. (2013). Word-of-mouth communication and percolation in social networks. *American Economic Review*, 103:2466–2498.

Campbell, A., Thornton, D., and Zenou, Y. (2024a). Strategic diffusion: Public goods vs. public bads. Technical Report DP19443, CEPR.

Campbell, A., Ushchev, P., and Zenou, Y. (2024b). The network origins of entry. *Journal of Political Economy*, 132(11):3867–3916.

Candogan, O. (2022). Persuasion in networks: Public signals and cores. *Operations Research*, 70(4):2264–2298.

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, 329(5996):1194–1197.

Centola, D., Becker, J., Brackbill, D., and Baronchelli, A. (2018). Experimental evidence for tipping points in social convention. *Science*, 360(6393):1116–1119.

Centola, D. and Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–734.

Chandrasekhar, A. G., Chaudhary, V., Golub, B., and Jackson, M. O. (2025). Multiplexing in networks and diffusion. Working paper.

Christakis, N. A. and Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4):370–379.

Christakis, N. A. and Fowler, J. H. (2008). The collective dynamics of smoking in a large social network. *New England journal of medicine*, 358(21):2249–2258.

Chwe, M. S.-Y. (2000). Communication and coordination in social networks. *The Review of Economic Studies*, 67(1):1–16.

Dasaratha, K. (2023). Innovation and strategic network formation. *The Review of Economic Studies*, 90(1):229–260.

Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–3153.

Elliott, M., Golub, B., and Leduc, M. V. (2022). Supply network formation and fragility. *American Economic Review*, 112(8):2701–2747.

Enikolopov, R., Makarin, A., and Petrova, M. (2020). Social media and protest participation: Evidence from russia. *Econometrica*, 88(4):1479–1514.

Gagnon, J. and Goyal, S. (2017). Networks, markets, and inequality. *American Economic Review*, 107(1):1–30.

Galeotti, A., Golub, B., and Goyal, S. (2020). Targeting interventions in networks. *Econometrica*, 88(6):2445–2471.

Galeotti, A. and Goyal, S. (2010). The law of the few. *American Economic Review*, 100(4):1468–1492.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., and Yariv, L. (2010). Network games. *The review of economic studies*, 77(1):218–244.

González, F. (2020). Collective action in networks: Evidence from the chilean student movement. *Journal of Public Economics*, 188:104220.

Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443.

Jackson, M. O. (2008). *Social and economic networks*. Princeton University Press.

Jackson, M. O. and Storms, E. C. (2025). Behavioral communities and the atomic structure of networks. *American Economic Journal: Microeconomics*. Forthcoming.

Jackson, M. O. and Yariv, L. (2006). Diffusion on social networks. *Economie publique/Public economics*, 1(16).

Janson, S. (2009). On percolation in random graphs with given vertex degrees. *Electronic Communications in Probability*, 14:86–118.

Janson, S. and Łuczak, M. (2008). Asymptotic normality of the k-core in random graphs. *Annals of Applied Probability*, 18(3):1085–1137.

Janson, S. and Luczak, M. J. (2007). A simple solution to the k-core problem. *Random Structures & Algorithms*, 30(1-2):50–62.

Kinateder, M. and Merlino, L. P. (2017). Public goods in endogenous networks. *American Economic Journal: Microeconomics*, 9(3):187–212.

Kinateder, M. and Merlino, L. P. (2022). Local public goods with weighted link formation. *Games and economic behavior*, 132:316–327.

King, G., Pan, J., and Roberts, M. E. (2013). How censorship in china allows government criticism but silences collective expression. *American Political Science Review*, 107(2):326–343.

Langtry, A. (2025). More connection, less community: network formation and local public goods provision. *arXiv preprint arXiv:2504.06872*.

Langtry, A., Taylor, S., and Zhang, Y. (2024). Network threshold games. Technical report, arXiv. org.

Larson, J. M., Nagler, J., Ronen, J., and Tucker, J. A. (2019). Social networks and protest participation: Evidence from 130 million twitter users. *American Journal of Political Science*, 63(3):690–705.

Leister, C. M., Zenou, Y., and Zhou, J. (2022). Social connectedness and local contagion. *The Review of Economic Studies*, 89(1):372–410.

López-Pintado, D. (2006). Contagion and coordination in random networks. *International Journal of Game Theory*, 34(3):371–381.

López-Pintado, D. (2008). Diffusion in complex social networks. *Games and Economic Behavior*, 62(2):573–590.

Madestam, A., Shoag, D., Veuger, S., and Yanagizawa-Drott, D. (2013). Do political protests matter? evidence from the tea party movement. *Quarterly Journal of Economics*, 128(4):1633–1685.

Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, pages 1255–1277.

Morris, S. (2000). Contagion. *The Review of Economic Studies*, 67(1):57–78.

Munshi, K. and Myaux, J. (2006). Social norms and the fertility transition. *Journal of development Economics*, 80(1):1–38.

Pittel, B., Spencer, J., and Wormald, N. (1996). Sudden Emergence of a Giant$k$-Core in a Random Graph. *Journal of Combinatorial Theory, Series B*, 67(1):111–151.

Reich, B. (2023). The architecture of social networks and the diffusion of innovations. *Working Paper*.

Rogers, L. C. and Veraart, L. A. (2013). Failure and rescue in an interbank network. *Management Science*, 59(4):882–898.

Sadler, E. (2020). Diffusion games. *American Economic Review*, 110(1):225–270.

Sadler, E. and Golub, B. (2021). Games on endogenous networks. *arXiv preprint arXiv:2102.01587*.

Schelling, T. C. (1978). *Micromotives and macrobehavior*. WW Norton & Company.

Seidman, S. B. (1983). Network structure and minimum degree. *Social networks*, 5(3):269–287.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.

V-Dem Institute (2024). Internet shutdowns shutting down democracy. Technical report, V-Dem Institute Policy Brief No. 40.

Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771.

Zhang, Y. (2025). Price competition with network spillovers: Cohesiveness, entry and interoperability. *unpublished manuscript*.

A Appendix A: Proofs

A.1 Proof of Theorem 1 (Participation)

Proof. Let D_i denote the degree of agent i. Conditional on W_i,

$$D_i \mid W_i \sim \text{Poisson}(rW_i)$$

independently across i. Unconditionally, the D_i are i.i.d. with mixed-Poisson law

$$p_j(r; F) \equiv \mathbb{P}(D_i = j) = \mathbb{E}\left[e^{-rW} \frac{(rW)^j}{j!}\right], \qquad j \geq 0,$$

and mean $\mathbb{E}[D_i] = r$. Let $n_j^{(n)} \equiv \#\{i \leq n : D_i = j\}$. By the law of large numbers,

$$\frac{n_j^{(n)}}{n} \to p_j(r; F) \qquad \text{for every } j \geq 0, \qquad \frac{1}{n}\sum_{i=1}^{n} D_i \to r$$

in probability. Thus the random degree sequence satisfies Condition 2.1 of Janson (2009) in probability. Following Section 2 of Janson (2009), a Skorohod coupling lets us work on a probability space on which that condition holds almost surely. Since Proposition 4.1 of Janson (2009) is stated conditional on the realized degree sequence, its conclusions then apply to the multigraph configuration model. Under Assumption 1, the configuration model is simple with probability bounded away from zero, so convergence in probability transfers to the simple-graph-conditioned model we use in Section 3.

Let D denote a random variable with the limiting degree law $(p_j(r; F))_{j \geq 0}$, and let D_p denote its p-thinning (i.e. $D_p \sim \text{Bin}(D, p)$). Conditional on W,

$$D_p \mid W \sim \text{Poisson}(rWp).$$

Therefore

$$h_1(p) \equiv \mathbb{P}(D_p \geq k) = \mathbb{E}[\psi_k(rWp)] = A_k(rp),$$

and

$$\begin{aligned} h(p) &\equiv \mathbb{E}[D_p \, \mathbf{1}_{\{D_p \geq k\}}] \\ &= \mathbb{E}[rWp \, \psi_{k-1}(rWp)] \\ &= rp \, \Phi_k(rp). \end{aligned}$$

The root equation in Janson's (2009) Proposition 4.1, in our notation, is $h(p) = rp^2$, which is equivalent to

$$rp\,\Phi_k(rp) = rp^2.$$

For $p = 0$ this is trivial. For $p > 0$, setting $x = rp$ gives

$$x = r\Phi_k(x).$$

Thus positive roots correspond exactly to positive solutions of the fixed-point equation $x = r\Phi_k(x)$, and for such a solution $h_1(p) = A_k(x)$ with $p = x/r$.

If $r < c_k(F)$, then $x > r\Phi_k(x)$ for every $x > 0$, so there are no positive roots. Proposition 4.1(i) of Janson (2009) then gives

$$\frac{v(\mathrm{Core}_k)}{n} \xrightarrow{p} 0.$$

Now suppose $r > c_k(F)$. Under Assumption 2, the function $g(x) = x/\Phi_k(x)$ has a unique minimizer $x_k^\star(F)$ and is strictly increasing on $[x_k^\star(F), \infty)$ (see Section A.2). Let $x(r)$ denote the largest positive solution to $x = r\Phi_k(x)$. Then $x(r) > x_k^\star(F)$. Choose $\delta \in (0, x(r) - x_k^\star(F))$. For every $x \in (x(r) - \delta, x(r))$, strict increase of g gives $g(x) < g(x(r)) = r$, i.e. $x < r\Phi_k(x)$. Writing $p = x/r$ and $\rho = x(r)/r$, we obtain

$$h(p) - rp^2 = rp\big(\Phi_k(rp) - p\big) > 0 \qquad \text{for all } p \in (\rho - \delta/r, \rho).$$

By Janson (2009, Remark 4.3), this is exactly the non-local-maximum condition required in Proposition 4.1(ii). Therefore

$$\frac{v(\mathrm{Core}_k)}{n} \xrightarrow{p} h_1(\rho) = A_k(r\rho) = A_k(x(r)) > 0.$$

By Remark 1, in every realized network the fraction of participating agents equals $v(\mathrm{Core}_k)/n$. Hence the above convergence is exactly the law of large numbers for equilibrium participation, giving (6) for $r > c_k(F)$.

At the exact critical cut-off $r = c_k(F)$, Proposition 4.1 of Janson (2009) does not apply. In the Erdős–Rényi case, Pittel et al. (1996) show that the k-core at criticality has size $o(n)$, so $a_k(c_k(\delta_1); \delta_1) = 0$. More generally, we adopt the convention $a_k(c_k(F); F) = 0$ and define the jump $J_k(F)$ as the right-limit from above. □

A.2 Proof of Proposition 1 (Jump size)

Proof. We first show that $g(x) = x/\Phi_k(x)$ attains its minimum on $(0, \infty)$. The map $x \mapsto \Phi_k(x)$ is continuous on $(0, \infty)$ by dominated convergence, and it is strictly positive there because $\mathbb{P}(W > 0) > 0$ and $\psi_{k-1}(y) > 0$ for every $y > 0$. Hence g is continuous on $(0, \infty)$. By Lemma 2, $\Phi_k(x)/x \to 0$ as $x \downarrow 0$ (for $k \geq 3$ and $\mathbb{E}[W^2] < \infty$), so $g(x) \to \infty$ as $x \downarrow 0$. As $x \to \infty$, $\Phi_k(x) \leq \mathbb{E}[W] = 1$, hence

$g(x) \geq x \to \infty$. Therefore g attains its minimum on a compact subinterval of $(0, \infty)$, and $\arg\min g$ is nonempty.

Now assume the maintained single-jump regularity: g has a unique minimizer $x^\star$ and is strictly increasing on $[x^\star, \infty)$. For each $r > c_k(F) = g(x^\star)$, the equation $r = g(x)$ has a unique solution on $[x^\star, \infty)$; this solution is the largest fixed point, which we denote by $x_k(r)$. Since g is continuous and strictly increasing on $[x^\star, \infty)$, its inverse on that interval is continuous, so

$$x_k(r) \downarrow x^\star \qquad \text{as } r \downarrow c_k(F).$$

Because $a_k(r; F) = A_k(x_k(r))$ and A_k is continuous,

$$J_k(F) = \lim_{r \downarrow c_k(F)} a_k(r; F) = A_k(x^\star).$$

Moreover, this value is strictly positive: since $x^\star > 0$, $\mathbb{P}(W > 0) > 0$, and $\psi_k(y) > 0$ for every $y > 0$, the expectation

$$A_k(x^\star) = \mathbb{E}[\psi_k(W x^\star)]$$

is strictly positive. Hence $J_k(F) > 0$. □

A.3 Proof of Theorem 2 (Comparative statics)

Proof. (i) Monotonicity in the interaction rate. Fix the participation threshold k, and let $x_k(r)$ denote the largest solution to $x = r\Phi_k(x)$. By Theorem 1,

$$a_k(r; F) = A_k(x_k(r)).$$

Write $T_r(x) \equiv r\Phi_k(x)$. For each r, the map T_r is increasing in x and maps $[0, r]$ into itself because $\Phi_k(x) \leq 1$. Starting the iteration from a sufficiently large initial condition—for instance $x_0 = r$—produces a decreasing sequence

$$x_{m+1} = T_r(x_m)$$

that converges to the largest fixed point $x_k(r)$. If $r_2 > r_1$, then $T_{r_2}(x) \geq T_{r_1}(x)$ for every $x \geq 0$. Starting both iterations from the same large initial condition, say $x_0 = r_2$, induction gives $x_m^{(r_2)} \geq x_m^{(r_1)}$ for every m, and passing to the limit yields

$$x_k(r_2) \geq x_k(r_1).$$

Since A_k is increasing, $a_k(r; F)$ is weakly increasing in the interaction rate r.

Now let $r_2 > r_1 > c_k(F)$. Then $x_k(r_1) > 0$, and $\Phi_k(x_k(r_1)) > 0$ because $\mathbb{P}(W > 0) > 0$ and

$\psi_{k-1}(y) > 0$ for every $y > 0$. Hence

$$T_{r_2}(x_k(r_1)) = r_2\Phi_k(x_k(r_1)) > r_1\Phi_k(x_k(r_1)) = x_k(r_1).$$

Monotone iteration for T_{r_2} starting from $x_k(r_1)$ therefore converges to a fixed point strictly above $x_k(r_1)$, so $x_k(r_2) > x_k(r_1)$. Because $\mathbb{P}(W > 0) > 0$ and $\psi_k(y)$ is strictly increasing for every $y > 0$, the map A_k is strictly increasing on $(0, \infty)$. It follows that

$$a_k(r_2; F) > a_k(r_1; F).$$

(ii) Monotonicity in the participation threshold. Fix the interaction rate r. For every $x > 0$,

$$\Phi_{k+1}(x) = \mathbb{E}[W\psi_k(Wx)] < \mathbb{E}[W\psi_{k-1}(Wx)] = \Phi_k(x),$$

because $\mathbb{P}(W > 0) > 0$ and $\psi_k(y) < \psi_{k-1}(y)$ for every $y > 0$. Therefore the fixed-point maps satisfy

$$T_{k+1,r}(x) \equiv r\Phi_{k+1}(x) \leq r\Phi_k(x) \equiv T_{k,r}(x) \qquad \text{for all } x \geq 0,$$

and the same monotone-iteration argument gives

$$x_{k+1}(r) \leq x_k(r).$$

Using again the representation from Theorem 1,

$$\begin{aligned} a_{k+1}(r; F) &= \mathbb{E}[\psi_{k+1}(Wx_{k+1}(r))] \\ &\leq \mathbb{E}[\psi_k(Wx_{k+1}(r))] \\ &\leq \mathbb{E}[\psi_k(Wx_k(r))] \\ &= a_k(r; F). \end{aligned}$$

This proves weak monotonicity in k.

If $a_{k+1}(r; F) > 0$, then necessarily $x_{k+1}(r) > 0$ because $A_{k+1}(0) = 0$. Since $\mathbb{P}(W > 0) > 0$ and $\psi_{k+1}(y) < \psi_k(y)$ for every $y > 0$, the first inequality above is then strict, and hence

$$a_{k+1}(r; F) < a_k(r; F).$$

This is exactly the strict inequality stated in Theorem 2(ii). □

A.4 Proof of Theorem 3 (Jump size and near-homogeneous stability)

Proof. Fix $k \geq 3$ and let $x^\star \equiv x_k^{\star,ER}$ be the unique minimizer of $g^{ER}(x) = x/\psi_{k-1}(x)$, with minimum value $c^{ER} \equiv c_k^{ER}$.

Part (i): the homogeneous benchmark. When $W \equiv 1$, $\Phi_k(x) = \psi_{k-1}(x)$, $A_k(x) = \psi_k(x)$, and the minimizer $x^\star$ satisfies the first-order condition $\psi_{k-1}(x^\star) = x^\star \, \psi'_{k-1}(x^\star)$. Using the standard Poisson identities $\psi'_m(x) = p_{m-1}(x)$, where $p_m(x) \equiv \mathbb{P}(\text{Poisson}(x) = m)$, and $\psi_{k-1} = \psi_k + p_{k-1}$, we obtain

$$
\begin{aligned}
J_k^{ER} &= \psi_k(x^\star) = \psi_{k-1}(x^\star) - p_{k-1}(x^\star) = x^\star \, p_{k-2}(x^\star) - p_{k-1}(x^\star) \\
&= (k-1)\, p_{k-1}(x^\star) - p_{k-1}(x^\star) = (k-2)\, p_{k-1}(x^\star),
\end{aligned}
$$

where the penultimate step uses the Poisson recursion $x\, p_{k-2}(x) = (k-1)\, p_{k-1}(x)$. For $k = 3$, $x^\star$ solves $e^x = 1 + x + x^2$, and $J_3^{ER} = p_2(x^\star) \approx 0.2676 > 0.26$.

Part (ii): limited heterogeneity.

Step 1: compactness of the minimization problem near ER. Because $g^{ER}(x) \to \infty$ as $x \downarrow 0$ and as $x \to \infty$, there exist $0 < a < x^\star < b < \infty$ such that

$$
\inf_{x \in (0,a] \cup [b,\infty)} g^{ER}(x) \;\geq\; c^{ER} + 3\eta. \tag{14}
$$

We will show that for F sufficiently close to δ_1, the minimizer of $g^F(x) \equiv x/\Phi_k(x)$ must lie in $[a, b]$ as well.

Fix b as above. Note that for every F, $\Phi_k(x) \leq \mathbb{E}[W] = 1$ (since $\psi_{k-1} \leq 1$), so $g^F(x) = x/\Phi_k(x) \geq x$ for all $x \geq 0$. Thus, for the chosen b, $g^F(x) \geq x \geq b$ for all $x \geq b$. In particular, if $b \geq c^{ER} + 3\eta$ (we may enlarge b to ensure this), then

$$
\inf_{x \geq b} g^F(x) \;\geq\; c^{ER} + 3\eta \qquad \text{for all } F. \tag{15}
$$

Next, fix $a > 0$ as in (14) and (if needed) shrink it further so that $a \leq 1/2$. By Lemma 4, for any interaction profile F with dispersion $\sigma(F) = \sigma$ and any $x \in (0, a]$,

$$
g^F(x) = \frac{x}{\Phi_k(x)} \;\geq\; \frac{1}{4x + 6\sigma^2} \;\geq\; \frac{1}{4a + 6\sigma^2}.
$$

Choose $\varepsilon_1 \in (0, 1]$ small enough (and, if needed, shrink a once more) so that

$$
\frac{1}{4a + 6\varepsilon_1^2} \;\geq\; c^{ER} + 3\eta.
$$

Then for all F with $\sigma(F) \le \varepsilon_1$,

$$\inf_{x\in(0,a]} g^F(x) \;\ge\; c^{ER} + 3\eta. \tag{16}$$

Combining (15) and (16), for all F with $\sigma(F) \le \varepsilon_1$: if the global minimum of g^F is below $c^{ER} + 3\eta$, then every minimizer must lie in the compact interval $[a, b]$. We now show that the minimum is indeed achieved below this level.

Step 2: uniform closeness of g^F to g^{ER} on $[a, b]$. On $[a, b]$, ψ_{k-1} is bounded below by

$$m \equiv \min_{x\in[a,b]} \psi_{k-1}(x) > 0.$$

By Lemma 5, there exists $\varepsilon_2 > 0$ such that $\sigma(F) \le \varepsilon_2$ implies

$$\sup_{x\in[a,b]} |\Phi_k(x) - \psi_{k-1}(x)| \le \frac{m}{2},$$

hence $\Phi_k(x) \ge m/2$ for all $x \in [a, b]$. Therefore, for such F and all $x \in [a, b]$,

$$|g^F(x) - g^{ER}(x)| = \left|\frac{x}{\Phi_k(x)} - \frac{x}{\psi_{k-1}(x)}\right| = x \cdot \frac{|\Phi_k(x) - \psi_{k-1}(x)|}{\Phi_k(x)\psi_{k-1}(x)} \le b \cdot \frac{\sup_{[a,b]} |\Phi_k - \psi_{k-1}|}{(m/2)\cdot m}.$$

Hence $\sup_{x\in[a,b]} |g^F(x) - g^{ER}(x)| \to 0$ as $\sigma(F) \to 0$.

Step 3: continuity of the cutoff map. Let $c(F) \equiv c_k(F) = \min_{x>0} g^F(x)$; by Step 1 the minimum is achieved on $[a, b]$ when $\sigma(F)$ is small. Let $x^F \in \arg\min_{x\in[a,b]} g^F(x)$. Then

$$c(F) = g^F(x^F) \le g^F(x^\star) \le g^{ER}(x^\star) + \sup_{[a,b]} |g^F - g^{ER}| = c^{ER} + o(1),$$

and also

$$c^{ER} = g^{ER}(x^\star) \le g^{ER}(x^F) \le g^F(x^F) + \sup_{[a,b]} |g^F - g^{ER}| = c(F) + o(1).$$

Thus $|c(F) - c^{ER}| \to 0$ as $\sigma(F) \to 0$, proving the cutoff continuity claim.

Step 4: uniqueness near the homogeneous benchmark. Define

$$H^F(x) \equiv x\,\Phi_k'(x) - \Phi_k(x), \qquad H^{ER}(x) \equiv x\,\psi_{k-1}'(x) - \psi_{k-1}(x),$$

so that by Lemma 3(ii),

$$(g^F)'(x) = -\frac{H^F(x)}{\Phi_k(x)^2}, \qquad (g^{ER})'(x) = -\frac{H^{ER}(x)}{\psi_{k-1}(x)^2}.$$

Step 4a: the Erdős–Rényi sign pattern comes from an explicit H^{ER} calculation (not from uniqueness alone). Since $\psi'_{k-1}(x) = p_{k-2}(x)$, we have

$$H^{ER}(x) = x\, p_{k-2}(x) - \psi_{k-1}(x).$$

Differentiate and use $p'_m(x) = p_{m-1}(x) - p_m(x)$:

$$H^{ER\prime}(x) = p_{k-2}(x) + x\, p'_{k-2}(x) - \psi'_{k-1}(x) = x\big(p_{k-3}(x) - p_{k-2}(x)\big).$$

Moreover $p_{k-3}(x)/p_{k-2}(x) = (k-2)/x$, so $p_{k-3}(x) - p_{k-2}(x)$ is strictly positive for $x < k-2$ and strictly negative for $x > k-2$. Hence H^{ER} is strictly increasing on $(0, k-2]$ and strictly decreasing on $[k-2, \infty)$. Since $H^{ER}(x) \to 0$ as $x \downarrow 0$ and H^{ER} is strictly increasing on $(0, k-2]$, we have $H^{ER}(k-2) > 0$. Also $H^{ER}(x) \to -1$ as $x \to \infty$ because $p_{k-2}(x) \to 0$ and $\psi_{k-1}(x) \to 1$. Therefore H^{ER} has a unique zero on $(0, \infty)$, necessarily at some $x^\star > k-2$. In particular, $H^{ER}(x) > 0$ for $x \in (0, x^\star)$ and $H^{ER}(x) < 0$ for $x > x^\star$, so g^{ER} is strictly decreasing on $(0, x^\star)$ and strictly increasing on $(x^\star, \infty)$.

Step 4b: transfer the single-crossing of H^{ER} to H^F on the compact $[a, b]$. Fix $\delta > 0$ so small that $x^\star - \delta > k-2$ and $x^\star + \delta < b$. By continuity and the strict sign pattern above,

$$m_L \equiv \min_{x \in [a, x^\star - \delta]} H^{ER}(x) > 0, \qquad m_R \equiv \min_{x \in [x^\star + \delta, b]} \big(- H^{ER}(x)\big) > 0.$$

Next note the deterministic bound, for all $x \in [a, b]$,

$$|H^F(x) - H^{ER}(x)| \le x\, |\Phi'_k(x) - \psi'_{k-1}(x)| + |\Phi_k(x) - \psi_{k-1}(x)| \le b \sup_{[0,b]} |\Phi'_k - \psi'_{k-1}| + \sup_{[0,b]} |\Phi_k - \psi_{k-1}|.$$

By Lemma 5 (and $\mathbb{E}|W-1| \le \sigma(F)$, $\mathbb{E}|W^2 - 1| \le \sigma(F)\sqrt{4 + \sigma(F)^2}$), the right-hand side tends to 0 as $\sigma(F) \to 0$. Hence, for $\sigma(F)$ sufficiently small,

$$H^F(x) \ge \frac{m_L}{2} \ \text{ for all } x \in [a, x^\star - \delta], \qquad H^F(x) \le -\frac{m_R}{2} \ \text{ for all } x \in [x^\star + \delta, b]. \tag{17}$$

In particular $H^F(x^\star - \delta) > 0$ and $H^F(x^\star + \delta) < 0$, so by continuity H^F has at least one zero in $(x^\star - \delta, x^\star + \delta)$.

Step 4c: uniqueness via monotonicity of H^F on $[x^\star - \delta, b]$. For $x > 0$, differentiating $\Phi'_k(x) = \mathbb{E}[W^2 p_{k-2}(Wx)]$ under the expectation (as in Lemma 3(i)) gives

$$\Phi''_k(x) = \mathbb{E}\big[W^3\,(p_{k-3}(Wx) - p_{k-2}(Wx))\big]\,.$$

Let $q(z) \equiv p_{k-3}(z) - p_{k-2}(z) = \psi''_{k-1}(z)$. Because $x^\star - \delta > k-2$, q is continuous and strictly negative on $[x^\star-\delta,b]$, so

$$m'' \equiv \min_{x\in[x^\star-\delta,b]} \big(-q(x)\big) > 0.$$

For each $x\in[x^\star-\delta,b]$, define $f_x(w)\equiv w^3 q(wx)$. Using the explicit Poisson form $p_m(t)=e^{-t}t^m/m!$, every term in $f_x'(w)$ is a polynomial in w times e^{-wx}; since $x\ge x^\star-\delta>0$, this implies

$$L \;\equiv\; \sup_{x\in[x^\star-\delta,b]}\ \sup_{w\ge 0}\ |f_x'(w)| \;<\; \infty,$$

because $\sup_{w\ge0} w^m e^{-w(x^\star-\delta)}<\infty$ for each fixed m. Therefore, for all $x\in[x^\star-\delta,b]$,

$$|\Phi_k''(x)-q(x)| = |\mathbb{E}[f_x(W)-f_x(1)]| \le L\,\mathbb{E}|W-1| \le L\,\sigma(F).$$

Choosing $\sigma(F)$ small enough that $L\sigma(F)\le m''/2$ yields $\Phi_k''(x)\le -m''/2<0$ on $[x^\star-\delta,b]$. Consequently, by Lemma 3(iii),

$$(H^F)'(x) = x\,\Phi_k''(x) \;<\; 0 \qquad \text{for all } x\in[x^\star-\delta,b],$$

so H^F is strictly decreasing on $[x^\star-\delta,b]$. Combined with (17), this implies that H^F has a *unique* zero in $(x^\star-\delta,x^\star+\delta)$; denote it by $x_k^\star(F)$. Since $(g^F)'(x)=-H^F(x)/\Phi_k(x)^2$, we have $(g^F)'(x)<0$ for $x<x_k^\star(F)$ and $(g^F)'(x)>0$ for $x\in(x_k^\star(F),b]$, hence g^F is strictly decreasing on $[a,x_k^\star(F)]$ and strictly increasing on $[x_k^\star(F),b]$. By Step 1, every global minimizer of g^F lies in $[a,b]$, so $x_k^\star(F)$ is the unique global minimizer of g^F. Moreover, since $\delta>0$ was arbitrary, this shows $x_k^\star(F)\to x^\star$ as $\sigma(F)\to 0$.

Step 4d: strict increase on $[x_k^\star(F),\infty)$. The preceding argument gives strict increase on $[x_k^\star(F),b]$. To extend this to $[b,\infty)$, enlarge b once more (if needed) so that $\psi_{k-1}(b/2)\ge 1/2$ and

$$\frac{k^k e^{-k}}{(k-2)!\,b} \;\le\; \frac{1}{8}.$$

Then for any F satisfying Assumption 1 and any $x\ge b$,

$$\Phi_k(x) \ge \mathbb{E}\big[W\,\psi_{k-1}(Wx)\mathbf{1}_{\{W\ge 1/2\}}\big] \ge \psi_{k-1}(b/2)\,\mathbb{E}\big[W\mathbf{1}_{\{W\ge1/2\}}\big] \ge \frac12\cdot\frac12 = \frac14,$$

where we used $Wx\ge x/2\ge b/2$ on $\{W\ge 1/2\}$ and $\mathbb{E}[W\mathbf{1}_{\{W<1/2\}}]\le 1/2$. Also, using $p_{k-2}(t)=e^{-t}t^{k-2}/(k-2)!$ and $\sup_{w\ge0} w^k e^{-wx}=(k/x)^k e^{-k}$, let

$$C_k \equiv \frac{k^k e^{-k}}{(k-2)!}.$$

Then

$$x\Phi_k'(x) = x\,\mathbb{E}[W^2 p_{k-2}(Wx)] = \frac{x^{k-1}}{(k-2)!}\,\mathbb{E}\Big[W^k e^{-Wx}\Big] \le \frac{x^{k-1}}{(k-2)!}\cdot\frac{k^k e^{-k}}{x^k} = \frac{C_k}{x} \le \frac{C_k}{b} \le \frac{1}{8}.$$

Thus $\Phi_k(x) - x\Phi_k'(x) \ge 1/8 > 0$ for all $x \ge b$, so $(g^F)'(x) > 0$ on $[b,\infty)$. Therefore g^F is strictly increasing on $[x_k^\star(F),\infty)$, i.e. the hypotheses of Proposition 1 (single-jump regularity) hold for F sufficiently close to δ_1.

Step 5: continuity of the jump. Under single-jump regularity, $J_k(F) = A_k(x_k^\star(F))$ and $J_k^{ER} = \psi_k(x^\star)$. Because $x_k^\star(F) \to x^\star$ by Step 4, and because (by Lemma 5) $A_k(x) \to \psi_k(x)$ uniformly on $[a,b]$ as $\sigma(F) \to 0$, we obtain

$$J_k(F) = A_k(x_k^\star(F)) = \psi_k(x_k^\star(F)) + o(1) = \psi_k(x^\star) + o(1) = J_k^{ER} + o(1).$$

Thus $|J_k(F) - J_k^{ER}| \to 0$ as $\sigma(F) \to 0$, proving (ii). □

A.5 Proof of Remark 3 (Bernoulli dilution scaling)

Proof. Let $W \sim F$ and $B \sim \text{Bernoulli}(\nu)$ be independent and define $W^{(\nu)} \equiv (B/\nu)W$.

(i) Conditional on W,

$$\mathbb{E}[W^{(\nu)} \mid W] = \mathbb{E}[B/\nu]\cdot W = W.$$

Hence for any convex φ with $\mathbb{E}|\varphi(W^{(\nu)})| < \infty$, Jensen gives

$$\varphi(W) = \varphi(\mathbb{E}[W^{(\nu)} \mid W]) \le \mathbb{E}[\varphi(W^{(\nu)}) \mid W],$$

and taking expectations yields $\mathbb{E}[\varphi(W)] \le \mathbb{E}[\varphi(W^{(\nu)})]$, i.e. $F \le_{\text{cx}} F^{(\nu)}$.

(ii) Using $B \in \{0,1\}$,

$$\Phi_k^{(\nu)}(x) = \mathbb{E}\left[\frac{B}{\nu}W\,\psi_{k-1}\left(\frac{B}{\nu}Wx\right)\right] = \nu\cdot\mathbb{E}\left[\frac{1}{\nu}W\,\psi_{k-1}\left(\frac{1}{\nu}Wx\right)\right] = \Phi_k(x/\nu).$$

Also $\psi_k(0) = 0$ for $k \ge 1$, so

$$A_k^{(\nu)}(x) = \mathbb{E}\left[\psi_k\left(\frac{B}{\nu}Wx\right)\right] = \nu\,\mathbb{E}\left[\psi_k\left(\frac{1}{\nu}Wx\right)\right] = \nu\,A_k(x/\nu).$$

Thus $g^{(\nu)}(x) = x/\Phi_k^{(\nu)}(x) = \nu g(x/\nu)$ and taking infima yields $c_k(F^{(\nu)}) = \nu c_k(F)$.

(iii) If F satisfies Assumption 2, then Part (ii) gives

$$g^{(\nu)}(x) = \nu g(x/\nu).$$

Hence $g^{(\nu)}$ has unique minimizer $\nu x_k^\star(F)$. Moreover, for any $x_2 > x_1 \geq \nu x_k^\star(F)$,

$$g^{(\nu)}(x_2) - g^{(\nu)}(x_1) = \nu\Big(g(x_2/\nu) - g(x_1/\nu)\Big) > 0,$$

so $g^{(\nu)}$ is strictly increasing on $[\nu x_k^\star(F), \infty)$. Thus $F^{(\nu)}$ also satisfies Assumption 2, and

$$J_k(F^{(\nu)}) = A_k^{(\nu)}(\nu x_k^\star(F)) = \nu A_k(x_k^\star(F)) = \nu J_k(F).$$

The two-class (core–periphery) family F_ν is the special case $F = \delta_1$: setting $W \equiv 1$ above gives $c_k(F_\nu) = \nu\, c_k^{ER}$ and $J_k(F_\nu) = \nu\, J_k^{ER}$. □

A.6 Call-price representation and proof of Theorem 4

We first state a general identity that underpins the comparison theorem.

Lemma 1 (Call-price identity)**.** *Let W be any nonnegative integrable random variable with distribution F, and define $C_F(t) \equiv \mathbb{E}[(W-t)_+]$ for $t \geq 0$. Let $h\colon [0,\infty) \to \mathbb{R}$ be C^2 with $h(0) = h'(0) = 0$ and $\int_0^\infty |h''(t)|\, dt < \infty$. Then*

$$\mathbb{E}[h(W)] = \int_0^\infty C_F(t)\, h''(t)\, dt.$$

Proof. For any fixed $w \geq 0$,

$$\int_0^\infty (w-t)_+ h''(t)\, dt = \int_0^w (w-t) h''(t)\, dt.$$

Integrating by parts,

$$\int_0^w (w-t) h''(t)\, dt = \Big[(w-t)h'(t)\Big]_0^w + \int_0^w h'(t)\, dt = -wh'(0) + h(w) - h(0).$$

Since $h(0) = h'(0) = 0$, this equals $h(w)$. Therefore

$$h(w) = \int_0^\infty (w-t)_+ h''(t)\, dt.$$

Now take expectations. Because $W \geq 0$ and $\mathbb{E}[W] < \infty$,

$$0 \leq C_F(t) = \mathbb{E}[(W-t)_+] \leq \mathbb{E}[W] < \infty.$$

Hence

$$\int_0^\infty C_F(t)\, |h''(t)|\, dt \leq \mathbb{E}[W] \int_0^\infty |h''(t)|\, dt < \infty,$$

so Fubini's theorem applies:

$$\mathbb{E}[h(W)] = \int_0^\infty \mathbb{E}[(W-t)_+]\, h''(t)\, dt = \int_0^\infty C_F(t)\, h''(t)\, dt. \qquad \square$$

Proof of Theorem 4. Write

$$p_m(z) \equiv e^{-z}\frac{z^m}{m!}, \qquad \psi_m(z) \equiv \mathbb{P}(\text{Poisson}(z) \geq m).$$

Recall the standard identities $\psi_m'(z) = p_{m-1}(z)$ and $p_m'(z) = p_{m-1}(z) - p_m(z)$.

Define, for fixed $x > 0$,

$$h_{k,x}(w) \equiv w\, \psi_{k-1}(xw), \qquad \ell_{k,x}(w) \equiv \psi_k(xw).$$

Because $k \geq 3$, we have $h_{k,x}(0) = h_{k,x}'(0) = 0$ and $\ell_{k,x}(0) = \ell_{k,x}'(0) = 0$. Also both second derivatives are polynomial-exponential functions and hence absolutely integrable on $[0,\infty)$.

For $h_{k,x}$, a direct calculation gives

$$h_{k,x}'(w) = \psi_{k-1}(xw) + xw\, p_{k-2}(xw),$$

hence, using $xw\, p_{k-3}(xw) = (k-2)p_{k-2}(xw)$,

$$h_{k,x}''(w) = \frac{x^{k-1}}{(k-2)!} w^{k-2} e^{-xw}(k - xw).$$

For $\ell_{k,x}$,

$$\ell_{k,x}'(w) = x\, p_{k-1}(xw),$$

so

$$\ell_{k,x}''(w) = \frac{x^k}{(k-1)!} w^{k-2} e^{-xw}(k - 1 - xw).$$

Applying Lemma 1 to $h_{k,x}$ and $\ell_{k,x}$ gives

$$\Phi_{k,F}(x) = \mathbb{E}[h_{k,x}(W)] = \frac{x^{k-1}}{(k-2)!}\int_0^\infty C_F(t)\, t^{k-2} e^{-xt}(k - xt)\, dt,$$
$$A_{k,F}(x) = \mathbb{E}[\ell_{k,x}(W)] = \frac{x^k}{(k-1)!}\int_0^\infty C_F(t)\, t^{k-2} e^{-xt}(k - 1 - xt)\, dt.$$

The one-sign-change claim is immediate: $t^{k-2}e^{-xt} > 0$ for all $t > 0$, while $k - xt$ crosses zero exactly once at $t = k/x$ and $k - 1 - xt$ crosses zero exactly once at $t = (k-1)/x$.

For part (i), subtract the formulas for F_1 and F_2. Since $D(t) = C_{F_2}(t) - C_{F_1}(t)$, we obtain (**??**).

For part (ii), assume (10). Then by part (i),

$$\Phi_{k,F_2}(x_1^\star) \geq \Phi_{k,F_1}(x_1^\star),$$

hence

$$g_{F_2}(x_1^\star) \leq g_{F_1}(x_1^\star) = c_k(F_1).$$

Therefore

$$c_k(F_2) = \inf_{x>0} g_{F_2}(x) \leq g_{F_2}(x_1^\star) \leq c_k(F_1).$$

If moreover $D(t) = 0$ for all $t > k/x_1^\star$, then every nonzero contribution to the integral in (10) comes from the region where $k - x_1^\star t \geq 0$. Since also $D(t) \geq 0$, the integral is nonnegative. □

Proof of the call-price identity for Bernoulli dilution. Let $W^{(\nu)} = (B/\nu)W$ where $B \sim \text{Bernoulli}(\nu)$ is independent of W. Then

$$C_{F^{(\nu)}}(t) = \mathbb{E}\left[\left(\frac{B}{\nu}W - t\right)_+\right] = \nu\,\mathbb{E}\left[\left(\frac{W}{\nu} - t\right)_+\right] = \mathbb{E}[(W - \nu t)_+] = C_F(\nu t).$$

Substituting into the call-price representation for Φ_k and using the change of variables $u = \nu t$ gives $\Phi_{k,F^{(\nu)}}(x) = \Phi_{k,F}(x/\nu)$. Similarly, $A_{k,F^{(\nu)}}(x) = \nu\, A_{k,F}(x/\nu)$. Hence $g_{F^{(\nu)}}(x) = \nu\, g_F(x/\nu)$, so $c_k(F^{(\nu)}) = \nu\, c_k(F)$. If $x_k^\star(F)$ is the unique minimizer of g_F, then $\nu x_k^\star(F)$ is the unique minimizer of $g_{F^{(\nu)}}$, and

$$J_k(F^{(\nu)}) = A_{k,F^{(\nu)}}(\nu x_k^\star(F)) = \nu\, A_{k,F}(x_k^\star(F)) = \nu\, J_k(F).$$ □

A.7 Proof of Proposition 2 (The missing middle)

Proof. Let $c \equiv c_k(F)$ and $J \equiv J_k(F) > 0$, and write

$$a(r) \equiv a_k(r; F), \qquad u(r) \equiv \alpha r - \frac{1}{2}r^2, \qquad \pi_\beta(r; F) \equiv u(r) - \beta a(r).$$

Because $\alpha > c$, the unconstrained optimum is $r^{\text{naive}} = \alpha > c$.

For $r \leq c$, Theorem 1 gives $a(r) = 0$, so $\pi_\beta(r; F) = u(r)$. Since $u'(r) = \alpha - r > 0$ on $[0, c]$, we have

$$\pi_\beta(r; F) < \pi_\beta(c; F) = u(c) \qquad \text{for all } r < c.$$

Define

$$f(\beta) \equiv \sup_{r>c} \pi_\beta(r; F).$$

Because $a(\cdot)$ is nondecreasing and $a(r) \geq \lim_{s \downarrow c} a(s) = J$ for every $r > c$, for any $0 \leq \beta_1 < \beta_2$ and any $r > c$,

$$\pi_{\beta_1}(r; F) = \pi_{\beta_2}(r; F) + (\beta_2 - \beta_1)a(r) \geq \pi_{\beta_2}(r; F) + (\beta_2 - \beta_1)J.$$

Taking suprema yields

$$f(\beta_1) \geq f(\beta_2) + (\beta_2 - \beta_1)J > f(\beta_2),$$

so f is strictly decreasing. It is also finite and convex, because it is the supremum of affine functions of β, and hence continuous on $(0, \infty)$.

Now $f(0) = \sup_{r>c} u(r) = u(r^{\text{naive}}) > u(c)$, while

$$f(\beta) \leq u(r^{\text{naive}}) - \beta J \to -\infty \qquad \text{as } \beta \to \infty.$$

Therefore there exists a unique $\bar{\beta}(\alpha; F) \in (0, \infty)$ such that

$$f(\bar{\beta}) = u(c), \qquad f(\beta) > u(c) \text{ for } \beta < \bar{\beta}, \qquad f(\beta) < u(c) \text{ for } \beta > \bar{\beta}.$$

Proof of part (i). If $\beta > \bar{\beta}$, then for every $r > c$,

$$\pi_\beta(r; F) \leq f(\beta) < u(c) = \pi_\beta(c; F).$$

Together with the strict inequality $\pi_\beta(r; F) < \pi_\beta(c; F)$ for all $r < c$, this implies that $r = c$ is the unique optimizer.

Knife-edge case $\beta = \bar{\beta}$. At $\beta = \bar{\beta}$, we have $f(\bar{\beta}) = u(c) = \pi_{\bar{\beta}}(c; F)$, so c is optimal. We next show that there is also an optimal interaction rate above the critical cut-off. Choose $\varepsilon \in (0, r^{\text{naive}} - c)$ small enough that

$$u(r) \leq u(c) + \frac{\bar{\beta}J}{4} \qquad \text{for all } r \in [c, c + \varepsilon].$$

Since $a(r) \geq J$ for all $r > c$, for every $r \in (c, c + \varepsilon]$,

$$\pi_{\bar{\beta}}(r; F) \leq u(c) + \frac{\bar{\beta}J}{4} - \bar{\beta}J = u(c) - \frac{3\bar{\beta}J}{4} < u(c) = f(\bar{\beta}).$$

So the supremum over $r > c$ is attained away from c. For $r > r^{\text{naive}}$, we have $u(r) < u(r^{\text{naive}})$ and $a(r) \geq a(r^{\text{naive}})$, hence

$$\pi_{\bar{\beta}}(r; F) < \pi_{\bar{\beta}}(r^{\text{naive}}; F).$$

Therefore maximizing over $r > c$ is equivalent to maximizing over the compact interval $[c+\varepsilon, r^{\text{naive}}]$. By the upper-semicontinuity argument given in the footnote after Equation (11), the map $r \mapsto \pi_{\bar{\beta}}(r; F)$ is upper semicontinuous on this interval: participation can only jump upward, so the induced jumps in $\pi_{\bar{\beta}}$ are downward. An upper semicontinuous function on a compact interval attains its maximum, so there exists $\tilde{r} \in [c + \varepsilon, r^{\text{naive}}]$ such that

$$\pi_{\bar{\beta}}(\tilde{r}; F) = f(\bar{\beta}) = u(c).$$

Thus $\tilde{r} > c$ is also optimal.

Proof of part (ii). If $\beta < \bar{\beta}$, then $f(\beta) > u(c) = \pi_\beta(c; F)$, so no optimizer can equal c. As shown above, no optimizer can lie below c either. Hence every optimizer $r^\star$ satisfies $r^\star > c$.

Now fix any optimal $r^\star > c$. Since c is feasible, $\pi_\beta(r^\star; F) \geq \pi_\beta(c; F)$, so

$$\alpha r^\star - \frac{1}{2}(r^\star)^2 - \beta a(r^\star) \geq \alpha c - \frac{1}{2}c^2.$$

Using $a(r^\star) \geq J$, a necessary condition is

$$\alpha r^\star - \frac{1}{2}(r^\star)^2 - \beta J \geq \alpha c - \frac{1}{2}c^2.$$

Writing $r^\star = c + \varepsilon$ with $\varepsilon > 0$ gives

$$\frac{1}{2}\varepsilon^2 - (\alpha - c)\varepsilon + \beta J \leq 0.$$

Hence ε must be at least the smaller root,

$$\varepsilon_- = (\alpha - c) - \sqrt{(\alpha - c)^2 - 2\beta J}.$$

Using $\sqrt{x^2 - y} \leq x - \frac{y}{2x}$ for $x > 0$, we obtain

$$\varepsilon_- \geq \frac{\beta J}{\alpha - c}.$$

Therefore

$$r^\star \geq c + \frac{\beta J_k(F)}{\alpha - c} = c_k(F) + d_k(F).$$

□

A.8 Proof of Proposition 3 (Comparative statics and tipping points)

Proof. Throughout, write $c \equiv c_k(F)$ and $a(r) \equiv a_k(r; F)$.

Part (i): Monotonicity. Fix $0 \leq \beta_1 < \beta_2$ and choose $r_i \in R(\beta_i)$. Optimality implies

$$\pi_{\beta_1}(r_1; F) \geq \pi_{\beta_1}(r_2; F), \qquad \pi_{\beta_2}(r_2; F) \geq \pi_{\beta_2}(r_1; F).$$

Adding these inequalities and canceling the common terms $\alpha r - \frac{1}{2}r^2$ gives

$$(\beta_2 - \beta_1)\big(a(r_2) - a(r_1)\big) \leq 0,$$

so $a(r_2) \leq a(r_1)$.

If $r_2 \le c$, then on $[0, c]$ we have $a(r) = 0$ and $\pi_{\beta_2}(r; F) = \alpha r - \frac{1}{2}r^2$, which is strictly increasing because $\alpha > c$. Hence any optimizer in $[0, c]$ must equal c, so $r_2 = c$. Since c is feasible at β_1, optimality of r_1 gives $\pi_{\beta_1}(r_1; F) \ge \pi_{\beta_1}(c; F)$, and therefore $r_1 \ge c = r_2$.

If instead $r_2 > c$, then Theorem 1(i) gives $a(r_2) > 0$. Since $a(r_2) \le a(r_1)$, we also have $a(r_1) > 0$, so again by Theorem 1(i) we must have $r_1 > c$. On (c, ∞), the map $a(\cdot)$ is strictly increasing by Theorem 2(i). Therefore $a(r_2) \le a(r_1)$ implies $r_2 \le r_1$.

Part (ii): Strict monotonicity. Under the maintained single-jump regularity (Assumption 2) and the transversality condition in Lemma 6, participation $a(\cdot)$ is C^2 and strictly increasing on (c, ∞). Assume $0 \le \beta_1 < \beta_2 < \bar{\beta}$ and that each $R(\beta_i)$ is the singleton $\{r_i\}$. By Proposition 2(ii), both optimizers satisfy $r_1, r_2 > c$. Part (i) already gives $r_2 \le r_1$.

Suppose, for contradiction, that $r_2 = r_1 \equiv \hat{r}$. Then the same interior point $\hat{r} > c$ is optimal for both β_1 and β_2. By Lemma 6, the map $r \mapsto \pi_\beta(r; F)$ is differentiable on (c, ∞), so $\hat{r}$ must satisfy the first-order condition

$$0 = \frac{\partial}{\partial r}\pi_\beta(\hat{r}; F) = \alpha - \hat{r} - \beta a'(\hat{r})$$

at both values of β. Since $a'(\hat{r}) > 0$ by Lemma 6, this equation determines β uniquely at a given $\hat{r}$, a contradiction. Hence $r_2 < r_1$. Because $a(\cdot)$ is strictly increasing on (c, ∞), it follows that

$$a(r_2) < a(r_1).$$

Part (iii): Local derivative. Let $\beta < \bar{\beta}(\alpha; F)$ be such that the maximizer is unique, interior, and nondegenerate, and denote it by $r^*(\beta) > c$. Define

$$H(\beta, r) \equiv \frac{\partial}{\partial r}\pi_\beta(r; F) = \alpha - r - \beta a'(r).$$

Then

$$H(\beta, r^*(\beta)) = 0, \qquad \frac{\partial H}{\partial r}(\beta, r^*(\beta)) = -1 - \beta a''(r^*(\beta)) = \frac{\partial^2}{\partial r^2}\pi_\beta(r^*(\beta); F) < 0.$$

By the implicit function theorem, $r^*(\cdot)$ is differentiable at β, with

$$\frac{\mathrm{d}r^*}{\mathrm{d}\beta} = -\frac{\partial H/\partial \beta}{\partial H/\partial r} = -\frac{-a'(r^*)}{-1 - \beta a''(r^*)} = -\frac{a'(r^*)}{1 + \beta a''(r^*)} < 0,$$

where the final inequality uses $a'(r^*) > 0$ and $-1 - \beta a''(r^*) < 0$. Writing $\bar{a}^* = a(r^*(\beta))$, we then have

$$\frac{\mathrm{d}\bar{a}^*}{\mathrm{d}\beta} = a'(r^*)\,\frac{\mathrm{d}r^*}{\mathrm{d}\beta} < 0.$$

□

Supporting lemmas

Lemma 2 (Small-x behavior of Φ_k under finite variance)**.** *Fix $k \geq 3$ and let $W \sim F$ satisfy $\mathbb{E}[W] = 1$, $\mathbb{P}(W > 0) > 0$, and $\mathbb{E}[W^2] < \infty$. Then*

$$\lim_{x \downarrow 0} \frac{\Phi_k(x)}{x} = 0.$$

Consequently, $g(x) = x/\Phi_k(x) \to \infty$ as $x \downarrow 0$, so any minimizer of g is interior.

A.9 Proof of Lemma 2 (Small-x behavior)

Proof. Because $k \geq 3$, $\psi_{k-1}(y) \leq \psi_2(y)$ for all $y \geq 0$. Moreover, for $y \geq 0$, $\psi_2(y) = 1 - e^{-y}(1+y) \leq \min\{1, y^2/2\}$. Hence

$$\Phi_k(x) = \mathbb{E}[W\psi_{k-1}(Wx)] \leq \mathbb{E}\big[W \min\{1, (Wx)^2/2\}\big] .$$

Split at $W \leq 1/x$:

$$\Phi_k(x) \leq \frac{x^2}{2}\mathbb{E}[W^3 \mathbf{1}_{\{W \leq 1/x\}}] + \mathbb{E}[W \mathbf{1}_{\{W > 1/x\}}].$$

For the tail term,

$$\mathbb{E}[W \mathbf{1}_{\{W > 1/x\}}] \leq x\, \mathbb{E}[W^2 \mathbf{1}_{\{W > 1/x\}}] = o(x),$$

since $\mathbb{E}[W^2] < \infty$ implies $\mathbb{E}[W^2 \mathbf{1}_{\{W > t\}}] \to 0$ as $t \to \infty$.

For the truncated third moment, let $T = 1/x$. By integration by parts,

$$\mathbb{E}[W^3 \mathbf{1}_{\{W \leq T\}}] = 3 \int_0^T t^2\, \mathbb{P}(W > t)\, dt - T^3 \mathbb{P}(W > T) \;\leq\; 3 \int_0^T t^2\, \mathbb{P}(W > t)\, dt.$$

Since $\mathbb{E}[W^2] < \infty$, we have $t^2 \mathbb{P}(W > t) \leq \mathbb{E}[W^2 \mathbf{1}_{\{W > t\}}] \to 0$. Therefore for any $\varepsilon > 0$ there exists T_0 such that $t^2 \mathbb{P}(W > t) \leq \varepsilon$ for all $t \geq T_0$, implying

$$\mathbb{E}[W^3 \mathbf{1}_{\{W \leq T\}}] \leq 3 \int_0^{T_0} t^2 \mathbb{P}(W > t)\, dt + 3\varepsilon(T - T_0) = O(1) + o(T).$$

Thus $\mathbb{E}[W^3 \mathbf{1}_{\{W \leq T\}}] = o(T)$, so $(x^2/2)\mathbb{E}[W^3 \mathbf{1}_{\{W \leq 1/x\}}] = o(x)$. Combining terms gives $\Phi_k(x) = o(x)$, i.e. $\Phi_k(x)/x \to 0$, and hence $x/\Phi_k(x) \to \infty$. □

A.10 Derivative identities for g and the H-criterion

Lemma 3 (H-criterion for extrema of g)**.** *Fix $k \geq 3$ and an interaction profile F satisfying Assumption 1. Let*

$$\Phi_k(x) \equiv \mathbb{E}[W\, \psi_{k-1}(Wx)] , \qquad g(x) \equiv \frac{x}{\Phi_k(x)} \quad (x > 0),$$

and define

$$H(x) \equiv x\,\Phi_k'(x) - \Phi_k(x).$$

Then:

(i) Φ_k *is twice continuously differentiable on* $(0,\infty)$*, with derivatives*

$$\Phi_k'(x) = \mathbb{E}\big[W^2\, p_{k-2}(Wx)\big]\,, \qquad \Phi_k''(x) = \mathbb{E}\big[W^3\,(p_{k-3}(Wx) - p_{k-2}(Wx))\big]\,,$$

where $p_m(z) \equiv \mathbb{P}(\mathrm{Poisson}(z) = m)$*.*

(ii) g *is continuously differentiable on* $(0,\infty)$ *and*

$$g'(x) = \frac{\Phi_k(x) - x\Phi_k'(x)}{\Phi_k(x)^2} = -\frac{H(x)}{\Phi_k(x)^2}.$$

Consequently, $x > 0$ *is a critical point of* g *if and only if* $H(x) = 0$*.*

(iii) H *is continuously differentiable on* $(0,\infty)$ *and satisfies the identity*

$$H'(x) = x\,\Phi_k''(x).$$

Proof. Part (i): For each fixed $w \geq 0$,

$$\frac{\partial}{\partial x}\Big[w\,\psi_{k-1}(wx)\Big] = w^2\, p_{k-2}(wx).$$

Because $p_{k-2}(z) \leq 1$ for all $z \geq 0$ and $\mathbb{E}[W^2] < \infty$, dominated convergence yields

$$\Phi_k'(x) = \mathbb{E}\big[W^2\, p_{k-2}(Wx)\big]\,.$$

To differentiate once more, fix a compact interval $I = [a,b] \subset (0,\infty)$. For every integer $m \geq 0$ and every $z \geq 0$,

$$z\,p_m(z) = (m+1)p_{m+1}(z) \leq m+1.$$

Hence, for every $x \in I$,

$$W^3 p_m(Wx) = W^2\,\frac{Wx\,p_m(Wx)}{x} \leq \frac{m+1}{a}\,W^2.$$

Applying this with $m = k-3$ and $m = k-2$, we obtain

$$W^3\big|p_{k-3}(Wx) - p_{k-2}(Wx)\big| \leq \frac{2k-3}{a}\,W^2,$$

which is integrable by Assumption 1. Since

$$\frac{\partial}{\partial x}\Big[W^2 p_{k-2}(Wx)\Big] = W^3\big(p_{k-3}(Wx) - p_{k-2}(Wx)\big),$$

dominated convergence gives

$$\Phi_k''(x) = \mathbb{E}\big[W^3\big(p_{k-3}(Wx) - p_{k-2}(Wx)\big)\big]\,, \qquad x \in I.$$

The same domination is uniform on I, so Φ_k'' is continuous on I. Since $I \subset (0,\infty)$ was arbitrary, $\Phi_k \in C^2((0,\infty))$.

Part (ii): This is the quotient rule:

$$g'(x) = \frac{\Phi_k(x) - x\Phi_k'(x)}{\Phi_k(x)^2} = -\frac{H(x)}{\Phi_k(x)^2}.$$

Hence $g'(x) = 0$ if and only if $H(x) = 0$.

Part (iii): By definition,

$$H(x) = x\Phi_k'(x) - \Phi_k(x),$$

so, using Part (i),

$$H'(x) = \Phi_k'(x) + x\Phi_k''(x) - \Phi_k'(x) = x\Phi_k''(x).$$

□

A.11 Uniform small-x bound under bounded dispersion

Lemma 4 (Uniform small-x bound under bounded dispersion)**.** *Fix $k \geq 3$ and let $W \sim F$ satisfy $\mathbb{E}[W] = 1$ and $\mathrm{Var}(W) = \sigma^2 < \infty$. Then for all $x \in (0, 1/2]$,*

$$\Phi_k(x) = \mathbb{E}[W\psi_{k-1}(Wx)] \;\leq\; 4x^2 + 6\sigma^2 x. \tag{18}$$

Consequently, for all $x \in (0, 1/2]$,

$$g(x) = \frac{x}{\Phi_k(x)} \;\geq\; \frac{1}{4x + 6\sigma^2}. \tag{19}$$

In particular, for any $a \in (0, 1/2]$,

$$\inf_{0<x\leq a} g(x) \;\geq\; \frac{1}{4a + 6\sigma^2}.$$

Proof. Because $k \geq 3$ we have $\psi_{k-1}(y) \leq \psi_2(y)$ for all $y \geq 0$. Moreover, if $N \sim \mathrm{Poisson}(y)$ then

$1_{\{N\ge 2\}} \le N(N-1)/2$, so

$$\psi_2(y) = \mathbb{P}(N \ge 2) \le \frac{\mathbb{E}[N(N-1)]}{2} = \frac{y^2}{2},$$

and hence $\psi_2(y) \le \min\{1, y^2/2\}$. Therefore,

$$\Phi_k(x) \le \mathbb{E}\left[W \min\left\{1, \frac{(Wx)^2}{2}\right\}\right] \le \frac{x^2}{2}\mathbb{E}\big[W^3 \mathbf{1}_{\{W\le 1/x\}}\big] + \mathbb{E}\big[W \mathbf{1}_{\{W>1/x\}}\big]\,.$$

Assume $x \le 1/2$, so $1/x \ge 2$. For the first term, split at $W \le 2$:

$$\mathbb{E}\big[W^3\mathbf{1}_{\{W\le 1/x\}}\big] \le \mathbb{E}\big[W^3\mathbf{1}_{\{W\le 2\}}\big] + \mathbb{E}\big[W^3\mathbf{1}_{\{2<W\le 1/x\}}\big] \le 8 + \frac{1}{x}\mathbb{E}\big[W^2\mathbf{1}_{\{W>2\}}\big]\,,$$

since $W^3 \le (1/x)W^2$ on $\{W \le 1/x\}$. On $\{W > 2\}$ we have $W-1 \ge 1$, so Chebyshev implies $\mathbb{P}(W > 2) \le \sigma^2$ and hence

$$\mathbb{E}\big[W^2\mathbf{1}_{\{W>2\}}\big] \le 2\mathbb{E}[(W-1)^2] + 2\mathbb{P}(W>2) \le 4\sigma^2.$$

Thus

$$\frac{x^2}{2}\mathbb{E}\big[W^3\mathbf{1}_{\{W\le 1/x\}}\big] \le 4x^2 + 2\sigma^2 x.$$

For the second term, write $W = 1 + (W-1)$ and use Cauchy–Schwarz and Chebyshev: for $t > 1$,

$$\mathbb{E}[W\mathbf{1}_{\{W>t\}}] \le \mathbb{E}[(W-1)\mathbf{1}_{\{W>t\}}] + \mathbb{P}(W>t) \le \sigma\sqrt{\mathbb{P}(W>t)} + \mathbb{P}(W>t) \le \frac{\sigma^2}{t-1} + \frac{\sigma^2}{(t-1)^2}.$$

Taking $t = 1/x$ yields

$$\mathbb{E}\big[W\mathbf{1}_{\{W>1/x\}}\big] \le \frac{\sigma^2 x}{1-x} + \frac{\sigma^2 x^2}{(1-x)^2} \le 4\sigma^2 x,$$

where the last inequality uses $x \le 1/2$. Combining bounds gives (18). Dividing by x yields (19). □

A.12 Uniform approximation on compacts

Lemma 5 (Uniform approximation on compacts)**.** *Fix $k \ge 3$ and $X > 0$, and let $W \sim F$ satisfy $\mathbb{E}[W] = 1$. Then*

$$\sup_{0\le x\le X} \big|\Phi_k(x) - \psi_{k-1}(x)\big| \le (1+X)\,\mathbb{E}|W-1|, \tag{20}$$

$$\sup_{0\le x\le X} \big|A_k(x) - \psi_k(x)\big| \le X\,\mathbb{E}|W-1|. \tag{21}$$

If additionally $\mathbb{E}[W^2] < \infty$*, then*

$$\sup_{0 \le x \le X} \left|\Phi_k'(x) - \psi_{k-1}'(x)\right| \le \mathbb{E}|W^2 - 1| + X\,\mathbb{E}|W - 1|. \tag{22}$$

Proof. We use only that ψ_m and p_m are 1-Lipschitz in their argument: since $\psi_m'(y) = p_{m-1}(y) \in [0,1]$ for $y \ge 0$, we have $|\psi_m(y) - \psi_m(z)| \le |y - z|$; similarly $p_m'(y) = p_{m-1}(y) - p_m(y) \in [-1,1]$ so $|p_m(y) - p_m(z)| \le |y - z|$.

Bound (20)*.* Write

$$\Phi_k(x) - \psi_{k-1}(x) = \mathbb{E}\big[W\psi_{k-1}(Wx) - \psi_{k-1}(x)\big].$$

Add and subtract $\psi_{k-1}(Wx)$:

$$W\psi_{k-1}(Wx) - \psi_{k-1}(x) = (W-1)\psi_{k-1}(Wx) + \big(\psi_{k-1}(Wx) - \psi_{k-1}(x)\big).$$

Since $\psi_{k-1}(\cdot) \in [0,1]$, the first term has absolute value at most $|W - 1|$. By Lipschitzness, the second term is bounded by $|Wx - x| = x|W - 1|$. Taking expectations yields $|\Phi_k(x) - \psi_{k-1}(x)| \le (1+x)\mathbb{E}|W-1|$, and taking $x \le X$ gives (20).

Bound (21)*.* By Lipschitzness of ψ_k,

$$|A_k(x) - \psi_k(x)| = \big|\mathbb{E}[\psi_k(Wx) - \psi_k(x)]\big| \le \mathbb{E}[|Wx - x|] = x\,\mathbb{E}|W-1| \le X\,\mathbb{E}|W-1|.$$

Bound (22)*.* Differentiate Φ_k:

$$\Phi_k'(x) = \mathbb{E}\big[W^2\,\psi_{k-1}'(Wx)\big] = \mathbb{E}\big[W^2\,p_{k-2}(Wx)\big], \qquad \psi_{k-1}'(x) = p_{k-2}(x).$$

Then

$$\Phi_k'(x) - p_{k-2}(x) = \mathbb{E}\big[(W^2-1)p_{k-2}(Wx)\big] + \mathbb{E}\big[p_{k-2}(Wx) - p_{k-2}(x)\big].$$

The first term is bounded by $\mathbb{E}|W^2 - 1|$ since $p_{k-2} \in [0,1]$. The second term is bounded by $\mathbb{E}|Wx - x| = x\mathbb{E}|W-1| \le X\mathbb{E}|W-1|$ by Lipschitzness of p_{k-2}. This yields (22). □

A.13 Smoothness and strict monotonicity of participation above the cutoff

Lemma 6 (Smoothness of participation above the cutoff)**.** *Assume the single-jump regularity in Assumption 2: the function* $g(x) = x/\Phi_k(x)$ *has a unique minimizer* $x_k^\star(F)$ *and is strictly increasing on* $[x_k^\star(F), \infty)$*. Assume moreover the transversality condition*

$$g'(x) > 0 \qquad \text{for all } x > x_k^\star(F)$$

(equivalently, $\Phi_k(x) > x\Phi_k'(x)$ for all $x > x_k^\star(F)$). For each $r > c \equiv c_k(F)$, let $x(r)$ be the unique solution to $x = r\Phi_k(x)$ in $[x_k^\star(F), \infty)$, and recall $a(r) \equiv a_k(r; F) = A_k(x(r))$.

Then $r \mapsto x(r)$ and $r \mapsto a(r)$ are C^2 on (c, ∞). Moreover, for all $r > c$,

$$a'(r) = A_k'(x(r))\, x'(r) > 0.$$

In particular, $a(\cdot)$ is strictly increasing on (c, ∞), and for each $\beta > 0$ the map $r \mapsto \pi_\beta(r; F) = \alpha r - \frac{1}{2}r^2 - \beta\, a(r)$ is C^2 on (c, ∞).

Proof. Write $p_m(z) \equiv \mathbb{P}(\mathrm{Poisson}(z) = m) = e^{-z} z^m / m!$.

Φ_k is C^2 on $(0, \infty)$. By Lemma 3(i), $\Phi_k \in C^2((0, \infty))$ with derivatives $\Phi_k'(x) = \mathbb{E}[W^2\, p_{k-2}(Wx)]$ and $\Phi_k''(x) = \mathbb{E}[W^3(p_{k-3}(Wx) - p_{k-2}(Wx))]$.

A_k is C^2 on $(0, \infty)$. Since $\psi_k'(z) = p_{k-1}(z)$ and $|p_{k-1}(z)| \le 1$, dominated convergence (with integrable dominator W) gives, for every $x > 0$,

$$A_k'(x) = \mathbb{E}[W\, p_{k-1}(Wx)]\,.$$

For the second derivative, $\frac{\partial}{\partial x}[W\, p_{k-1}(Wx)] = W^2(p_{k-2}(Wx) - p_{k-1}(Wx))$, and

$$W^2 \big| p_{k-2}(Wx) - p_{k-1}(Wx) \big| \le 2W^2,$$

which is integrable by $\mathbb{E}[W^2] < \infty$ (Assumption 1). Dominated convergence therefore yields

$$A_k''(x) = \mathbb{E}\big[W^2 \big(p_{k-2}(Wx) - p_{k-1}(Wx)\big)\big]\,,$$

and A_k'' is continuous by the same domination, so $A_k \in C^2((0, \infty))$.

For $r > c$, the single-jump regularity implies that the equation $r = g(x)$ has a unique solution $x = x(r) \in [x_k^\star(F), \infty)$, and this $x(r)$ is exactly the largest-fixed-point branch solving $x = r\Phi_k(x)$. Since $c = g(x_k^\star(F))$ and g is strictly increasing on $[x_k^\star(F), \infty)$, we have $x(r) > x_k^\star(F)$ for all $r > c$. By transversality, $g'(x(r)) > 0$. Now

$$g'(x) = \frac{\Phi_k(x) - x\Phi_k'(x)}{\Phi_k(x)^2},$$

so $g'(x(r)) > 0$ implies $\Phi_k(x(r)) > x(r)\Phi_k'(x(r))$. Using $x(r) = r\Phi_k(x(r))$, this is equivalent to

$$1 - r\Phi_k'(x(r)) > 0.$$

Hence the implicit function theorem applied to $H(r, x) \equiv x - r\Phi_k(x) = 0$ gives that $r \mapsto x(r)$ is C^2 on (c, ∞), with derivative

$$x'(r) = \frac{\Phi_k(x(r))}{1 - r\Phi_k'(x(r))} > 0.$$

Since A_k is C^2, $a(r) = A_k(x(r))$ is C^2 on (c, ∞), and

$$a'(r) = A'_k(x(r))\, x'(r).$$

Finally, $A'_k(x(r)) > 0$ because $\mathbb{P}(W > 0) > 0$ (Assumption 1) and $p_{k-1}(Wx(r)) > 0$ whenever $W > 0$ and $x(r) > 0$. Thus $a'(r) > 0$ for all $r > c$. $\square$

A.14 Generic uniqueness of the principal's optimum

Lemma 7 (Generic uniqueness)**.** *Fix $\alpha > 0$ and define the value function*

$$V(\beta) \equiv \sup_{r \geq 0} \Big(\alpha r - \tfrac{1}{2}r^2 - \beta\, a_k(r; F)\Big).$$

Then V is convex in β, hence differentiable for all but countably many β. Moreover, if V is differentiable at β, then $\arg\max_{r \geq 0} \pi_\beta(r; F)$ *is a singleton.*

Proof. For each fixed r, $\beta \mapsto \pi_\beta(r; F)$ is affine with slope $-a_k(r; F)$, so $V(\beta) = \sup_r \pi_\beta(r; F)$ is convex as the pointwise supremum of affine functions. A convex function on an interval is differentiable except at a countable set of points.

Now fix β and suppose there are two distinct global maximizers $r_1 < r_2$. If $r_2 \leq c_k(F)$, then $a_k(r_1; F) = a_k(r_2; F) = 0$ and $\pi_\beta(r; F) = \alpha r - \frac{1}{2}r^2$ on $[0, c_k(F)]$, which has a unique maximizer on $[0, c_k(F)]$ (equal to $\min\{\alpha, c_k(F)\}$), a contradiction. Hence $r_2 > c_k(F)$, and by strict monotonicity of $a_k(\cdot; F)$ on $(c_k(F), \infty)$ (Theorem 2(i)), we have $a_k(r_1; F) \neq a_k(r_2; F)$. Therefore the two affine functions $\beta \mapsto \pi_\beta(r_1; F)$ and $\beta \mapsto \pi_\beta(r_2; F)$ have different slopes, so V cannot be differentiable at β. $\square$

B Online Appendix

This appendix collects supplementary material: microfoundations for equilibrium selection and the payoff specification, intuition for the discontinuous jump in participation, and several technical remarks and extensions that support but are not required for the main-text proofs.

B.1 Micro-founding the largest equilibrium

In Section 3 we focus on the largest equilibrium of the second stage (i.e. the one with the most participating agents). As discussed there, there are two natural interpretations of this focus, depending on the application. In entrenched-behavior settings, best-response dynamics from the all-participate profile select the largest equilibrium as a literal prediction. In coordination settings, the largest equilibrium characterizes the coordination capacity that the principal designs against.

Here, we formalize these two interpretations by providing two micro-foundations for the focus on the largest equilibrium. In the first, agents are boundedly rational and play myopic best responses. In the second, agents are fully rational and reach equilibrium using only local information and a short communication phase. Both micro-foundations yield the same outcome: the set of participating agents is the k-core of the realized graph.

Unsophisticated agents: myopic best-response dynamics

Set-up. Suppose that the game is played over $n+1$ periods. At $t=1$, the principal chooses $r \geq 0$, and Nature realizes the graph G (as in our baseline model). Additionally, at $t=1$, all agents participate *non-strategically* (i.e. $a_{i,t=1}=1$ for all i). In each subsequent period $t \geq 2$, all agents best respond myopically to actions in the preceding period $t-1$. Agent i's action in period $t \geq 2$ is thus determined by the best response

$$BR_{i,t}^{(n)}(M_{i,t-1}) = \begin{cases} 1, & \text{if } M_{i,t-1} \geq k, \\ 0, & \text{if } M_{i,t-1} < k, \end{cases} \tag{23}$$

where $M_{i,t-1} = \sum_{j \neq i} G_{ij} a_{j,t-1}$. For convenience, we write $a_i \equiv a_{i,n+1}$ for i's action in the final period and assume that the principal cares about participation at the end of the game, not during it.

The result. Exactly as in our baseline model, the unique limit of the above dynamic process is that each agent participates if and only if he belongs to the k-core of the graph G. In other words, Remark 1 holds unchanged.

Discussion. In addition to myopia, this set-up assumes that all agents start off by participating. From a technical standpoint, this guarantees convergence of the best-response dynamics to the largest equilibrium of the static game. The logic is very similar to that laid out after Remark 1: agents iteratively drop out when they have fewer than k participating neighbors, and the process terminates at the k-core. The only difference is that agents are updating myopically rather than performing analogous reasoning in their heads.

Economically, this set-up imposes $a=1$ as the pre-existing default. The myopic-updating view therefore best fits settings where the principal wants to eradicate or reduce pre-existing behaviors. This is natural when considering the adoption of new technologies or new social norms: the old versions were the status quo before the start of the game, so everyone must have been using the old technology or adhering to the old norm. Note that in this interpretation our terminology becomes less clean: "participate" means to use the *old* technology or norm, and "abstain" means to adopt the *new* one.

Sophisticated agents: local communication

At $t = 1$, the principal chooses $r \geq 0$, Nature realizes the graph G, and all agents announce $m_{i,1} \equiv \text{P}$. This is followed by a "communication phase" of n periods (periods 2 through $n+1$) and an "action phase" at $t = n+2$. In each period $t \geq 2$, agents observe only (i) how many neighbors they have and (ii) what each of those neighbors announced in the previous period.

Communication phase (periods 2 to n+1). During the communication phase, every agent i announces a *cheap talk* message $m_{i,t} \in \{P, A\}$, interpreted as "I Plan to participate" or "I will Abstain." Suppose agents adopt the rule:

$$m_{i,t} = \begin{cases} \text{A} & \text{if strictly fewer than } k \text{ neighbours sent P at } t-1, \\ \text{P} & \text{otherwise.} \end{cases}$$

Action phase (period n+2). Consider the strategy s_i given by: play $a_i = 1$ if and only if at least k of i's neighbors sent P in the final communication round (period $n+1$).

The result. The message rule implements the usual k-core peeling process: after (at most) n periods, no further agents switch to sending A's and the only agents who still announce P are the vertices in the k-core of G.

Given the strategy s_i for the action phase, any agent who still announces P in the final communication round has at least k neighbors who also announced P. These agents will therefore choose $a_i = 1$ in the action phase. Hence playing $a_i = 1$ is a best response for every agent who announces P at time $n+1$, that is, for every agent in the k-core. Any agent who observes fewer than k neighbors sending P in the final round plays $a_i = 0$, which is his best response since fewer than k of his neighbors will participate. Thus (while not strictly required by Nash equilibrium) $s = (s_i)_{i \in N}$ has the advantage of being sequentially rational in the action phase.

In the communication phase, no agent can profit (in any period) by unilaterally deviating from the message rule. Sending P when the rule prescribes A cannot push the number of remaining P-neighbors back to k, so it cannot raise the deviator's final payoff. Sending A when the rule prescribes P can only reduce the number of P-senders among the deviator's neighbors (by triggering further removals), weakly lowering his payoff. Therefore the prescribed message profile $m = (m_{i,t})_{i \in N, 1 \leq t \leq n+1}$ is itself incentive-compatible.

Consequently, (m, s) is a Nash equilibrium of the extensive-form game. The agents who play $a_i = 1$ in the terminal period are exactly those in the k-core of G.

Discussion. Agents here are fully rational and use the communication phase to perform the same iterated elimination that the "unsophisticated" agents achieve through myopic play. Crucially, each agent needs only *local* information: his own degree and the most recent messages of his neighbors. The procedure demonstrates that the k-core is a natural characterization of equilibrium even when agents are fully sophisticated, as in our baseline model (see Remark 1).

B.2 The discontinuous emergence of the k-core: intuition

Suppose we choose a participating vertex at random in the graph and follow it to one of its neighbors, say i. Consider the $k = 2$ case. For i to be part of a group which sustains the behavior, we only need to find one *additional* participating neighbor. This is a relatively weak requirement: it allows for a small cluster of agents to be self-sustaining. If the probability that a random neighbor participates is very small (say, ρ), an agent's chance of finding at least one participating neighbor is also small (roughly proportional to ρ). This allows for a stable outcome where a tiny fraction of agents can mutually sustain each other, and this fraction grows continuously as the network becomes denser.

Now consider the $k = 3$ case. For i to be part of a group which sustains the behavior, we need to find two other participating neighbors. Because the probability of finding multiple participating neighbors compounds, this is a much stronger requirement. If ρ is very small, the probability of any given agent finding two participating neighbors is drastically smaller (roughly proportional to ρ^2). A tiny self-sustaining group is therefore impossible; the demanding requirement for participation cannot be met if the pool of potential supporters is itself vanishingly small.

The system cannot be "born small." For a k-core with $k \geq 3$ to exist at all, the network must already be dense enough to support a large concentration of participants, leading to the discontinuous jump from zero to a substantial fraction of the population.

This logic can be formalized through branching process approximations of the local network structure, and it applies to any degree distribution with finite variance—not just the Erdős–Rényi case. In our mixed-Poisson model, the local structure of the network is a multi-type branching process indexed by the interaction weight W, but the key mechanism is the same: the fixed-point equation for the survival probability ξ involves ξ^{k-1}, which is flat near the origin for $k \geq 3$, preventing continuous emergence of participation.

B.3 Heavy-tail edge case (outside Assumption 1)

Lemma 2 is exactly where the finite-variance restriction is used to force $g(x) = x/\Phi_k(x) \to \infty$ as $x \downarrow 0$. If one leaves this moment class (for example, sufficiently heavy-tailed profiles with $\mathbb{E}[W^2] = \infty$), that boundary behavior can fail. Then the minimizer of g need not be interior, and participation can emerge continuously with $J_k(F) = 0$. We do not analyze that tail regime in this paper. All main-text results assume Assumption 1.

B.4 A stage-0 microfoundation for the interaction profile

Proposition 4 (Binary access microfoundation implies softened core–periphery)**.** *Assume each agent draws an i.i.d. cost $C_i \sim G$ and chooses channel access $X_i \in \{0,1\}$ at stage 0, with $X_i = 1$ if and only if $C_i \le B$, where B is a common benefit. Let $\nu \equiv \mathbb{P}(X_i = 1) = G(B) \in (0,1)$, and fix $\epsilon \in [0,1)$. Define interaction weights*

$$W_i = \begin{cases} \epsilon, & X_i = 0, \\ \dfrac{1-\epsilon(1-\nu)}{\nu}, & X_i = 1. \end{cases}$$

Then (W_i) are i.i.d. with distribution $F_{\nu,\epsilon}$, $\mathbb{E}[W_i] = 1$, and $\mathbb{E}[W_i^2] < \infty$. Moreover, as $\epsilon \downarrow 0$, $F_{\nu,\epsilon}$ converges to the sharp two-class profile F_ν.

Proof. By construction, $\mathbb{P}(W_i = \epsilon) = 1 - \nu$ and $\mathbb{P}\Big(W_i = \frac{1-\epsilon(1-\nu)}{\nu}\Big) = \nu$. Hence

$$\mathbb{E}[W_i] = (1-\nu)\epsilon + \nu\frac{1-\epsilon(1-\nu)}{\nu} = 1.$$

The support has two finite points, so $\mathbb{E}[W_i^2] < \infty$. Taking $\epsilon \downarrow 0$ yields the sharp two-class profile F_ν. □

B.5 Congestion-based payoff microfoundation

The baseline payoff in Equation (1) is reduced-form:

$$\pi(r, F, \bar{a}) = \alpha r - \frac{1}{2}r^2 - \beta\bar{a}.$$

One microfoundation that adds an F-dependent curvature term is the following. Suppose the principal's value from agent i's realized degree D_i is

$$v(D_i) = bD_i - \frac{\kappa}{2}D_i^2, \qquad b > 0,\ \kappa \ge 0,$$

and $D_i \mid W_i \sim \text{Poisson}(rW_i)$ with $\mathbb{E}[W] = 1$. Then

$$\mathbb{E}[v(D_i) \mid W_i] = \Big(b - \frac{\kappa}{2}\Big) rW_i - \frac{\kappa}{2}(rW_i)^2.$$

Averaging over W gives

$$\mathbb{E}[v(D)] = \Big(b - \frac{\kappa}{2}\Big) r - \frac{\kappa}{2}r^2\mathbb{E}[W^2].$$

Adding a direct infrastructure cost $\frac{1}{2}r^2$ and participation cost $\beta\bar{a}$ yields

$$\pi(r, F, \bar{a}) = \alpha r - \frac{1}{2}r^2 - \frac{\kappa}{2}r^2\mathbb{E}_F[W^2] - \beta\bar{a}, \qquad \alpha \equiv b - \frac{\kappa}{2}.$$

To match the maintained assumption $\alpha > 0$ in the main text, this microfoundation requires $b > \kappa/2$. When F is fixed, $1 + \kappa\mathbb{E}_F[W^2]$ is constant, so this extension is equivalent to rescaling the quadratic coefficient on r^2 and leaves the qualitative missing-middle logic unchanged.

B.6 Interpretation of the H-criterion

Remark 4 (Interpretation: when do we get a single jump?)**.** *Lemma 3 provides a simple diagnostic for the single-jump regularity condition used in Proposition 1. Since* $g'(x) = -H(x)/\Phi_k(x)^2$ *and* $\Phi_k(x)^2 > 0$*, every local extremum of* g *occurs at a zero of* H *on* $(0, \infty)$*. In particular, if* H *crosses zero exactly once (equivalently,* g *decreases and then increases once), then* g *has a unique minimizer and the giant* k*-core emerges through a single discontinuous jump. If* H *has multiple zeros (which can occur in highly heterogeneous, bimodal interaction profiles), then* g *can have multiple local minima, and the* k*-core may exhibit multiple branches or multiple discontinuities.*

B.7 Convex order does not imply monotonicity of the jump

The following example shows that ordering interaction profiles by the convex order does not guarantee monotonicity of the jump.

Remark 5 (Convex order counterexample for $k = 3$)**.** *Define two interaction profiles, each with* $\mathbb{E}[W] = 1$*:*

$$F_1: \quad W = \begin{cases} \frac{30}{49} & \text{w.p. } 0.7, \\ \frac{60}{49} & \text{w.p. } 0.1, \\ \frac{110}{49} & \text{w.p. } 0.2; \end{cases} \qquad\qquad F_2: \quad W = \begin{cases} \frac{10}{81} & \text{w.p. } 0.5, \\ \frac{130}{81} & \text{w.p. } 0.4, \\ \frac{240}{81} & \text{w.p. } 0.1. \end{cases}$$

One can verify that $F_1 \leq_{\text{cx}} F_2$ *(i.e.,* F_2 *is a mean-preserving spread of* F_1*) by checking the stop-loss order:* $\mathbb{E}_{F_1}[(W-t)^+] \leq \mathbb{E}_{F_2}[(W-t)^+]$ *for all* $t \geq 0$.[27] *The variances are* $\text{Var}(W) \approx 0.420$ *under* F_1 *and* $\text{Var}(W) \approx 0.916$ *under* F_2*. However, direct computation gives*

$$J_3(F_1) \approx 0.1003 \quad < \quad J_3(F_2) \approx 0.1040.$$

The more dispersed profile has the larger *jump.*

The obstruction is that $J_k(F) = A_k(x_k^\star(F))$ *depends on both the argmin* $x_k^\star(F)$ *of* $g^F(x) = x/\Phi_k^F(x)$ *and the composition with* A_k^F*. Neither the argmin map nor the composition preserves*

[27] Both distributions have $\mathbb{E}[W] = 1$. Direct computation of $\mathbb{E}[(W-t)^+]$ for each distribution confirms the inequality for all $t \geq 0$, which is equivalent to the convex order.

convex-order monotonicity.

B.8 Robustness to a positive-degree periphery

This section shows that the comparative-statics results derived for the two-class family F_ν (Section 4.4) are robust to giving the periphery strictly positive interaction weights.

The softened family. For $\epsilon \in [0,1)$, define

$$F_{\nu,\epsilon}: \quad W = \begin{cases} \epsilon & \text{w.p. } 1-\nu, \\ w_c(\nu,\epsilon) \equiv \dfrac{1-\epsilon(1-\nu)}{\nu} & \text{w.p. } \nu, \end{cases}$$

so that $\mathbb{E}[W] = 1$. Equivalently, $W = \epsilon + (1-\epsilon)X$ where $X \sim F_\nu$: every agent receives a uniform baseline weight ϵ, and the remaining mass is distributed according to the sharp two-class family. When $\epsilon = 0$, $F_{\nu,0} = F_\nu$. When $\epsilon > 0$, all agents have strictly positive expected degree (ϵr for the periphery, $w_c r$ for the core).

Key objects. Writing $\Phi_{k,\epsilon}(x) \equiv \mathbb{E}_{F_{\nu,\epsilon}}[W\psi_{k-1}(Wx)]$, $A_{k,\epsilon}(x) \equiv \mathbb{E}_{F_{\nu,\epsilon}}[\psi_k(Wx)]$, and $g_\epsilon(x) \equiv x/\Phi_{k,\epsilon}(x)$, the critical cutoff is $c_k(F_{\nu,\epsilon}) = \inf_{x>0} g_\epsilon(x)$ and the jump is $J_k(F_{\nu,\epsilon}) = A_{k,\epsilon}(x^\star_\epsilon)$, where $x^\star_\epsilon$ is the minimizer of g_ϵ.

Proposition 5 (Positive-degree periphery robustness)**.** *Fix $k \geq 3$ and $\nu \in (0,1)$. Assume $g_0(x) = x/\Phi_{k,0}(x)$ has a unique minimizer $x^\star_0$ (i.e., the single-jump condition holds for F_ν). For the rate claim in part (ii), assume in addition that $g''_0(x^\star_0) > 0$. Then:*

(i) (Continuity.) *$c_k(F_{\nu,\epsilon}) \to \nu\, c_k^{ER}$ and $J_k(F_{\nu,\epsilon}) \to \nu\, J_k^{ER}$ as $\epsilon \downarrow 0$.*

(ii) (Rates.) *The critical cutoff satisfies*

$$c_k(F_{\nu,\epsilon}) = \nu\, c_k^{ER} + O(\epsilon).$$

The jump is first-order stable:

$$J_k(F_{\nu,\epsilon}) = \nu\, J_k^{ER} + O(\epsilon),$$

where the implicit constants in $O(\epsilon)$ depend only on k and ν.

Proof. The argument proceeds in three steps.

Step 1: Uniform convergence on compacts. Since ψ_m is C^∞ and 1-Lipschitz, and the support of $F_{\nu,\epsilon}$ is bounded uniformly in ϵ (for fixed ν), $\Phi_{k,\epsilon}$ and $A_{k,\epsilon}$ converge uniformly to $\Phi_{k,0}$ and $A_{k,0}$

on every compact $[0, b]$:

$$\sup_{x\in[0,b]} \left|\Phi_{k,\epsilon}(x) - \Phi_{k,0}(x)\right| = O(\epsilon), \qquad \sup_{x\in[0,b]} \left|A_{k,\epsilon}(x) - A_{k,0}(x)\right| = O(\epsilon).$$

The bound follows from bounding the periphery term $(1-\nu)\epsilon\, \psi_{k-1}(\epsilon x) \leq (1-\nu)\epsilon$ and the shift in core weight $|w_c - 1/\nu| = O(\epsilon)$. The same holds for first derivatives.

Step 2: Minimizers stay bounded. For any fixed ν, $g_\epsilon(x) \to \infty$ as $x \downarrow 0$ (since $\Phi_{k,\epsilon}(x) = O(x^{k-1})$ for $k \geq 3$) and as $x \to \infty$ (since $\Phi_{k,\epsilon}(x) \to 1$), both uniformly in ϵ. Hence all minimizers of g_ϵ lie in a fixed compact $[a, b]$ for ϵ small.

Step 3: Argmin stability (parts (i) and (ii)). By Steps 1–2 and the uniqueness of $x_0^\star$, the Berge maximum theorem (applied to $-g_\epsilon$ on $[a, b]$) gives $x_\epsilon^\star \to x_0^\star$ and $c_k(F_{\nu,\epsilon}) = g_\epsilon(x_\epsilon^\star) \to g_0(x_0^\star) = \nu\, c_k^{ER}$. The Lipschitz bound from Step 1 gives the $O(\epsilon)$ rate for the cutoff.

For the jump, decompose

$$J_k(F_{\nu,\epsilon}) - \nu J_k^{ER} = \left[A_{k,\epsilon}(x_\epsilon^\star) - A_{k,0}(x_\epsilon^\star)\right] + \left[A_{k,0}(x_\epsilon^\star) - A_{k,0}(x_0^\star)\right].$$

The first bracket is $O(\epsilon)$ by Step 1 (uniform convergence of $A_{k,\epsilon}$ to $A_{k,0}$ on $[a, b]$). For the second bracket, use the nondegeneracy $g_0''(x_0^\star) > 0$. Since $g_0'(x_0^\star) = 0$, this yields a locally unique smooth critical-point branch under small C^1 perturbations of g_0'; combined with Step 1 (which gives $g_\epsilon' = g_0' + O(\epsilon)$ on $[a, b]$), we obtain $x_\epsilon^\star - x_0^\star = O(\epsilon)$. Because $A_{k,0}'$ is bounded on $[a, b]$, the second bracket is also $O(\epsilon)$. Hence

$$J_k(F_{\nu,\epsilon}) = \nu\, J_k^{ER} + O(\epsilon).$$

□