Strategic Influence: The Diffusion of Prosocial and Antisocial Behaviors*

Arthur Campbell† D.J. Thornton‡ Yves Zenou§

December 10, 2025

Abstract

A growing body of empirical evidence reveals a fundamental asymmetry in the diffusion of social behaviors: prosocial behaviors often exhibit strategic substitutability ("bystander effects"), while antisocial behaviors exhibit strategic complementarities ("licensing effects"). Moreover, even a single behavior (e.g., protest participation) can be a complement in some settings and a substitute in others. To unify these findings, we develop a model of strategic diffusion on networks with positive (prosocial) or negative (antisocial) spillovers. Our results rely on a novel conception of *influence*, capturing the causal impact of an individual's adoption on others. Prosocial behaviors are complementary in sparse networks but substitutable in dense ones, while antisocial behaviors exhibit the reverse pattern. Our model predicts that prosocial behaviors emerge continuously, while antisocial behaviors exhibit a discontinuous, sudden emergence. Effective policies can target the network density or perceptions of the extent of spillovers to encourage prosocial behaviors and inhibit antisocial ones.

Keywords: Diffusion, spillovers, influence, externalities, narratives, networks, bystander effect, licensing effect, mob-rule, prosocial behavior, antisocial behavior.

JEL classification codes: D43, D85, L13.

*This paper has been previously circulated under the title "Strategic Diffusion: Public Goods vs. Public Bads." We are grateful to Leeat Yariv, three anonymous referees, the participants at the 2023 Brown Bag seminar at Monash University, the 2024 Australasian Economic Theory Workshop at the University of Melbourne, the 2024 Network Science and Economics Conference, and the 2024 Barcelona School of Economics Summer Forum Networks Workshop, for helpful comments. The financial support from the Australian Research Council (DP200102547) is gratefully acknowledged.

†Monash University, Australia. E-mail: arthur.campbell@monash.edu

‡University of New South Wales, Australia. E-mail: d.thornton@unsw.edu.au

§Monash University, Australia, and CEPR. E-mail: yves.zenou@monash.edu

1 Introduction

The diffusion of social behavior—good or bad—is central to economic and social life.[1] Understanding the forces governing its spread is therefore a first-order question for welfare and policy design. Yet, the *strategic* nature of contagious behavior is context-dependent. For example, participation in Hong Kong's 1 July marches was found to be a strategic *substitute* in 2016 (Cantoni et al., 2019), but a strategic *complement* in 2017-18 (Bursztyn et al., 2021). How can the same collective action exhibit opposite strategic incentives?

This context-dependency in protest points to a more fundamental asymmetry in how *prosocial* versus *antisocial* behaviors spread. In a laboratory experiment Tsvetkova and Macy (2014) find that observing low prevalence of generosity raises willingness to give, but at high prevalence a "bystander effect" emerges: additional observed generosity reduces individuals' propensity to contribute. By contrast, Tsvetkova and Macy (2015a) find that observing low prevalence of antisocial behavior (taking from others) reduces willingness to harm relative to observing no information.

We propose a model in which strategic incentives are not intrinsic, but are shaped by agents' beliefs about the environment. Central to our analysis is the interaction between network density and an agent's *influence*—the discounted number of subsequent actions they expect their action to cause. Network density has two countervailing effects. First, more connections provide more potential observers of an agent's action (a "local effect" that amplifies influence). Second, a denser network creates more redundant paths for information to spread, reducing any single individual's pivotality over others' awareness (a "global effect" that diminishes influence). The tension between these local and global effects is the core mechanism driving our results. Our notion of influence extends Bénabou et al. (2020), whose model captures only the local effect. To formally analyze this, we model the strategic interaction as a game on a random graph with an arbitrary degree distribution. This allows us to characterize a phase transition between nonviral equilibria, where diffusion is contained to a negligible fraction of the population, and viral equilibria, where actions may spread to a nontrivial fraction.

This local-global mechanism is consistent with the observed empirical puzzles. A key insight is that individuals' attitudes towards influence depend on the nature of the act. For prosocial behaviors, spillovers are positive, so agents are *influence-seeking*. In sparse (subcritical) networks, the local effect dominates, influence is high, and actions are strategic complements. In dense

[1] For prosocial actions, evidence of contagion spans behaviors such as blood donations (Schröder et al., 2023), protest participation (Myers, 2000), political donations (Traag, 2016), and generosity (Tsvetkova and Macy, 2014). Similarly, antisocial behaviors exhibiting contagion include bullying (Fei et al., 2024), riots (Myers, 2011), gun violence (Fagan et al., 2007), information sharing among competitors (Stein, 2008; Boldrin and Levine, 2008; Pool et al., 2015), and even participation in Ponzi schemes (Rantala, 2019).

(supercritical) networks, the global effect dominates, influence is low, individuals free-ride, and behavior exhibits strategic substitutability—the "bystander effect." For antisocial behaviors, spillovers are negative, so agents are *influence-averse*. They are deterred by high influence in sparse networks, but are "licensed" by the diffusion of responsibility (low influence) in dense networks, which generates strategic complementarity.

This asymmetry matches the experimental evidence in Tsvetkova and Macy (2014, 2015b), where participants observed behavior prevalence but not network structure. We interpret the observation treatment as shifting beliefs about one's pivotality (expected influence), rather than structural density. In the prosocial setting (Tsvetkova and Macy, 2014), interaction occurred in large groups, and observed cumulative generosity increased helping when prevalence was low (high perceived pivotality) but reduced it when prevalence was high (a bystander effect). In the antisocial chain design (Tsvetkova and Macy, 2015b), exposure followed short linear paths, and observing *low* prevalence decreased harm relative to no information, with no evidence of an observation-induced surge. This "unbroken windows" deterrence aligns with our model's non-viral prediction: influence-averse agents are discouraged from antisocial actions at high perceived pivotality, and higher observed prevalence alone does not generate licensing along sparse, non-viral exposure paths.

Our model also provides a new lens for empirical work outside the laboratory: observing whether a given behavior is a strategic complement or substitute can, in turn, provide information on agents' collective belief about whether the effective network density ($\lambda\sigma$) is in a subcritical or supercritical state. While this presents an identification challenge—separating beliefs about network density λ from the equilibrium strategy σ—it offers a clear path to using observed strategic incentives to infer unobserved beliefs about the environment.

For example, while the Hong Kong protest experiments by Cantoni et al. (2019) and Bursztyn et al. (2021) do not directly vary the actual network structure, they can be interpreted as recalibrating agents' *beliefs* about the network density and, consequently, the expected diffusion of participation. In our model, strategic incentives are driven by beliefs about influence, which is determined by the perceived network density. Under this interpretation, we can reconcile the behavior in the Hong Kong protests with our local-global mechanism.

Cantoni et al. (2019) study the 2016 protest, an event where subjects held beliefs of high expected participation (an average posterior belief of $\approx$ 143,000). They find that informing individuals about high citywide turnout decreased their willingness to protest, that is, participation is a strategic substitute. We interpret this as behavior in a perceived supercritical network: when a protest is believed to be large (the network is dense), individuals perceive their own marginal influence as low, which activates the "bystander effect" they document.

By contrast, Bursztyn et al. (2021) study the smaller 2017–2018 protests (described as "modestly sized" at $\approx$ 50,000 participants). They find that incentivizing participation among peers increased participation—a strategic complement. We interpret this as behavior in a perceived subcritical network: when a protest is believed to be smaller (sparse), individuals perceive their potential to influence others as high, leading to complementarity. Thus, our model provides a unified framework that reconciles these opposite findings. The parameter that drives strategic incentives (λ in our model) is the perceived density, which is determined by beliefs about expected turnout. The documented shift in beliefs from high (143,000) to modest (50,000) provides evidence of a change in the perceived network state—from supercritical (generating substitutes) to subcritical (generating complements).

A key consequence of our equilibrium analysis is the stark difference in the emergence of viral prosocial versus antisocial behaviors. Our model predicts that prosocial behaviors emerge continuously as a function of network density. In contrast, antisocial behaviors exhibit a discontinuous emergence once network density crosses a critical threshold. This discontinuity is also present in other model parameters; hence, seemingly small changes in the environment—e.g., a small increase in the private benefit of the antisocial action—can lead to a sudden emergence of antisocial behavior. Moreover, conditional on becoming viral, antisocial cascades propagate more rapidly than prosocial ones.

Finally, we analyze the problem of a social planner who can choose the network's density to encourage prosocial outcomes and inhibit antisocial ones. We show that the optimal network density is just below the critical threshold for the emergence of viral antisocial behavior. Our framework also demonstrates that effective policies can target the discount factor—the decay of externalities with social distance.

1.1 Related Literature

Closest to us is the work of Sadler (2020, 2025). Sadler (2020) develops a general class of "diffusion games" with strategic complementarities between direct neighbors, and highlights the role of the giant component. We instead study global (indirect) externalities: an agent's utility depends on the total number of adopters in the network, regardless of distance, so adoption need not be complementary.[2] Sadler (2025) introduces a *seed multiplier* in multi-type configuration networks to study optimal seeding. Our notion of *influence* plays an analogous role at the individual level and lets us show how complementarity and substitutability differ around the viral threshold. Our emphasis on the asymmetry between prosocial versus antisocial diffusion is also new to this literature.

[2]In Sadler (2020), adoption externalities only arise when a person also adopts. Hence, the number of adopting neighbors when an individual does not adopt is an irrelevant strategic consideration.

Network games and diffusion. Our framework differs from the large literature on games on *fixed* networks, where links between players typically represent a direct strategic externality (see Jackson and Zenou, 2015; Bramoullé and Kranton, 2016; Jackson et al., 2017, for reviews). In our model, connections are for observation, while externalities are global. Many papers in this literature study games of strategic complements, often in the form of threshold models—where the incentive to take action 1 rises with the fraction of neighbors doing so (Granovetter, 1978; Schelling, 1978; Blume, 1993; Ellison, 1993; Morris, 2000; Brock and Durlauf, 2001; Leister et al., 2022; Langtry et al., 2024). Unlike these papers, the presence of global externalities in our model rules out any neighbor-threshold rule—decision turn on the expected *global* impact. Strategic substitutes have been considered in the form of privately provided public goods—a type of prosocial behavior (Bramoullé and Kranton, 2007; Allouch, 2015, 2017).[3] Work allowing both strategic complements and substitutes with continuous actions includes Ballester et al. (2006); Bramoullé et al. (2014). These papers are similar to ours, but by working with random graphs, we are able to leverage the regularities that emerge in the limit of a large population.

Our use of random graphs and phase transitions relates to work that applies these tools to diffusion in economics (e.g., Watts, 2002; Campbell, 2013; Akbarpour et al., 2023; Langtry, 2023; Campbell et al., 2024; Dasaratha, 2023).[4] Our binary-action environment is also connected to threshold models on random graphs (Jackson and Yariv, 2005; López-Pintado, 2006, 2008; Jackson and Yariv, 2007, 2011; Jackson and López-Pintado, 2013; Campbell et al., 2025; Langtry and Thornton, 2025). Our novel definition of *influence* and conceptual focus on the asymmetry between prosocial and antisocial behaviors is what separates us from this literature.

Collective action. The literature on collective action typically assumes either *strategic complementarity*, or *strategic substitutability*, but not both. For example, in the context of protests one view is that turnout lowers costs or raises the value of participation (e.g., Kuran, 1989; Chwe, 2000; Edmond, 2013), while another is that turnout creates incentives for free-riding (e.g., Olson, 1965; Palfrey and Rosenthal, 1984; Shadmehr and Bernhardt, 2011). Our model endogenizes which force dominates by tying it to network density via influence, showing how the same collective action can be a complement in sparse (subcritical) environments yet a substitute in dense (supercritical) ones. While most collective-action models abstract from networked diffusion (Chwe, 2000 is an exception), we show precisely how the network environment mediates strategic incentives.

The remainder of the paper proceeds as follows. Section 2 presents a special case of the model.

[3]Elliott and Golub (2019) study public goods in observational networks with global positive externalities; their main results are nonparametric and do not commit to substitutes or complements.

[4]For an overview, see Newman et al. (2001); Kleinberg (2007); Vega-Redondo (2007); Jackson (2008); Easley and Kleinberg (2010).

Section 3 characterizes influence. Section 4 provides the equilibrium analysis and welfare/policy implications. Section 5 offers generalizations, and Section 6 concludes.

2 The Model

2.1 Outline

In this section, we present a model of diffusion on an Erdös-Rényi random graph. We extend our model to graphs with an arbitrary degree distribution in Section 5.1, but the qualitative insights of our model remain the same as in this foundational case. Let $\{G(n,p)\}_{n\in\mathbb{N}}$ be a sequence of Erdös-Rényi (or *binomial*) random graphs, where $p = p(n) = \frac{\lambda}{n}$ for some $\lambda \geq 0$. We endow $\{G(n,p)\}_{n\in\mathbb{N}}$ with the structure of a game $\mathcal{G}^{(n)} = (G(n,p), \mathcal{A}, u)$ as follows. Each vertex $i \in \{1,\ldots,n\}$ represents an agent, and each agent faces a binary adoption decision $a_i \in \mathcal{A} = \{0,1\}$. If player i chooses action 1, we say that i is an *adopter*.[5]

The game $\mathcal{G}^{(n)}$ is played over $n+1$ time periods t. Denote by $a_i(t)$ the action player i at time t. At $t = 0$, every player has action $a_i(0) = 0$. At $t = 1$, nature makes three moves:

1. First, nature draws a graph $G = G(n,p)$ by including each possible edge independently with probability p.

2. Second, nature chooses a "seed" uniformly at random and,

3. Third, the seed adopts the prosocial or antisocial behavior $\beta \in \{+,-\}$.[6]

For all $t \geq 2$, neighbors of an adopter in the previous period are *exposed* and make a once and for all decision to adopt or not. Hence, the adoption decision occurs when an agent is first exposed to the behavior, but not again if they are exposed by a different neighbor in the future. As such, we often write a_i rather than $a_i(t)$ for the action i would take upon exposure.

2.2 Payoffs

At the end of the game, an agent i's payoff is a function of their own action a_i and the actions a_{-i} of the other agents in the network. The utility that agent i receives places a larger weight on the actions of those close to i in the network. Let ℓ_{ij} denote the length of the shortest path from i to j (or $+\infty$ if no such path exists) and let $\delta \in (0,1)$ be a common decay factor.[7] The

[5]This model along with its more general analogue are similar to that of a *single-type diffusion game* (Sadler, 2020). Sadler's framework offers a convenient way to analyze the role of influence in the diffusion of behaviors.

[6]One can think of the seed as a non-strategic individual who exogenously wishes to engage in the behavior.

[7]The introduction of a discount factor is motivated by many real-world settings where spillovers decrease with geographical distance. For instance, the negative effects of violence or the mobilizing effect of a protest

decay factor δ captures how local-versus-global spillovers are: higher values of δ correspond to spillovers from actions that are more far-reaching.[8] While our results are robust to alternative utility specifications (as we discuss in Appendix F.1.1), for our main analysis we specify agent i's utility as:

$$u_\beta(a_i, a_{-i}) = \begin{cases} \overbrace{(v-c)a_i}^{\text{intrinsic cost}} + \overbrace{v\sum_{j\neq i}\delta^{\ell_{ij}}a_j}^{\text{externality}}, & \text{if } \beta = + \\ (c-v)a_i - v\sum_{j\neq i}\delta^{\ell_{ij}}a_j, & \text{if } \beta = -, \end{cases} \tag{1}$$

where v captures the common value to all agents from each one of them who adopts and $c > v$ is the private cost (benefit) to an individual from adopting a prosocial (antisocial) behavior β. One can also interpret v as capturing a first-order approximation of the marginal increase in "better" expected outcomes when adoption affects the realization of probabilistic events.[9] An agent's utility can be written as a combination of an intrinsic payoff from the action and an externality from the actions of others. We assume that the intrinsic payoff has the opposite sign to the externality; hence, the intrinsic payoff is negative (positive) in the case of a prosocial (antisocial) behavior.

2.3 Strategies

We make a number of simplifying assumptions about the information set available to a player at the time when they are first exposed, and we discuss how one would relax the most crucial of these assumptions in Appendix F.1.3. First, agents do not know their degree but believe correctly that it is distributed according to a binomial distribution with parameter $p = \lambda/n$. Second, they know that at least one of their neighbors has chosen to adopt but do *not* know which ones. Third, they do not know exactly how much time has passed since the seed adopted and therefore the calendar time t at which they act. These guarantee that each agent acts at a single information set which we denote by h_i ($i \in \{1, \ldots, n\}$), all agents are symmetrically informed at the moment each acts, and each agent i's strategy may be characterized by a (potentially mixed) strategy over a single action $\sigma_i = \mathbb{P}(a_i = 1) \in [0, 1]$. The strategy σ_i represents the probability that agent i takes the action if they are exposed. We confine our attention to *symmetric equilibria*, that is, where all players play the same strategy $\sigma_i^* = \sigma^*$ for all i.

are strongest on those geographically or socially proximate. Extensive evidence confirms that these and similar behaviors exhibit contagion, e.g. Myers (2000, 2011); Fagan et al. (2007); Keizer et al. (2008).

[8] A special case arises when $\delta = 1$, which we discuss in Online Appendix B.1.

[9] We thank an anonymous reviewer for highlighting this interpretation.

2.4 Influence

For any agent i (adopter or otherwise), we can consider the random number of other agents who adopt if i were to adopt, but do not have the opportunity to adopt otherwise (i.e., the information set where those agents have the opportunity to take an action is not reached). This is what we refer to as an agent i's influence: the number of agents that agent i *causes* to adopt by choosing $a_i = 1$, discounted by their distance from i. Consider the example in Figures 1 and 2, which shows the connected component containing the seed in a potential adopter network. The seed is colored blue, agents colored black have zero influence, and agents colored green have nonzero influence over the agents colored red that pass through the green node on all paths back to the seed.[10] In this example, if any of the green colored nodes chose to not adopt then the red nodes over which they have influence would become disconnected from the component containing the seed. In total, the reduction in the number of individuals adopting would be equal to the green node plus the red nodes that become disconnected from the seed when the green node does not adopt. Furthermore, to illustrate how fewer network links may increase the influence of agents, consider how influence changes after removing the links associated with a single vertex (as we do in Figure 2). In Figure 1, the first green node to the right of the seed has influence over only a single agent. This is because the other node adjacent to it is on an alternate path between the green node and the seed. But once we remove the second path, we increase the influence of this agent from 1 to 3. This is an example of what we earlier referred to as the global effect of density: decreasing the connectivity of a graph can *increase* the influence of agents within that graph. This happens precisely because removing links may disconnect loops thereby increasing the influence of agents on those loops. This idea is formalized in Proposition 3.

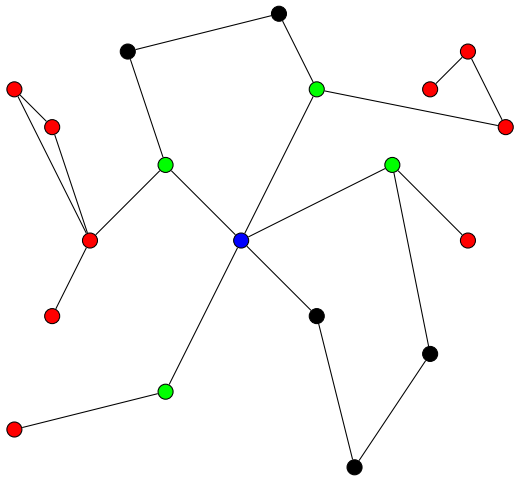

Figure 1: Influence

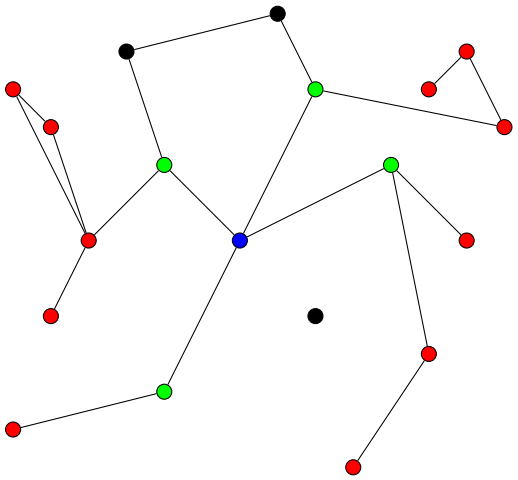

Figure 2: Removing a node

[10] Some of the red nodes also have nonzero influence, but for illustration we are focusing on the neighbors of the seed.

In expectation, the number of agents over which i has influence depends on the strategy σ which i expects others to play, and on the structure of the network. We write $\mathbb{E}_{\sigma}^{(n)}$ for an expectation taken in $\mathcal{G}^{(n)}$ assuming players follow the strategy σ, and we define $\mathbb{E}_{\sigma} = \lim_{n\to\infty} \mathbb{E}_{\sigma}^{(n)}$. Similarly, define $\mathbb{E}_{\mathcal{D}_n}^{(n)}$ to be an expectation taken in $\mathcal{G}^{(n)}$ over the degree distribution $\mathcal{D}_n = \text{Bin}(n, p)$, and define $\mathbb{E}_{\mathcal{D}}$ as its limit, where $\mathcal{D} = \text{Po}(\lambda)$—the Poisson distribution. We then define $\mathcal{I}_{\lambda}^{(n)}(\sigma)$ and $\mathcal{I}_{\lambda}(\sigma)$ (read "the expected influence under σ") by

$$\mathcal{I}_{\lambda}^{(n)}(\sigma) \equiv \mathbb{E}_{\mathcal{D}_n}^{(n)}\left[\mathbb{E}_{\sigma}^{(n)}[\sum_{j\neq i}\delta^{\ell_{ij}}a_j|a_i=1,h_i] - \mathbb{E}_{\sigma}^{(n)}[\sum_{j\neq i}\delta^{\ell_{ij}}a_j|a_i=0,h_i]\right], \quad \text{and} \tag{2}$$

$$\mathcal{I}_{\lambda}(\sigma) \equiv \lim_{n\to\infty}\mathcal{I}_{\lambda}^{(n)}(\sigma) = \mathbb{E}_{\mathcal{D}}\left[\mathbb{E}_{\sigma}[\sum_{j\neq i}\delta^{\ell_{ij}}a_j|a_i=1,h_i] - \mathbb{E}_{\sigma}[\sum_{j\neq i}\delta^{\ell_{ij}}a_j|a_i=0,h_i]\right]. \tag{3}$$

Suppose all players are playing strategy σ. We normalize the utility of not adopting to 0 and we denote the expected change in utility from adopting by $u_{\beta}^{(n)}(\sigma) \equiv \mathbb{E}_{\sigma}^{(n)}[u_{\beta}(1,a_{-i}) - u_{\beta}(0,a_{-i})]$, and $u_{\beta}(\sigma)$ in the limit as $n\to\infty$.[11] Now, using the notation in (3) and the utility function (1), we can write the expected difference in utility to an agent from adopting $a_i = 1$ versus not adopting $a_i = 0$ at the information set where they have the opportunity to act—in particular, when they hold imperfect information about the graph—as:

$$\mathbb{E}_{\mathcal{D}}^{(n)}[u_{\beta}^{(n)}(\sigma)] = \begin{cases} v - c + v\mathcal{I}_{\lambda}^{(n)}(\sigma), & \beta = + \\ c - v - v\mathcal{I}_{\lambda}^{(n)}(\sigma), & \beta = -. \end{cases} \tag{4}$$

Observe that our definition of (expected) influence $\mathcal{I}_{\lambda}(\sigma)$ captures the *causal effect* of adoption on others' adoption—that is, the expected influence of agent i is the (random, discounted) number of other agents who adopt if i adopts but would not adopt otherwise.

3 Analysis

3.1 Existence of Equilibrium

A symmetric perfect Bayesian equilibrium in the game $\mathcal{G}^{(n)}$ is a strategy $\sigma^* \in [0,1]$ satisfying:

$$\begin{aligned} \sigma^* = 1 &\implies \mathbb{E}_{\mathcal{D}_n}^{(n)}[u_{\beta}^{(n)}(\sigma^*)] \geq 0 \\ \sigma^* \in (0,1) &\implies \mathbb{E}_{\mathcal{D}_n}^{(n)}[u_{\beta}^{(n)}(\sigma^*)] = 0 \\ \sigma^* = 0 &\implies \mathbb{E}_{\mathcal{D}_n}^{(n)}[u_{\beta}^{(n)}(\sigma^*)] \leq 0, \end{aligned}$$

[11] The expectation here is taken over a given fixed graph so that influence is well defined.

and beliefs over nodes in the extensive form where an agent is exposed under the strategy σ^* that are consistent with nature's choice of a random graph with degree distribution $\mathcal{D}_n$. Beliefs are only payoff relevant for determining the expected influence of each agent and, given the symmetry of our game, this will be identical for each agent. It will be convenient to avoid specifying the beliefs themselves and rather specify the calculation of expected influence as a function of σ^* and $\mathcal{D}_n$ with the understanding this is through a set of beliefs formed via Bayes rule for each agent. The calculation of influence in the finite n case is discussed in Online Appendix B.

Proposition 1. *There exists a symmetric perfect Bayesian equilibrium in the game $\mathcal{G}^{(n)}$.*

We now focus on symmetric equilibria in the limit as $n \to \infty$.

Definition 1. *Let σ' and σ be strategies. We say that σ' is a* limit best-reply *to σ, if*

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{E}^{(n)}_{\mathcal{D}_n}[u^{(n)}_\beta(\sigma)] \geq 0 \ \textit{whenever } \sigma' = 1,\\
&\lim_{n\to\infty} \mathbb{E}^{(n)}_{\mathcal{D}_n}[u^{(n)}_\beta(\sigma)] = 0 \ \textit{whenever } \sigma' \in (0,1), \ \textit{and}\\
&\lim_{n\to\infty} \mathbb{E}^{(n)}_{\mathcal{D}_n}[u^{(n)}_\beta(\sigma)] \leq 0 \ \textit{whenever } \sigma' = 0.
\end{aligned}$$

The strategy σ' is a limit equilibrium *if σ' is a limit best-reply to itself.*

A limit equilibrium is an epsilon-equilibrium for any $\epsilon > 0$ and all sufficiently large n. Furthermore, we show in Online Appendix D that generically any limit equilibrium can be found as a limit of a sequence of symmetric perfect Bayesian equilibria as $n \to \infty$ in the finite player game $\mathcal{G}^{(n)}$.

3.2 Potential Adopter Network

In a symmetric equilibrium, all potential adopters are determined by the realization of their strategy σ^*. Following the realization of the graph $G(n,p)$, one can imagine realizing n independent Bernoulli random variables with parameter σ^*, one for every individual in the network. The realization of these variables establishes the subgraph of G containing all "potential adopters", which we call the *potential adopter network*.[12] The location of the seed determines who actually adopts. The number of adopters is determined by the size of the component containing the seed in the potential adopter network. Hence, central to our analysis will be the component structure of the potential adopter network, which itself is a random graph with independent edge probabilities σp.

[12]This is commonly known as the *site-percolation network* in random-graph theory.

At this point, we need to introduce the notion of a *giant component* in a random graph. A classic result of Erdős and Rényi (1959) is the phase transition in the size of the largest component in $G(n, \lambda/n)$ around $\lambda = 1$. The result states that as $n \to \infty$,

1. *Sub-critical region* $\lambda < 1$, with high probability the largest component of $G(n, \lambda/n)$ has a size which is at most $\underline{c}\log(n)$ for some constant $\underline{c}$.

2. *Super-critical region* $\lambda > 1$, with high probability the largest component of $G(n, \lambda/n)$ contains a fraction $\overline{c}n$ of the vertices for some constant $\overline{c}$. Moreover, the second largest component has a size which is at most $\underline{c}\log(n)$ for some constant $\underline{c}$.

A component which contains a constant fraction of vertices is called a giant component. The above result states that the binomial random graph contains a giant component (with high probability) if and only if $\lambda > 1$. Moreover, when it exists the giant component is unique. Similarly, the potential adopter network contains a giant component if and only if $\lambda\sigma > 1$.

Our analysis is concerned with the limit of large networks i.e., where $n \to \infty$. In this limit, it is well known that the Erdös-Rényi random graph $G(n, \lambda/n)$ converges locally in probability to a Poisson branching process with mean offspring λ (Van der Hofstad, 2023b, see, e.g., Thm 2.18).[13] In the case where $\lambda < 1$, the process stops with probability 1; however, in the case where $\lambda > 1$, there is a non-zero probability the process never goes extinct. Hence when $\lambda > 1$, there are two potential outcomes of following a randomly chosen link to one of its ends in the Erdös-Rényi random graph and finding all subsequent connected nodes. The first outcome is that the link leads to a component with a finite expected number of people. This occurs with probability ρ_σ (called the *extinction* probability). The second outcome is that the link leads to an infinite path (as $n \to \infty$) that connects to the giant component with probability. This outcome occurs with probability $1 - \rho_\sigma$ (the *survival* probability).[14] An individual with d connections has d independent realizations of these events (following a randomly chosen link to one of its ends). If any one of these connections leads to the giant component, then the individual is also part of it. Hence, the probability this person is in the giant component is $1 - \rho_\sigma^d$. Moreover, conditional on this event (at least one connection leads to the giant component) the distribution of extinction events over the $d-1$ remaining connections is given by $\text{Bin}(d-1, \rho_\sigma)$. In our potential adopter network it is well known that the extinction probability ρ_σ is given by the smallest positive solution to:

$$\rho_\sigma = e^{-\lambda\sigma(1-\rho_\sigma)}. \tag{5}$$

Finally, the expected number of people reached by following a link and conditioning on extinction

[13] We describe Poisson branching processes in more detail in Online Appendix A.1.

[14] See Van der Hofstad 2023b, Thm 2.28 and its application to Erdös-Rényi random graphs Thm 2.34. For an arbitrary degree distribution in the configuration model see also Thm 4.9 and the discussion thereafter.

is $\frac{1}{1-\lambda\sigma\rho_\sigma}$.[15].

3.3 Inference and Calculation of Influence

The total number of adopters in an equilibrium corresponds to the number of people in the same component as the seed in the potential adopter network. The expected influence in equilibrium corresponds to the expected discounted number of individuals that are connected-to versus disconnected-from the seed's component when an individual does versus does not adopt. In our model, exposure confers information about an individual's component, in particular, the agent is in the same component as the seed. This is informative of an agent's expected influence. In the finite network case this inference is complicated and intractable. However, a consequence of local convergence of the random graph to a branching process in large networks is that the expected influence converges to a particularly simple form.

To compute the expected influence of an arbitrary agent i, we partition the other agents in i's component into those in the forward components found through each of an i's links.[16] This construction means that agent i is essential for constructing a path between agents in different elements of the forward component partition. In the limit of a large random graph, these forward components (and an individual's component) are random objects and, as discussed earlier, characterized by a Poisson branching process.

First, consider an equilibrium in which $\lambda\sigma^* < 1$ (the sub-critical region). Each forward component of an individual is finite with probability 1 and independent of each other. Hence, the probability an individual with d friends is in the same component as the seed is in proportion to their connectivity. Conditional on exposure, the updated probability of having d friends is $\Pr(d|h_i, \lambda\sigma^* < 1) = \frac{dp_d}{\sum dp_d}$ where $\{p_d\}_{d\geq 0}$ are the prior probabilities. Furthermore, when $\{p_d\}_{d\geq 0}$ is Poisson then it is well known that $\mathbb{E}[d|h_i, \lambda\sigma^* < 1] = \lambda\sigma^* + 1$. Upon exposure an individual forms an expectation over the forward components found via following each one of its links excluding the link through which it was exposed. These components are independent realizations of a Poisson branching process with mean offspring $\lambda\sigma^*$ where the extinction probability ρ_σ is equal to 1. The expected influence of this individual is equal to their discounted expected number of neighbors $\delta\mathbb{E}[d-1|h_i, \lambda\sigma^* < 1] = \delta\lambda\sigma^*$ times the discounted expected forward component size given by $(1 + \delta\lambda\sigma^* + (\delta\lambda\sigma^*)^2 + \dots) = \frac{1}{1-\delta\lambda\sigma^*}$.[17] and hence expected influence is:

[15] This follows from Van der Hofstad 2023a, Thm 3.5 and Thm 3.15.

[16] Of course, some links may result in the same forward component in which case only a single copy is maintained in the partition; hence, there may be fewer forward components created than an agent has links.

[17] Exposure is informative of the size of component found via following the exposure link but, in the sub-critical region, these beliefs do not affect the calculation of expected influence $\mathcal{I}_\lambda(\sigma)$.

$$\mathcal{I}_\lambda(\sigma^*) = \frac{\lambda\sigma^*\delta}{1-\lambda\sigma^*\delta} \tag{6}$$

when $\lambda\sigma^* < 1$.

Second, consider an equilibrium where $\lambda\sigma^* > 1$ (the super-critical region). In the super-critical region, the network of potential adopters contains a unique giant component and upon exposure an agent believes, almost surely, that they are in the giant component.[18] In this case an agent updates their beliefs to reflect that their number of links follows the distribution of neighbors for individuals in the giant component, which is given by $\Pr(d = k|h_i, \lambda\sigma^* > 1) = \frac{p_k(1-\rho_{\sigma^*}^k)}{1-\rho_{\sigma^*}}$.[19] Conditional on this event (at least one connection leads to the giant component) the remaining $d-1$ links are either connected to finite forward components (this occurs with probability ρ_{σ^*}) or connected to a forward component containing everyone in the giant component other than the people contained in i's finite forward components (this occurs with probability $1-\rho_{\sigma^*}$). We call this forward component the giant forward component. Therefore, i's partition of the giant component consists of the giant forward component and the number of successes of $d-1$ independent realizations of Bernoulli(ρ_{σ^*}) random variables, one for each potential finite forward component. The probability that the seed is in any one of these forward components is of course dominated ($\Pr \to 1$) by the giant forward component and vanishing ($\Pr \to 0$) for the finite components.[20] The calculation of expected influence when the seed is in the giant forward component is simply the expected discounted number of people in the remaining finite forward components of the partition, this is given by $\frac{\delta\lambda\sigma^*\rho_{\sigma^*}}{1-\delta\lambda\sigma^*\rho_{\sigma^*}} - \delta\rho_{\sigma^*}\frac{1-\lambda\sigma^*\rho_{\sigma^*}}{1-\delta\lambda\sigma^*\rho_{\sigma^*}}$, where the first term is nearly identical to (6) and the second term is a correction that occurs due to agents' updating about their expected degree given that they are in the giant component.

We summarize our calculation of influence for both cases with in following lemma.

Lemma 2. *The expected influence function $\mathcal{I}_\lambda\colon [0,1] \to \mathbb{R}$ is given by*

$$\mathcal{I}_\lambda(\sigma) = \begin{cases} \frac{\delta\lambda\sigma}{1-\delta\lambda\sigma} & \text{if } \lambda\sigma < 1, \\ \frac{\delta\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} - \delta\rho_\sigma\frac{1-\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} & \text{if } \lambda\sigma > 1 \end{cases} \tag{7}$$

where ρ_σ is as in (5). In particular, $\mathcal{I}_\lambda$ is continuous.[21]

[18] The giant component contains a positive fraction of the population as $n \to \infty$ whereas the expected size of all other components $\to 0$. Hence, the probability of being in the same component as the seed is vanishing in the limit $n \to \infty$ for all components other than the giant component.

[19] This is termed "viral" inference in Sadler (2020).

[20] To be precise, we show in the proof of Lemma 11 (found in Appendix 7.10 of the paper) that the event where the seed is in a finite forward component contributes negligibly to influence.

[21] In fact, it is easily shown that $\mathcal{I}_\lambda$ is an analytic function at every point other than where $\lambda\sigma = 1$.

We further characterize how influence changes with agents' adoption probability σ and network density λ in the following proposition.

Proposition 3. *Let $\sigma \in [0, 1]$ be any strategy.*

1. *If σ is nonviral, then $\mathcal{I}_\lambda(\sigma)$ is strictly increasing in σ and strictly increasing in λ.*

2. *If σ is viral, then $\mathcal{I}_\lambda(\sigma)$ is strictly decreasing in σ and strictly decreasing in λ.*

Figure 3 captures the main idea behind Proposition 3, which is a key comparative static result for our analysis.

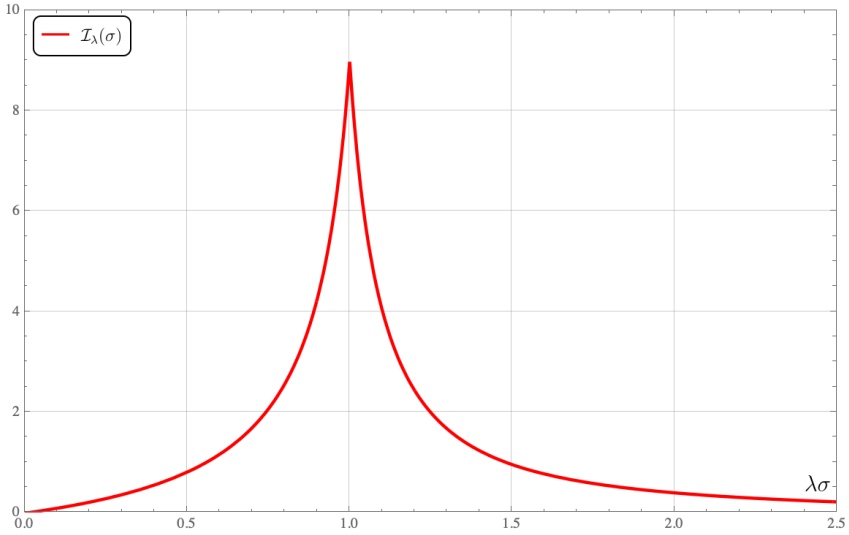

Figure 3: Expected influence is increasing when the network is sparse ($\lambda\sigma < 1$) and decreasing when the network is dense ($\lambda\sigma > 1$).

We see that in the sub-critical region influence reaches its maximum of $\delta/(1-\delta)$ at the point where the giant component emerges ($\lambda\sigma = 1$). Beyond this point it is decreasing as the giant component grows. This illustrates one of the consequences of Proposition 3, and a key contribution of our model: that adoption is a strategic complement when the graph is sparse, and a strategic substitute when it is dense.

This non-monotonic relationship between influence and density, illustrated in Figure 3, is the key to our paper's central empirical applications. As we discussed in the Introduction, this mechanism resolves the apparent paradox between Cantoni et al. (2019) and Bursztyn et al. (2021). Their conflicting findings (complements vs. substitutes) align with our model's prediction of a strategic flip when perceived network density—proxied by beliefs about turnout—crosses from a subcritical to a supercritical state.[22]

[22]Another interpretation, also consistent with our framework, is that the underlying network parameter λ was fixed, but agents' beliefs about the equilibrium strategy σ changed. In 2016, agents may have expected

Finally, we note that Bursztyn et al. (2021) offer a related interpretation for these differing results. They posit that strategic incentives depend on the type of tie: observing local peers (friends) triggers social utility and coordination effects, generating complements, whereas observing global population turnout (weak ties) triggers beliefs about public good provision or government crackdowns, leading to substitutes. Our model provides a theoretical mechanism that endogenizes this distinction. When perceived turnout is low, diffusion is contained and individual influence is high. In this state, the "local" social utility and coordination effects described by Bursztyn et al. (2021) are strategically pivotal, consistent with complementarity. Conversely, when perceived turnout is high (as in Cantoni et al. (2019)), influence is negligible and responsibility is diffused globally. In this state, the "global" public good aspect of the action dominates, and our model correctly predicts the emergence of strategic substitutability. Our framework thus provides a microfoundation for why the strategic focus shifts from local to global as beliefs about the scale of diffusion change.

4 Results

We analyze the behavior of expected influence as a function of the network density and proceed to fully characterize symmetric equilibria for prosocial and antisocial behaviors. Furthermore, we will be interested in the structure of the potential adopter network induced by an equilibrium strategy. We call a strategy *viral* if it induces a giant component in the potential adopter network, and *nonviral* otherwise. For expositional simplicity, we will focus throughout on the case where δ is sufficiently large, in particular, $\delta > 1 - \frac{v}{c}$.[23] In our setting, there are four types of symmetric equilibria:

1. No-adoption ($\sigma^* = 0$)
2. Viral or Nonviral full-adoption ($\sigma^* = 1$)[24]
3. Nonviral mixed ($\sigma^* \in (0,1),\ \lambda\sigma^* < 1$)
4. Viral mixed ($\sigma^* \in (0,1),\ \lambda\sigma^* > 1$)

All four types of equilibria can arise depending on the structure of the underlying graph (as

coordination on a high-participation (supercritical, $\lambda\sigma > 1$) equilibrium, while in 2017-18, they expected a low-participation (subcritical, $\lambda\sigma < 1$) equilibrium. Both interpretations rely on our central mechanism: the strategic flip occurs when the effective density $\lambda\sigma$ crosses the critical threshold. We find the interpretation in the main text, which maps beliefs about turnout directly to the environmental parameter λ, to be the more parsimonious one.

[23] When $\delta < 1 - \frac{v}{c}$, the utility from taking the action is always negative (positive) for prosocial (antisocial) behaviors and so the analysis is trivial— there is a unique equilibrium.

[24] When $\sigma = 1$, the potential adopter network is all of $G(n,p)$. Hence the full-adoption equilibrium is viral if and only if $\lambda > 1$.

determined by λ). Clearly there can not be a viral equilibrium when the underlying graph does not contain a giant component ($\lambda < 1$). We will show that there exists a critical value of $\lambda_\beta^{\text{crit}}$ that separates regions where viral equilibria exist ($\lambda > \lambda_\beta^{crit}$) from those where they do not ($\lambda < \lambda_\beta^{crit}$). There is a stark contrast between the transition from nonviral to viral equilibria around this threshold for prosocial behaviors when compared to antisocial ones. To this end, define $C_{1,\beta}(\lambda)$ as the fraction of agents (in the limit as $n \to \infty$) in the largest component of the potential adopter network, in the largest equilibrium of the game with prosocial behaviors ($\beta = +$) or antisocial behaviors ($\beta = -$). We will be interested in whether the fraction of agents who adopt in the largest equilibrium goes to zero or a limit bounded away from zero as the density approaches the critical density from above i.e. $\lim_{\lambda \downarrow \lambda_\beta^{crit}} C_{1,\beta}(\lambda)$.

Our focus on the largest equilibrium $\overline{\sigma}$ reflects the fact that nonviral equilibria are economically insignificant. In the case of prosocial behaviors, there turns out to be a unique viral equilibrium (whenever one exists). For antisocial behaviors, there are multiple viral equilibria, but only the largest is stable. To be precise, the largest equilibrium is always a generically an asymptotically stable point of the mean dynamic for the best response(see, e.g., Sandholm, 2010, Section 4.2.1). For example, under prosocial behaviors, for any $\epsilon > 0$ if $\sigma = \overline{\sigma} + \epsilon$ then $BR(\sigma) < \sigma$, and if $\sigma = \overline{\sigma} - \epsilon$ then $BR(\sigma) > \sigma$. The same best response property holds for the largest equilibrium in antisocial behaviors, but not for the second largest (viral-mixed) equilibrium.

4.1 Prosocial Behaviors

To begin, we focus on the case of prosocial behaviors ($\beta = +$). When other players play according to σ, expected utility from adoption is given by $\mathbb{E}_{\mathcal{D}}[u_+(\sigma)] = (v - c) + v\mathcal{I}_\lambda(\sigma)$. Hence, the best response correspondence BR_+ takes a particular simple form that depends on a threshold level of influence:

$$BR_+(\sigma) = \begin{cases} 1, & \mathcal{I}_\lambda(\sigma) > \frac{c}{v} - 1 \\ [0,1], & \mathcal{I}_\lambda(\sigma) = \frac{c}{v} - 1 \\ 0, & \mathcal{I}_\lambda(\sigma) < \frac{c}{v} - 1. \end{cases} \tag{8}$$

We characterize the structure of equilibria as a function of λ. We label each case in the proposition below with the "type" of the largest equilibrium—no-adoption, full-adoption, nonviral mixed, or viral mixed. It is helpful to first define two thresholds, $\underline{\lambda} \in (0,1)$ and $\overline{\lambda} \in (1,\infty)$ by

$$\underline{\lambda} = \frac{c - v}{\delta c}, \quad \overline{\lambda}\delta\rho_1 = \frac{c - v(1 - \delta\rho_1)}{c + v\rho_1}. \tag{9}$$

These threshold mark critical levels of network density (λ): $\underline{\lambda}$ is the minimum density required for equilibria beyond zero adoption to emerge (i.e., where influence can potentially cross the

best-response threshold), and $\overline{\lambda}$ is the maximum density for the full-adoption equilibrium to exist in the viral regime ($\lambda > 1$), beyond which only partial/mixed adoption is sustained at high densities due to the global effect of density on influence.

Proposition 4 (Characterization of equilibria for prosocial behaviors)**.** *For prosocial behaviors, the equilibrium structure depends on network density λ as follows:*

1. *(No-adoption) If $\lambda < \underline{\lambda}$, the unique equilibrium is $\sigma = 0$.*
2. *(Nonviral full-adoption) If $\underline{\lambda} \leq \lambda < 1$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium $\lambda\sigma = \underline{\lambda}$, and a nonviral full-adoption equilibrium $\sigma = 1$.*
3. *(Viral full-adoption) If $1 < \lambda \leq \overline{\lambda}$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium defined by $\lambda\sigma = \underline{\lambda}$ and a viral full-adoption equilibrium $\sigma = 1$.*
4. *(Viral mixed) If $\lambda > \overline{\lambda}$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium $\lambda\sigma = \underline{\lambda}$ and a viral mixed equilibrium defined by $\lambda\sigma = \overline{\lambda}$.*

The proposition can be understood by the relationship between (1) expected influence, (2) the threshold $\frac{c}{v} - 1$ in the best response and (3) the strategic complementarity/substitutability of adoption in the sub/super critical regions. In sparse networks ($\lambda < \underline{\lambda}$) or at low adoption levels ($\sigma < \underline{\lambda}/\lambda$), expected influence is below the best-response threshold $\frac{c}{v} - 1$, making zero adoption the best response. This guarantees that a no-adoption equilibrium exists for any λ, and this equilibrium is unique in sufficiently sparse networks (Part 1 of Proposition 4). In the sub-critical region ($\lambda < 1$), adoption is a strategic complement. As density λ increases, so does expected influence. For sufficiently high λ ($\lambda > \underline{\lambda}$), expected influence crosses the $\frac{c}{v} - 1$ threshold, enabling both the full-adoption and a nonviral mixed equilibrium alongside the no-adoption equilibrium (Part 2). Increasing density beyond $\lambda = 1$ enters the super-critical region. The three equilibria from Part 2 persist, but the full-adoption equilibrium becomes viral due to the emergence of a giant component in the network (Part 3). However, in the super-critical region, expected influence decreases with both adoption and density (as per Proposition 3). For sufficiently high λ ($\lambda > \overline{\lambda}$), influence at full-adoption falls *below* the $\frac{c}{v} - 1$ threshold, making full-adoption no longer a best response. Instead, the highest equilibrium becomes a viral mixed strategy, where the mixing level σ adjusts such that the resulting expected influence precisely matches the $\frac{c}{v} - 1$ threshold, sustaining the viral mix as a best response (Part 4).

A consequence of this equilibrium structure is that one of two forces will always limit the diffusion of a prosocial behavior. First, full adoption is only a viral equilibrium if the network is not *too* dense ($\lambda < \overline{\lambda}$); hence, even if everyone is prepared to adopt, viral diffusion is structurally limited by the size of the giant component. Second, in sufficiently dense networks ($\lambda > \overline{\lambda}$), diffusion

is limited by strategic considerations. The abundance of alternative paths reduces individual influence, creating a strong incentive to free-ride. This dynamic mirrors the "bystander effect" described by Tsvetkova and Macy (2014) and the strategic substitutability in Cantoni et al. (2019): when generosity appears common, the perceived marginal benefit of one's own action diminishes. Crucially, our model captures this mechanism without strictly requiring agents to know the network topology—they need only infer lower pivotality from higher observed prevalence. It is for this reason that our model serves as a natural framework for interpreting these empirical findings: it formally maps beliefs about others' behavior to the underlying network density, thereby rationalizing the shift in strategic incentives. Consequently, full adoption is no longer a best response in these dense regimes, limiting the connectedness of the potential adopter network.

We depict the relationship between influence and the best response in Figure 4. For exposition, we fix $\frac{c}{v} = 2$ and $\delta = .8$ in the figure. The top part of the figure shows expected influence in red and the threshold $\frac{c}{v} - 1$ in blue, around which the best response changes. The bottom part of the figure shows the best response correspondence in red, and the 45-degree line in black—the intersections of the red and black lines are equilibria. We fix $\lambda = 2.5$ so that the figure shows the entire range of possibilities for influence as a function of σ, the vertical orange line is the critical threshold where $\lambda\sigma = 1$, i.e. $\sigma = 0.4$. The maximum value of expected influence is $\delta/(1-\delta) = 4$ which occurs at this critical threshold. We also have $\underline{\lambda} = 0.625$, and $\overline{\lambda} \approx 1.37$, and these can be obtained by looking at the values of σ for which the expected influence (red) crosses the threshold $\frac{c}{v} - 1$ (blue). For example, the red and blue lines first intersect at $\sigma = 0.25$, so $\lambda\sigma = \underline{\lambda}$ implies that $\underline{\lambda} = 2.5 \times 0.25 = 0.625$, as claimed. Finally, because influence is small on the tails and large in the middle, best responses are 0 for small or large σ, and 1 for intermediate values.

To conclude this section, we translate Parts 2 and 3 of Proposition 4 into a statement about the size of diffusion. To unify our notation, we write $\lambda_\beta^{\text{crit}}$, for the critical threshold around which viral equilibria emerge (for $\beta \in \{+, -\}$). With this notation, $\lambda_+^{\text{crit}} = 1$, since viral equilibria in prosocial behaviors emerge alongside the emergence of the giant component. Finally, recall that we use $C_{1,\beta}(\lambda)$ to denote the fraction of agents in the largest component of the potential adopter network for $\beta \in \{+, -\}$ in the largest equilibrium. Then we have the following result.

Corollary 5 (Prosocial behaviors emerge continuously)**.** *The adoption of prosocial behaviors in the largest equilibrium is a continuous function of the network density λ, namely*

$$\lim_{\lambda \uparrow \lambda_+^{crit}} C_{1,+}(\lambda) = 0 = \lim_{\lambda \downarrow \lambda_+^{crit}} C_{1,+}(\lambda).$$

Figure 4: Expected influence and the best response correspondence for prosocial behaviors. Agents adopt only when expected influence is sufficiently large. Everyone adopts around the critical threshold where $\lambda\sigma = 1$. Maximal influence depends on λ, i.e. if $\lambda < \underline{\lambda}$ then expected influence is always below the blue line $\underline{\lambda}$ and no one ever adopts.

To understand Corollary 5, notice that Proposition 4 and Corollary 5 implies that around the phase transition in the underlying graph (i.e., $\lambda \in (1-\epsilon, 1+\epsilon)$), there is a parallel phase transition in equilibrium behavior: for $\lambda < 1$, all equilibria are necessarily nonviral, while for $\lambda > 1$ there exists an equilibrium which induces a giant component of potential adopters. Corollary 5 says that this transition admits a *continuous* change in the size of the largest component in the potential adopter network with respect to λ. We now turn to the case of the diffusion of antisocial behaviors, and demonstrate a stark contrast.

4.2 Antisocial Behaviors

We now consider the case of antisocial behaviors ($\beta = -$). Recall that when other players play according to σ, expected utility from adoption is given by $\mathbb{E}_{\mathcal{D}}[u_-(\sigma)] = -(v-c) - v\mathcal{I}_\lambda(\sigma)$. The best response correspondence again takes on a simple form characterized by a threshold level of

influence:

$$BR_{-}(\sigma) = \begin{cases} 1, & \mathcal{I}_\lambda(\sigma) < \frac{c}{v} - 1 \\ [0,1], & \mathcal{I}_\lambda(\sigma) = \frac{c}{v} - 1 \\ 0, & \mathcal{I}_\lambda(\sigma) > \frac{c}{v} - 1. \end{cases} \tag{10}$$

Note that although the utility from adopting an antisocial behavior is simply the negative of the utility from a prosocial behavior, and this "reverses" the best responses, we will see that there are meaningful differences in the structure of equilibria between the two cases. We now describe the structure of equilibria for antisocial behaviors, utilizing the thresholds $\underline{\lambda}$ and $\overline{\lambda}$ defined in (9). Each case below is labeled according to the "type" of the largest equilibrium (e.g. viral full-adoption).[25]

Proposition 6 (Characterization of equilibria for antisocial behaviors)**.** *For antisocial behaviors, the equilibrium structure depends on network density λ as follows:*

1. *(Nonviral full-adoption) When $\lambda < \underline{\lambda}$, the unique equilibrium is $\sigma = 1$.*
2. *(Nonviral mixed) When $\underline{\lambda} \leq \lambda \leq \overline{\lambda}$, there is a unique equilibrium determined by the nonviral mixed strategy defined by $\lambda\sigma = \underline{\lambda}$.*
3. *(Viral full-adoption) When $\lambda > \overline{\lambda}$, there is a viral full-adoption equilibrium $\sigma = 1$ along with a nonviral mixed equilibrium $\lambda\sigma = \underline{\lambda}$, and a viral mixed equilibrium $\lambda\sigma = \overline{\lambda}$.*

Similar to the prosocial case (Proposition 4), we may understand Proposition 6 through the relationship between (1) expected influence, (2) the threshold $\frac{c}{v} - 1$ in the best response and (3) the strategic substitutability/complementarity of adoption in the sub/super critical regions. When the potential adopter network is sparse ($\lambda < \underline{\lambda}$ or low levels of adoption $\sigma < \underline{\lambda}/\lambda$) then expected influence is small ($< \frac{c}{v} - 1$) and the best response is to adopt. Hence, full-adoption is a unique nonviral equilibrium in sufficiently sparse networks (Part 1). In the sub-critical region ($\lambda < 1$) adoption is a strategic substitute and as the network becomes more dense, expected influence increases. Hence for sufficiently high levels of adoption, influence exceeds the threshold $\frac{c}{v} - 1$, and full-adoption is no longer sustainable as an equilibrium (non-adoption is the best response to sufficiently high adoption). Rather, a nonviral mixed equilibrium emerges as the unique equilibrium (Part 2). As the density increases ($1 < \lambda < \overline{\lambda}$), the underlying network passes into the super-critical region and expected influence remains large. This means the best response to high levels of adoption is non-adoption, and the unique equilibrium remains the

[25] As with prosocial behaviors, the largest viral equilibrium is generically asymptotically stable. Moreover, the second largest (viral mixed) equilibrium is unstable (again, with respect to the mean dynamic for the best-response).

nonviral mixed strategy equilibrium. Hence, unlike the case of prosocial behaviors, as the giant component emerges in the underlying graph, no viral equilibrium emerges alongside it. In the super-critical region, expected influence is decreasing in adoption and network density, so for sufficiently dense networks ($\lambda > \bar{\lambda}$), and sufficiently high adoption ($\sigma > \overline{\lambda}/\lambda$) expected influence falls below the threshold $\frac{c}{v} - 1$ in the best response. Once this occurs, a full-adoption equilibrium and viral mixed equilibrium emerge (Part 3).

A key consequence of Proposition 6 is that *viral* equilibria emerge for antisocial behaviors only when the network density is sufficiently large ($\lambda > \overline{\lambda}$). The mechanism for this result is that antisocial agents are influence-averse: they wish to avoid causing a globally negative outcome. In sparse networks, individual influence is high, creating an accountability pressure that deters adoption. This aligns with the "unbroken windows" effect found by Tsvetkova and Macy (2015b), where observing low prevalence among agents positioned in a linear chains (a nonviral structure) deterred harm. We contend this is because it observing low prevalence increased perceived pivotality for the agents. However, in dense networks ($\lambda > \overline{\lambda}$), the global effect of density reduces individual influence, leading to a "diffusion of responsibility." This creates a *licensing effect*: agents perceive that the negative outcome will spread regardless of their action, lowering the barrier to adoption.

As with Figure 4, we depict the relationship between the expected influence and the best response correspondence for antisocial behaviors. Note that because the utility of antisocial behaviors is the negative of the utility of prosocial behaviors, the plot is identical to Figure 4 but with the best response correspondence "inverted". We see that viral equilibria do not emerge smoothly alongside the phase transition in the underlying graph because when λ is close to the critical threshold, influence becomes too large for agents to be willing to adopt.

Proposition 6 implies that there is a discontinuous jump in adoption under the largest equilibrium when there is diffusion of an antisocial behavior. This is because a viral equilibrium emerges only when the giant component is already well-established in the graph—that is, the critical threshold is $\lambda_{-}^{\text{crit}} = \overline{\lambda} > 1$. In terms of component sizes, we have the following corollary analogous to Corollary 5.

Corollary 7 (Antisocial behaviors emerge discontinuously)**.** *The adoption of antisocial behaviors in the largest equilibrium exhibits a jump discontinuity as a function of the network density, namely*

$$\lim_{\lambda \uparrow \lambda_{-}^{crit}} C_{1,-}(\lambda) = 0 < \lim_{\lambda \downarrow \lambda_{-}^{crit}} C_{1,-}(\lambda).$$

The contrast between Corollary 7 and Corollary 5 has clear consequences for the optimal net-

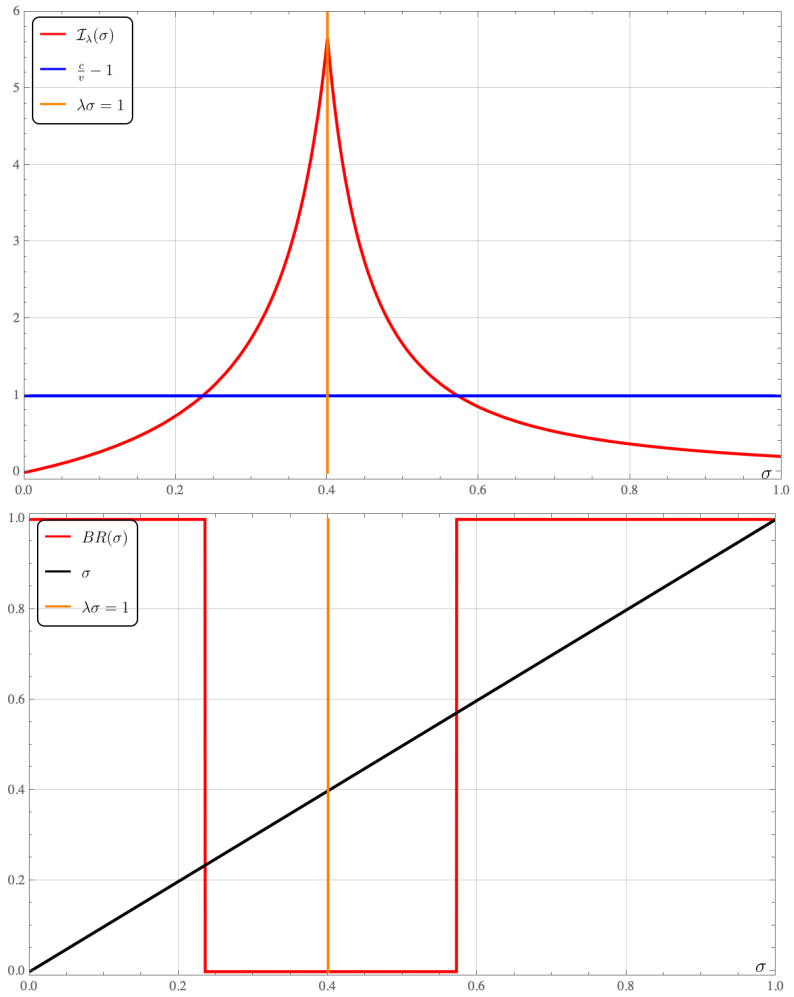

Figure 5: Expected influence and the best response correspondence for antisocial behaviors. Agents adopt only when expected influence is sufficiently small. No one adopts around the critical threshold where $\lambda\sigma = 1$.

work size. The next section synthesizes the findings from these corollaries, discusses welfare implications, presents further comparative statics, and explores policy insights stemming from our model's results.

4.3 Welfare and Policy Implications

From a welfare perspective, optimal policies maximize the spread of prosocial behaviors and minimize the spread of antisocial ones. To design such policies we must first characterize the *extent* of diffusion in the largest equilibrium, which represents the largest fraction of potential adopters who engage in the behavior.

A *large cascade* occurs if a nontrivial fraction of agents take the action in equilibrium. If there is a viral equilibrium σ, then the size of a large cascade is given by $\sigma(1-\rho_\sigma)$.[26] We plot the size

[26]This quantity represents the probability that a randomly chosen agent in the giant component of the potential

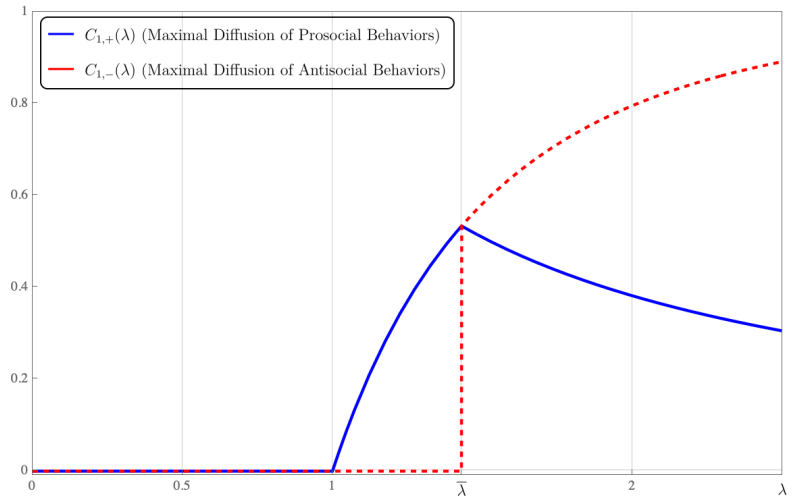

Figure 6: Size of a Large Cascade: Prosocial vs. Antisocial Behaviors. There is a discontinuous jump in maximal diffusion of antisocial behaviors at $\overline{\lambda}$.

of a large cascade for prosocial and antisocial behaviors in the largest equilibrium in Figure 6, and formalize the figure in Proposition 8.

Proposition 8 (Effect of Density on Large Cascade Size)**.** *Considering the largest equilibrium for prosocial and antisocial behaviors, the effect of an increase in the network density (λ) on the size of a large cascade are given in the table below (0 indicates no large cascades):*

λ	*prosocial behaviors*	*antisocial behaviors*
$\lambda < 1$	*0*	*0*
$1 < \lambda < \overline{\lambda}$	*strictly increasing*	*0*
$\lambda > \overline{\lambda}$	*strictly decreasing*	*strictly increasing*

Proposition 8 is a key comparative static result. It illustrates that the *bystander effect* for prosocial behaviors and the *licensing effect* for antisocial behaviors strengthen as network density (λ) increases beyond $\overline{\lambda}$. The result can be understood by decomposing the impact of increased density into two parts: a *direct effect* due to more potential adopters, and an *indirect effect* due to agents' adjustment of equilibrium strategies.

For prosocial behaviors, as density grows, individuals free-ride on others' influence which reduces the overall adoption— that is, the indirect effect dominates the direct effect. This is consistent with empirical findings that individual charitable donations decrease as the number of other contributors they observe increases (Tsvetkova and Macy, 2014). In contrast, for antisocial behaviors agent's equilibrium strategy is the same for all $\lambda > \overline{\lambda}$, so there is no indirect effect. Hence increased density only reinforces the "licensing effect" of the antisocial behavior, leading

adopter network takes the action (σ), multiplied by the size of the giant component in the potential adopter network ($1 - \rho_\sigma$).

to increased overall adoption in the larger network. In dense networks we would therefore expect that antisocial behaviors experience widespread diffusion while prosocial behaviors diffuse very little. This is consistent with the empirical finding of Ferrali (2025) that dishonest behavior diffuses while honest behavior does not.

This result also has clear implications for the speed of diffusion. In dense networks (where $\lambda > \overline{\lambda}$), equilibrium strategies in antisocial behaviors ($\sigma = 1$) are larger than equilibria in prosocial behaviors ($\sigma < 1$). Hence all else equal, antisocial cascades grow more rapidly than prosocial ones. This implication is consistent with empirical evidence accounts of "sudden surges" in crime (Fagan et al., 2007), the emergence of riots (Myers, 2011), and even the rapid spread of Ponzi schemes (Rantala, 2019). This implication is also consistent with the main result of Dimant (2019): that antisocial behaviors are more contagious than prosocial ones.

Next, consider a social planner who can choose the network density (λ) *ex-ante*, before knowing whether the behavior that emerges and diffuses on the network will be prosocial or antisocial. Let the planner's objective be to maximize welfare, which increases with the size of diffusion of prosocial behaviors and decreases with the size of diffusion of antisocial behaviors. We will focus on the case where the planner anticipates that the largest equilibrium will be played for whichever behavior eventually diffuses.[27] As a consequence of Proposition 8, the supremum of welfare is attained in the limit $\lim_{\lambda \uparrow \overline{\lambda}} \lambda$. However, due to the discontinuity in the size of large cascades for antisocial behaviors at $\overline{\lambda}$, there is no λ^* that maximizes welfare, though we can of course get arbitrarily close to the supremum.

Proposition 9 (Socially Optimal Network Density)**.** *For any arbitrarily small positive amount ($\varepsilon > 0$), there is a small interval ($I_\varepsilon = (\overline{\lambda} - \delta_\lambda, \overline{\lambda})$) below the viral threshold $\overline{\lambda}$ such that for any network density (λ) is in this interval (I_ε), the resulting social welfare is arbitrarily close to (within ε of) the supremum of welfare over all possible network densities (λ).*

The core insight of Proposition 9 is that a social planner faces a fundamental trade-off: while greater network density can facilitate the spread of prosocial behaviors, exceeding the critical threshold $\overline{\lambda}$ can lead to the abrupt emergence of antisocial behaviors. A converse implication is that if a network is large ($\overline{\lambda} > \lambda$), then a mob-rule can be overcome by lowering the connectivity below the critical threshold $\overline{\lambda}$. For example, in dense online environments contributing to the spread of misinformation, a planner may be able to eliminate a viral "misinformation equilibrium" by reducing the ease of information propagation across the network (e.g. through limits on sharing content) thereby increasing individual accountability for the content shared. Finally, Proposition 9 implies that for densities near $\overline{\lambda}$, the losses from the adoption of prosocial

[27] The analysis is virtually identical if the planner holds a probability distribution over the equilibrium which will be played and puts nonzero probability on the largest equilibrium.

behaviors are small relative to the gains from adoption of antisocial behaviors, so a planner who is uncertain about the exact threshold is better off "cutting more conservatively".

In our model, equilibrium is sensitive to three key parameters that govern agent incentives: the private value of acting (v), the private cost of acting (c), and the perceived reach of spillovers (δ). The interaction of these parameters determines the critical network density threshold $\overline{\lambda}$ around which viral antisocial equilibria emerge. Policies which target these parameters can improve outcomes by increasing the "weight" of individual influence.

Events like the mass protests in Iran following the death of Mahsa Amini in 2022 or the global "Black Lives Matter" movement sparked by the death of George Floyd in 2020 can be understood in our model as causing sharp shifts in parameter values. For example, such catalysts can simultaneously lower the perceived costs of participation (c), increase the moral or social value of acting (v), and heighten public belief that collective action will generate meaningful change that benefits everyone (δ). In our framework, such shifts can move the system from a regime where influence is insufficient to motivate action ($\lambda < \underline{\lambda}$) into the "high influence" regime ($\underline{\lambda} < \lambda < \overline{\lambda}$) where participation is a strategic complement, by changing individuals' beliefs about the number of potential adopters. Importantly, if the perceived density remains below the threshold ($\overline{\lambda}$) where the bystander effect dominates, incentives to participate are strong. By contrast, our model predicts that for antisocial behaviors, marginal shifts in these parameters can trigger discontinuous changes in the equilibrium level of adoption. The following proposition formalizes the impact of "nonstructural interventions" which target these parameters.

Proposition 10 (Impact of Spillover Reach (δ), Values (v) and Costs (c))**.** *For the largest equilibrium, a marginal increase in δ or v, or a marginal decrease in c has the following effects:*

1. ***Prosocial Behaviors:*** *No effect when $1 < \lambda \leq \overline{\lambda}$, strictly (and continuously) increases the size of large cascades when $\lambda > \overline{\lambda}$.*

2. ***Antisocial Behaviors:*** *No effect when $\lambda \neq \overline{\lambda}$, and a discontinuous decrease in the size of large cascades when $\lambda = \overline{\lambda}$.*

These findings suggest several interesting policy avenues. First, observe that Proposition 10 implies that nonstructural interventions are only effective for prosocial behaviors when there is already a bystander effect ($\lambda > \overline{\lambda}$), and for antisocial behaviors around the critical threshold ($\lambda \approx \overline{\lambda}$). When these conditions are not met, only changes to the network structure itself will affect equilibrium adoption. Since targeting the private value (v) or cost (c) of an action are traditional policy levers with fairly clear implications in our model, we will focus on discussing policies which target δ.

For prosocial behaviors, interventions that make the perceived extent of spillover reach (δ) more salient to agents can be effective. For example, vaccination campaigns emphasizing that community immunity makes the *individual* safer may encourage more individuals to participate. Similarly, in campaigns for volunteering emphasizing that a culture of community service eventually creates a more supportive local network for the individual volunteer can improve the individual's incentive to volunteer. By making the personal benefit from widespread adoption more apparent, these interventions raise agents' perceptions of δ. For antisocial behaviors, Proposition 10 implies that interventions are only effective around the critical threshold. Small changes in δ can create, or eliminate antisocial behavior. The strategic goal of policy in this case is to make agents aware of how widespread adoption— and their role in it— ultimately harms their own welfare.

Our analysis in this section demonstrates the rich set of policy implications that can be derived from our model. By disentangling the role of network density (λ) from the model parameters (δ, v, c), our model provides distinct levers for intervention. We demonstrate that optimal network connectivity is "large but not too large", and that increasing agents' perceptions about the extent of spillovers or weight of influence is always (weakly) beneficial, yielding discontinuous benefits for antisocial behaviors.

5 Extensions

5.1 Arbitrary Degree Distributions

When it comes to real-world networks, there are many properties that the binomial random graph $G(n, p)$ does not capture, for example "fat tails".[28] As such, it is natural to ask whether our results can be extended beyond the specialized class of Poisson networks. The object of this section is to answer this in the affirmative.

We can easily extend our model to a game on a network with an arbitrary degree distribution constructed using the *configuration model* (Bollobás, 1980; Wormald, 1978). In this setting, so long as the degree distribution is sufficiently sparse, the graph is locally tree-like and therefore can be approximated using branching processes (we discuss this in more detail in Online Appendix B).

An approach using generating functions works in this more general setting. Formally, define a game with n agents as a three-tuple $\Gamma^{(n)} = (\mathbf{d}^{(n)}, \mathcal{A}, u)$. The vector $\mathbf{d}^{(n)} = (\mathbf{d}_1^{(n)}, \ldots, \mathbf{d}_n^{(n)})$ is the

[28] See e.g., Jackson (2008).

degree sequence for the game and generalizes the binomial random graph of Section 2.[29] The actions and utility functions remain unchanged. We impose standard restrictions on the degree sequence so that the limiting degree distribution, which we denote by $\mathcal{D} \equiv \lim_{n\to\infty} \mathbf{d}^{(n)}$, is well behaved (in particular it must have a finite mean—see Online Appendix C for details).

The timing, payoffs and strategies in the more general model remain unchanged except that expectations will be different because they depend on the degree distribution. Let

$$G_0(z) = \sum_{k=0}^{\infty} p_k z^k$$

be the generating function for $\mathcal{D} = \{p_k\}_{k\geq 0}$. Then, the generating function for $\mathcal{D}'$ is given by

$$G_1(z) = \frac{G_0'(z)}{G_0'(1)},$$

where $G_0'(1) = \mathbb{E}(\mathcal{D})$. If all agents play the strategy σ, then the generating function for the forward adoption degree distribution is $G_1(1-\sigma+\sigma z)$. It follows that the *forward extinction probability* ρ_σ for the forward adoption degree distribution is the smallest solution in $[0,1]$ to the equation

$$\rho_\sigma = G_1(1-\sigma+\sigma\rho_\sigma). \tag{11}$$

Note that this implies

$$\frac{d}{dz} G_1(1-\sigma+\sigma z)\big|_{z=\rho_\sigma} = \sigma G_1'(1-\sigma+\sigma\rho_\sigma). \tag{12}$$

Equation (12) describes the expected offspring in the subcritical dual branching process with offspring distribution $\mathcal{D}$. In the Poisson model, the critical site percolation threshold was given by $\lambda\sigma^{\text{crit}} = 1$, i.e., $\sigma^{\text{crit}} = \lambda^{-1}$. In general, the critical site percolation threshold for a graph with an arbitrary degree distribution is known to be

$$\sigma^{\text{crit}} \equiv \frac{\mathbb{E}[\mathcal{D}]}{\mathbb{E}[\mathcal{D}(\mathcal{D}-1)]}.$$

We now give an explicit formula for the expected influence function. We have the following analogue of Lemma 2.

[29] This can be viewed as a special case of Sadler's single-type diffusion game.

Lemma 11. *The expected influence function* $\mathcal{I}_{\mathcal{D}}\colon [0,1] \to \overline{\mathbb{R}}$ *is given by*

$$\mathcal{I}_{\mathcal{D}}(\sigma) = \begin{cases} \frac{\delta\sigma G_1'(1)}{1-\delta\sigma G_1'(1)}, & \sigma < \sigma^{crit} \\ \frac{\delta\rho_\sigma(\widehat{G}_1'(1)-1)}{1-\delta\sigma G_1'(1-\sigma+\sigma\rho_\sigma)}, & \sigma > \sigma^{crit} \end{cases} \tag{13}$$

where ρ_σ *is as in (11), and*

$$\widehat{G}_1'(1) \equiv \frac{\sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma)}{1-G_0(1-\sigma+\sigma\rho_\sigma)}$$

is the expected number of neighbors who are potential adopters after updating due to viral inference. In particular, $\mathcal{I}_{\mathcal{D}}(\sigma)$ *is continuous.*

It turns out to be difficult in general to determine the behavior of (13) with respect to σ, despite there being a "discrete duality principle" (Molloy and Reed, 1998) analogous to the Poisson case.[30] However, we can provide a simple condition under which the case of an arbitrary degree distribution is analogous to the Poisson model.

Our results on the characterization of equilibria in Section 2 relied on a continuous parameterization of the underlying density of the graph (λ). The reason for this is that λ is a sufficient statistic for σ^{crit} and a giant component exists in the graph if and only if $\sigma^{\text{crit}} < 1$. Many distributions of interest can be parameterized in a similar way, for example any degree distribution that scales exponentially in the degree (e.g., a power law), or any family of mixed Poisson distributions. As such, we now restrict our attention to one-parameter families of degree distributions $\{\mathcal{D}_\theta\}_{\theta\in\Theta}$, for which $\theta \in \Theta \subseteq [0,\infty)$ is a sufficient statistic for $\frac{\mathbb{E}[\mathcal{D}_\theta]}{\mathbb{E}[\mathcal{D}_\theta(\mathcal{D}_\theta-1)]}$. We focus on distributions for which $\frac{\mathbb{E}[\mathcal{D}_\theta]}{\mathbb{E}[\mathcal{D}_\theta(\mathcal{D}_\theta-1)]}$ is strictly decreasing in θ, with $\inf_{\theta\in\Theta}\frac{\mathbb{E}[\mathcal{D}_\theta]}{\mathbb{E}[\mathcal{D}_\theta(\mathcal{D}_\theta-1)]} < 1$. This guarantees that there exists a critical threshold θ^c such that a giant component exists in the graph with degree distribution $\mathcal{D}_\theta$ if and only if $\theta > \theta^c$. By re-normalizing $\theta \mapsto \theta/\theta^c$ we my assume without loss of generality that $\theta^c = 1$.

For these degree distributions, we can offer a characterization of equilibria that is entirely analogous to Propositions 4 and 6, if the following condition is satisfied.[31]

Condition 1. Let $\sigma \in [0,1]$ be any strategy.

(i) If σ is nonviral, then $\mathcal{I}_\theta(\sigma)$ is strictly increasing in σ and strictly increasing in θ.

[30] We can characterize exactly what happens in the subcritical regime and around the critical threshold, but as the graph moves further into the supercritical regime we cannot say exactly what happens to influence for an arbitrary degree distribution. We discuss this more in Appendix F.1.

[31] We can weaken this condition slightly as we discuss in Online Appendix F.2.

(ii) If σ is viral, then $\mathcal{I}_\theta(\sigma)$ is strictly decreasing in σ and strictly decreasing in θ.

Condition 1 is identical to Proposition 3 but with λ replaced by the variable θ, which parameterizes the family of distributions. How restrictive is Condition 1? It is not difficult to show that (i) *always* holds, and that (ii) holds around the phase transition (i.e., when $\sigma \approx \sigma^{\text{crit}}$). Whether (ii) always holds away from the phase transition we do not know, however it does hold in the examples we provide in Section 5.2.[32]

We write $\underline{\theta}, \overline{\theta}$ for the solutions to

$$\frac{c}{v} - 1 = \mathcal{I}_\theta(1),$$

where $\underline{\theta} < 1 < \overline{\theta}$.[33] For any family of distributions satisfying Condition 1, we have the following analogue of Proposition 4.

Proposition 12 (Characterization of equilibria for prosocial behaviors)**.** *Let $\{\mathcal{D}_\theta\}_{\theta\in\Theta}$ be a family of degree distributions satisfying Condition 1. Then*

1. *(No-adoption) If $\theta < \underline{\theta}$, there is a unique equilibrium $\sigma = 0$.*
2. *(Nonviral full-adoption) If $\underline{\theta} < \theta < 1$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium $\sigma = \frac{c-v}{G_1'(1)\delta c}$, and a nonviral full-adoption equilibrium $\sigma = 1$.*
3. *(Viral full-adoption) If $1 < \theta < \overline{\theta}$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium $\sigma = \frac{c-v}{G_1'(1)\delta c}$ and a viral full-adoption equilibrium $\sigma = 1$.*
4. *(Viral mixed) If $\theta > \overline{\theta}$, there is a no-adoption equilibrium $\sigma = 0$, a nonviral mixed equilibrium $\sigma = \frac{1}{G_1'(1)}\left(1 - \frac{v}{c}\right)$, and a viral mixed equilibrium defined by the largest solution to $\mathcal{I}_\theta(\sigma) = \frac{c}{v} - 1$.*

The characterization of equilibria for antisocial behaviors is analogous to Proposition 6 and therefore omitted here. It is an immediate consequence of our analysis that the welfare and policy analysis provided in Section 4.3 also apply to graphs with an arbitrary degree distribution that satisfy Condition 1.

5.2 Examples

In this section we provide two examples of graphs with an arbitrary degree distribution— the Zipf and exponential distributions— for which we are able to derive explicit formulas for the

[32] Whether there exists a general class of distributions for which the condition holds is an open question.

[33] It is straightforward to show that $\underline{\theta}$ always exists. Whether $\overline{\theta}$ exists depends on c and v. If $\overline{\theta}$ does not exist then we set it equal to $+\infty$.

expected influence. We show that the behavior is qualitatively identical to the Poisson case, because both examples satisfy Condition 1. The only notable difference between the Poisson case and the two examples we now consider is that influence under the Zipf or exponential distribution has a fatter right-hand tail, so the upper threshold at which viral equilibria for antisocial behaviors emerge is higher than under the Poisson distribution— this can be seen in Figures 7 and 8.

5.2.1 Zipf Distribution

Consider a configuration model network with degree distribution

$$p_k = \frac{e^{\alpha \underline{k}}}{\Phi(e^{-\alpha}, 1, \underline{k})} \frac{e^{-\alpha k}}{k}, \tag{14}$$

where $\Phi(z, s, k)$ is the Lerch transcendent function and $\underline{k}$ is the smallest degree that occurs with nonzero probability. Equation (14) defines the so-called *Zipf* distribution. The Zipf distribution exhibits a power-law of the form $k^{-\gamma}$ with $\gamma = 1$, and an exponential tail controlled by the parameter $\alpha > 0$.[34] We focus on the case where $\underline{k} = 1$ since this is analytically tractable. We show in Online Appendix F.3.1 that expected influence can be computed explicitly using the parameterization $\theta = (e^{\alpha} - 1)^{-1}$ and is given by

$$\mathcal{I}_\theta(\sigma) = \begin{cases} \frac{\theta\sigma\delta}{1-\theta\sigma\delta}, & \theta\sigma < 1 \\ \frac{\delta(\theta\sigma+1)(\theta\sigma-1)}{\theta\sigma(\theta\sigma-\delta)\ln(\theta\sigma)} - \frac{\delta}{\theta\sigma-\delta}, & \theta\sigma > 1. \end{cases}$$

The parameterization θ also allows us the convenience of being able to plot expected influence under the Zipf degree distribution and under the Poisson degree distribution on the same axes with the same critical point. In Figure 7 we take $\sigma = 1$ and plot expected influence for the Zipf and Poisson degree distributions. Crucially, Condition 1 holds for the Zipf distribution, and therefore Proposition 12 provides a full characterization of the equilibrium structure as a function of θ. The thresholds $\underline{\theta}, \overline{\theta}$ are determined by the equations

$$\frac{\underline{\theta}\delta}{1 - \underline{\theta}\delta} = \frac{c}{v} - 1 \iff \underline{\theta} = \frac{c - v}{c\delta}, \tag{15}$$

$$\frac{\delta(\overline{\theta}\delta + 1)(\overline{\theta}\delta - 1)}{\overline{\theta}(\overline{\theta} - \delta)\ln(\overline{\theta})} - \frac{\delta}{\overline{\theta} - \delta} = \frac{c}{v} - 1, \tag{16}$$

and the second equation can be solved numerically for specific values of v, c and δ. E.g. when $c/v = 2$ and $\delta = 0.8$ we have $\overline{\theta} = 1.67$ (and $\underline{\theta} = 0.625$). In our next example the expected influ-

[34] The Zipf distribution arises in a number of real-world settings, including the distribution of city sizes. See e.g., Gabaix (1999), Ioannides and Overman (2003), and Arshad et al. (2018).

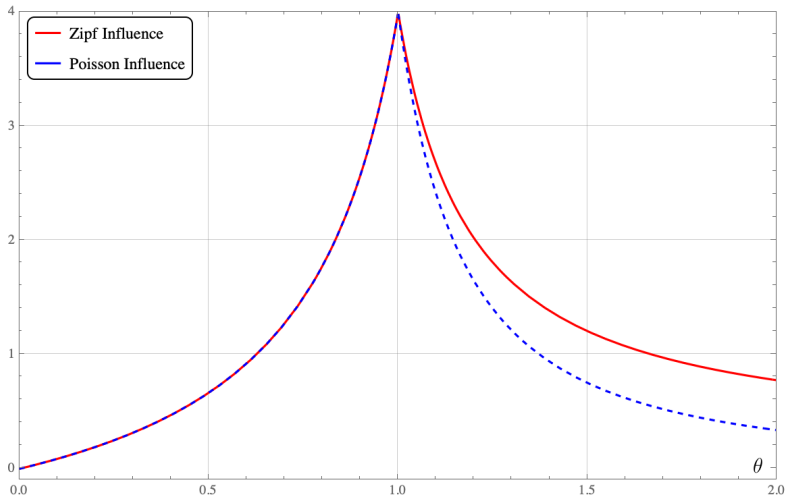

Figure 7: Zipf and Poisson Influence

ence can again be solved for analytically, and we use numerical methods to verify Condition 1.

5.2.2 Exponential Distribution

Next, we consider a configuration model network with degree distribution

$$p_k = Ae^{-\alpha k}, \tag{17}$$

where $A = e^{\alpha} - 1$. This is the discrete *exponential distribution* on $k \in \{1, 2 \dots\}$. We parameterize the distribution in terms of $\theta = 2(\mu - 1)$, where $\mu = E[\{p_k\}_{k \geq 0}] = (1 - e^{-\alpha})^{-1}$. Remarkably, we show in Online Appendix F.3.2 that the expected influence can be written in closed form as

$$\mathcal{I}_\theta(\sigma) = \begin{cases} \frac{\delta\theta\sigma}{1-\delta\theta\sigma}, & \theta\sigma < 1 \\ \frac{16\delta}{\theta\sigma(\theta\sigma+6)-16\delta+(2+\theta\sigma)\sqrt{\theta\sigma(\theta\sigma+8)}}, & \theta\sigma > 1. \end{cases}$$

We plot the expected influence under an exponential distribution against a Poisson distribution in Figure 8— we can see directly from the figure that the exponential distribution satisfies Condition 1. And so it follows that the qualitative results of our analysis hold for an exponential degree distribution. Moreover, since $\lim_{\theta \to \infty} \mathcal{I}_\theta(1) \to 0$, both thresholds $\underline{\theta}$ and $\overline{\theta}$ exist and determine the equilibrium.

5.3 Robustness and Discussion

Our theoretical analysis uses sparse random graphs and therefore abstracts from several features of real-world social networks such as clustering. Moreover, our utility specification takes a specific stance on how influence enters into agents' preferences which may not always apply. In this

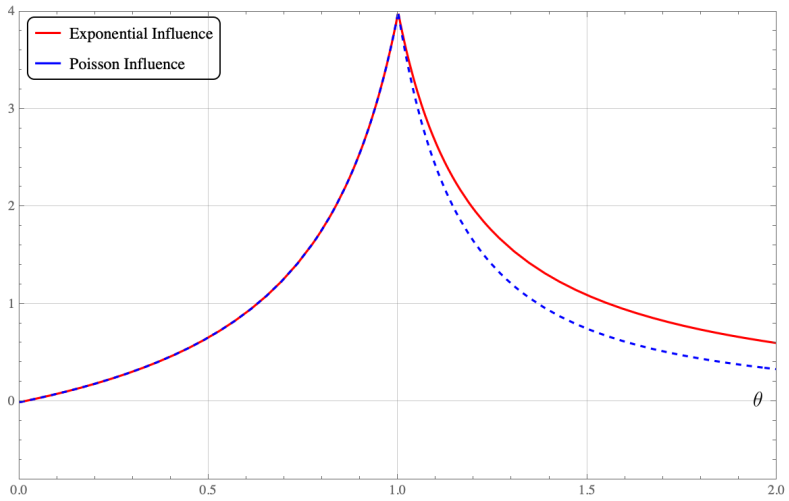

Figure 8: Influence for the Exponential Distribution

section, we highlight the limitations of our model and discuss the implications for applications to which our framework is a better (or worse) fit.

To begin, we use numerical simulations on finite networks to address concerns about our utility specification and the use of sparse random graphs.

5.3.1 Fractional Influence and Finite Populations.

A central concern in applications such as protests or voting is that agents often care about the *fraction* of the population participating (which determines the probability of success), rather than the absolute number of attendees. If the influence of a single agent on the aggregate participation rate vanishes as $n \to \infty$, one might worry that the strategic motives we identify become negligible in large populations.

To address this, we simulated our model on Erdős-Rényi graphs with a finite population of $n = 10{,}000$ agents. Figure 9 reports the average influence of an agent in the largest connected component (C_1) as a function of the mean degree of the potential adopter network. For illustrative purposes, we take $\delta = 0.95$.

Figure 9 demonstrates that influence remains quantitatively significant even when scaled by population size. At the critical threshold (mean degree ≈ 1), the average influence is approximately 9. In a population of 10,000, this represents a shift of nearly 0.1% of the total population, and a significantly larger fraction (3-4%) of the active component. The active component represents the set of agents who can observe the behavior and therefore might participate, so this is arguably the more important metric for applications.

More generally at the critical point, the giant component exhibits a tree-like structure where the

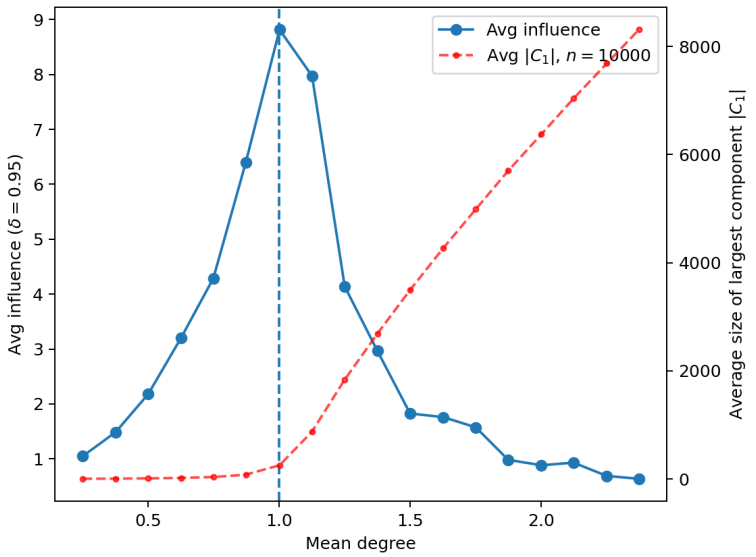

Figure 9: Average influence and the largest component size $|C_1|$ in finite populations. We sample $k = 200$ nodes from 200 graph realizations per σ, in intervals of 0.05, fixing $\lambda = 2.5$.

average path length scales as the square root of the number of vertices in the component (Rényi and Szekeres, 1967, Eq. 4.7). Since the influence of a node in a tree is proportional to the size of the branch defined by that node (which scales with the path length to the root), and the largest component at criticality scales with $n^{2/3}$, the average influence is expected to scale with $(n^{2/3})^{1/2} = n^{1/3}$ (or $n^{-1/3}$ as a fraction of the largest component). Hence for $n = 1{,}000{,}000$, our heuristic suggests that peak influence (when $\delta = 1$) should be around $n^{-1/3} = 100$. We ran additional simulations in a population of one-million and found that the results conform remarkably well with or estimate: peak influence was 99.36. This represents around $n^{-1/3} = 1\%$ of the active component. In settings where the probability of success is a threshold function of participation, a marginal shift of this magnitude by a single pivotal agent is economically non-negligible. That being said, we acknowledge that in arbitrarily large networks, the fraction of the population influenced by a single agent vanishes asymptotically. Consequently, the importance of our key mechanism diminishes as a function of the population size, if what matters to individuals is the fraction—rather than the number—of others they influence.

More importantly, Figure 9 confirms that the *shape* of the influence function in finite populations mirrors the asymptotic prediction: it is increasing in the subcritical regime and decreasing in the supercritical regime.[35] Consequently, even if utility is driven by the fraction of participants, the marginal incentives follow the exact non-monotonic pattern required to generate our central results: strategic complementarity in sparse networks and strategic substitutability in dense ones.

[35]We also verified that the shape of influence holds on average in small networks of $n = 100$.

5.3.2 Clustering and Network Topology.

Our analytical results use branching-process approximations that assume the network to be locally tree-like. Real-world social networks, however, often exhibit high clustering. This introduces local redundancy, which could theoretically dampen the "local effect" of influence, or accelerate the "global effect" of density.

To test the robustness of our mechanism to clustering, we simulated influence on two alternative network topologies on $n = 10{,}000$ vertices: (1) Watts-Strogatz (WS) small-world networks with (Watts and Strogatz, 1998), which allow us to tune clustering via a rewiring parameter β, and (2) geometric lattices (Square and Triangular), which exhibit highly redundant paths and a rigid spatial structure. In both cases we take $\delta = 0.95$.[36] Figure 10 displays the results.

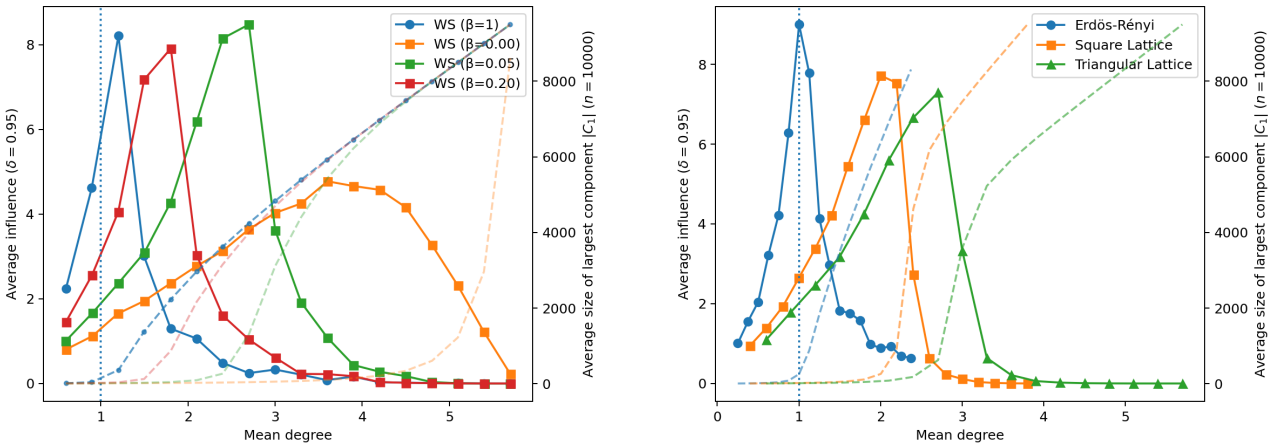

(a) Small-World Networks (varying β)

(b) Lattice Structures

Figure 10: Average influence and network topology. Panel (a) compares Watts-Strogatz small-world networks with ring-lattice degree $K = 6$, varying the rewiring probabilities β. Panel (b) compares geometric lattices. While the magnitude of peak influence varies with the structural properties of the topology, the fundamental non-monotonic shape is robust across all specifications.

In Panel (b), we see that in highly structured Square and Triangular lattices, the peak level of influence is dampened compared to the Erdös-Rényi graph. This is intuitive: spatial lattices induce high local clustering, meaning a neighbor is likely connected to another neighbor, creating redundant paths that render any single agent less pivotal.

However, Panel (a) reveals a more complex relationship in Small-World networks. As we decrease the rewiring parameter β from 1 (Random Graph) toward 0 (Regular Ring-Lattice), the peak influence does not decay monotonically. When $\delta < 1$, we find that the peak influence is maximized

[36] As with Figure 9, we sample $k = 200$ nodes from 200 graph realizations per σ, fixing the average degree. In the WS graph we report total rewiring ($\beta = 1$) in place of an Erdős-Rényi graph, while for the lattices the Erdős-Rényi graph is identical to Figure 9.

at intermediate values of β.[37] This suggests a distinct mechanism in the "small-world" regime: when β is small but non-zero, the network retains local clusters but gains long-range shortcuts. The agents controlling access to these shortcuts become disproportionately influential, serving as the sole "bridges" connecting otherwise distant parts of the network. This implies that while local clustering generally creates redundancy, the introduction of sparse shortcuts can actually amplify the strategic importance of pivotal agents.

Crucially, the fundamental non-monotonicity is preserved across all specifications. In both the Small-World networks and the geometric lattices, influence is strictly monotone increasing as the network approaches the critical density, and strictly monotone decreasing beyond it. This confirms that while topology alters the magnitude of influence, it does not alter the fundamental structural trade-off between local connectivity and global redundancy. Therefore, the "strategic flip" from complements to substitutes is robust to realistic network structures.

5.3.3 Limitations

We now discuss several limitations of our model.

One-shot, single-exposure decisions. We assume agents act once upon first exposure and observe only that at least one neighbor adopted. This "simple contagion" structure fits cases where a single observed act plausibly triggers a response (e.g., a costly prosocial gesture or a clear norm violation). It is a poorer fit for settings in which *multiple* simultaneous or repeated peer exposures are the driver of adoption ("complex contagion")—for instance, enterprise software rollouts, health behaviors that require reinforcement, or consumer technologies where several peers adopting is pivotal. In such environments, neighbor-threshold rules are a better description of utility and our influence-based best responses and the associated complementarity/substitutability flips may not apply without additional modeling of exposure counts and timing.

Broadcast, directionality, and exogenous seeding. We model undirected, neighborhood observation initiated by a single random seed. By contrast, environments with (i) broadcast or long-tie exposure (e.g., platform featuring, influencer posts), (ii) directed attention flows (algorithmic feeds/retweets), or (iii) *many* simultaneous seeds depart from the local, path-by-path propagation that underpins our influence calculus. Long directed ties expose far-flung nodes

[37] When $\delta = 1$, we find that the peak is strictly *decreasing* in β; that is, peak influence is maximized in the regular ring lattice. This suggests influence around the threshold is driven by the diameter of the active component. Since regular lattices exhibit large characteristic path lengths ($L \sim N$) while random graphs exhibit short ones ($L \sim \ln N$), removing the penalty for distance ($\delta = 1$) allows the "elongated" structure of the lattice to yield the highest influence. Conversely, $\delta < 1$ penalizes long paths, which is what delivers the nonmonotonicity in peaks present in panel (a) of Figure 10.

at once and many seeds raise the chance that adoption touches multiple large components irrespective of individual pivotality. In such settings, the prosocial/antisocial asymmetry we emphasize can be diluted: broadcast can spur high adoption of prosocial actions (e.g., a charity drive boosted across a platform) even when dense-network free-riding would otherwise make prosocial behavior substitutable, and it can likewise ignite antisocial cascades (e.g., coordinated brigading or doxxing) without relying on the low-influence licensing channel we model. These cases are better captured by augmented models with layered (local+broadcast) or directed attention networks; our local-observation mechanism is most informative when exposure is primarily neighbor-based and seeding is sparse.

Common-sign spillovers. Our baseline analysis presumes agreement among agents on the sign of spillovers for a given behavior (prosocial "+" or antisocial "−"). In many settings the same act can be perceived as beneficial by some and harmful by others—e.g., whistleblowing inside a firm (prosocial for organizational integrity or the public; antisocial from the perspective of implicated colleagues), or participation in a politically motivated boycott. Such *sign heterogeneity* induces interacting diffusion problems across types and can attenuate or overturn the clean bystander/licensing asymmetry and the sharp discontinuity for antisocial behaviors. Our results may still apply in networks among those who agree on the sign of the behavior, but addressing these environments in full would require a multi-type extension in which agents' payoffs depend on both their own type and the type distribution of adopters.

The limitations above do not negate the mechanism we highlight—the tension between local observability and global redundancy—but they clarify the conditions under which our sharp predictions are most credible and where additional modeling or empirical work may help.

6 Conclusion

We develop and analyze a model of behaviors with positive or negative spillovers that propagate through social networks via observational learning. We show that a key strategic consideration is influence: the *causal effect* of an agent's adoption decision on the subsequent choices of the others in the network. Our analysis reveals a non-trivial effect of greater density on behavior which arises due to two competing effects: a direct *positive* effect from increased observability, and a global *negative* effect as greater connectivity leads to more redundant paths. We establish that there are some robust properties in the way connectivity is related to influence: at first, the direct positive effect dominates and influence increases, then past a critical threshold of connectivity the negative global effect dominates and influence decreases. Our findings unify empirical evidence on the "bystander effect" for prosocial behaviors, and the "licensing effect"

for antisocial behaviors. Our model also explains why a single behavior can be a strategic complement in one setting, and a substitute in another.

We also demonstrate that there is a stark difference between the emergence of viral equilibria for prosocial versus antisocial behaviors. Prosocial behaviors exhibit a *continuous* emergence alongside the emergence of the giant component ($\lambda = 1$) in the underlying graph, while antisocial behaviors exhibit a *discontinuous* emergence in which equilibrium adoption appears suddenly at a higher critical threshold of density ($\lambda = \overline{\lambda} > 1$). This is consistent with empirical evidence on the rapid spread of antisocial behaviors relative to prosocial ones.

Leveraging these insights, our model yields clear and actionable policy recommendations. We demonstrate that socially optimal networks are highly connected, but not so connected as to permit viral antisocial equilibria. Furthermore, we highlight how non-structural interventions can be used effectively by targeting agents' beliefs about the reach or extent of spillovers to raise the weight they attach to influence. We suggest that this is one way to view the effects of catalyzing events such as those which trigger mass protests.

The robustness of our findings across different network topologies and alternative utility specifications underscores the generalizability of our insights. Our paper highlights the complex interplay between network structure, spillovers, and strategic influence, providing a helpful framework for understanding diffusion and improving welfare outcomes in complex systems. Our research also opens up a rich set of empirically testable questions for future research on the strategic diffusion of behaviors in social networks.

References

Akbarpour, M., S. Malladi, and A. Saberi (2023). Just a few seeds more: The inflated value of network data for diffusion. Unpublished manuscript, Stanford University.

Allouch, N. (2015). On the private provision of public goods on networks. *Journal of Economic Theory 157*, 527–552.

Allouch, N. (2017). The cost of segregation in (social) networks. *Games and Economic Behavior 106*, 329–342.

Arshad, S., S. Hu, and B. N. Ashraf (2018). Zipf's law and city size distribution: A survey of the literature and future research agenda. *Physica A: Statistical Mechanics and its Applications 492*, 75–92.

Athreya, K. and P. Ney (1972). *Branching Processes*. Springer, Heidelberg.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who's who in networks. wanted: the key player. *Econometrica 74*(5), 1403–1417.

Bénabou, R., A. Falk, and J. Tirole (2020). Narratives, imperatives, and moral persuasion. Unpublished manuscript, Princeton University.

Blume, L. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior 5*, 387–424.

Boldrin, M. and D. K. Levine (2008). Perfectly competitive innovation. *Journal of Monetary Economics 55*(3), 435–453.

Bollobás, B. (1980). A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics 1*(4), 311–316.

Bollobás, B. and O. Riordan (2015). An old approach to the giant component problem. *Journal of Combinatorial Theory, Series B 113*, 236–260.

Bramoullé, Y. and R. Kranton (2007). Public goods in networks. *Journal of Economic Theory 135*(1), 478–494.

Bramoullé, Y. and R. Kranton (2016). Games played on networks. In: Y. Bramoullé, A. Galeotti and B. Rogers (Eds.), *Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press.

Bramoullé, Y., R. Kranton, and M. D'Amours (2014, March). Strategic interaction and networks. *American Economic Review 104*(3), 898–930.

Brock, W. and S. Durlauf (2001). Discrete choice with social interactions. *Review of Economic Studies 68*, 235–260.

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

Callaway, D., M. Newman, S. Strogatz, and D. Watts (2000). Network robustness and fragility: Percolation on random graphs. *Physical Review Letters 85*(25), 5468–5471.

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

Campbell, A., C. M. Leister, P. Ushchev, and Y. Zenou (2025). Competing diffusions in a social network. Unpublished manuscript, Monash University.

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

Cantoni, D., D. Yang, N. Yuchtman, and Y. Zhang (2019). Protests as strategic games: Experimental evidence from Hong Kong's antiauthoritarian movement. *Quarterly Journal of Economics 134*(2), 1021–1077.

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

Dasaratha, K. (2023, 01). Innovation and strategic network formation. *The Review of Economic Studies 90*(1), 229–260.

Diestel, R. (2000). *Graph Theory, 2nd ed.* Berlin: Springer-Verlag.

Dimant, E. (2019). Contagion of pro- and anti-social behavior among peers and the role of social proximity. *Journal of Economic Psychology 73*, 66–88.

Easley, D. and J. Kleinberg (2010). *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge: Cambridge University Press.

Edmond, C. (2013). Information manipulation, coordination, and regime change. *The Review of Economic Studies 80*, 1422–1458.

Elliott, M. and B. Golub (2019). A network approach to public goods. *Journal of Political Economy 127*(2), 730–776.

Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica 61*, 1047–1071.

Erdős, P. and A. Rényi (1959). On random graphs. *Publicationes Mathematicae 6*, 290–297.

Fagan, J., D. Wilkinson, and G. Davies (2007). Social contagion of violence. In: Daniel J. Flannery, Alexander T. Vazsonyi & Irwin D. Waldman (Eds.)*, The Cambridge Handbook of Violent Behavior and Aggression,* Cambridge: Cambridge University Press.

Fei, G., X. Jiang, X. Zhang, Y. Yang, P. Lin, Q. Sun, Y. Xiao, Y. Sun, X. Li, Y. Qian, J. Gu, Y. Meng, X. Wang, J. Sheng, F. Liu, M. Li, L. Stallones, and H. Xiang (2024). The contagious spread of bullying among Chinese adolescents through large school-based social networks. *Computers in Human Behavior 158*, 108282.

Ferrali, R. (2025). The unequal diffusion of honesty and dishonesty in workplace networks. *Management Science,* forthcoming.

Gabaix, X. (1999). Zipf's law for cities: An explanation. *Quarterly Journal of Economics CXIV*, 739–767.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology 83*, 1420–1443.

Ioannides, Y. M. and H. G. Overman (2003). Zipf's law for cities: An empirical examination. *Regional Science and Urban Economics 33*, 127–137.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.

Jackson, M. O. and D. López-Pintado (2013). Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science 1*, 49–67.

Jackson, M. O., B. W. Rogers, and Y. Zenou (2017). The economic consequences of social-network structure. *Journal of Economic Literature 55*(1), 49–95.

Jackson, M. O. and L. Yariv (2005). Diffusion on social networks. *Economie Publique 16*, 69–82.

Jackson, M. O. and L. Yariv (2007). Diffusion of behavior and equilibrium properties in network games. *American Economic Review, Papers and Proceedings 97*(2), 92–98.

Jackson, M. O. and L. Yariv (2011). Diffusion, strategic interaction, and social structure. In: Benhabib, J., Bisin, A. and M.O. Jackson (Eds.), *Handbook of Social Economics, Volume 1A,* Amsterdam: Elsevier Science, pp. 645–678.

Jackson, M. O. and Y. Zenou (2015). Games on networks. In P. Young and S. Zamir (Eds.), *Handbook of Game Theory, Vol. 4*, pp. 34–61. Amsterdam: Elsevier Publisher.

Janson, S. (2009). On percolation in random graphs with given vertex degrees. *Electronic Journal of Probability 14*(5), 87–118.

Keizer, K., S. Lindenberg, and L. Steg (2008, 12). The spreading of disorder. *Science 322*(5908), 1681–1685.

Kleinberg, J. (2007). Cascading behavior in networks: Algorithmic and economic issues. In: N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani (Eds.), *Algorithmic Game Theory,* Cambridge: Cambridge University Press.

Kuran, T. (1989). Sparks and prairie fires: A theory of unanticipated political revolution. *Public Choice 61*, 41–74.

Langtry, A. (2023). More connection, less community: Network formation and local public goods provision. Unpublished manuscript, University of Cambridge.

Langtry, A., S. Taylor, and Y. Zhang (2024). Network threshold games. Unpublished manuscript, Cambridge University.

Langtry, A. and D. Thornton (2025). Controlling complex contagions. Unpublished manuscript, UNSW Sydney and University of Bristol.

Leister, C. M., Y. Zenou, and J. Zhou (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.

Molloy, M. and B. Reed (1998). The size of the giant component of a random graph with a given degree sequence. *Combinatorics, Probability and Computing 7*(3), 295–305.

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

Myers, D. (2011). Violent protest and heterogeneous diffusion processes: The spread of U.S. racial rioting from 1964 to 1971. *Mobilization: An International Quarterly 15*(3), 289–321.

Myers, D. J. (2000). The diffusion of collective violence: Infectiousness, susceptibility, and mass media networks. *American Journal of Sociology 106*(1), 173–208.

Newman, M., S. Strogatz, and D. Watts (2001). Random graphs with arbitrary degree distributions and their applications. *Physical Review E 64*(2), 026118.

Olson, M. (1965). *The Logic of Collective Action: Public Goods and the Theory of Groups*. New York: Schocken Books.

Palfrey, T. and H. Rosenthal (1984). Participation and the provision of discrete public goods: A strategic analysis. *Journal of Public Economics 24*(105), 171–193.

Pool, V. K., N. Stoffman, and S. E. Yonker (2015). The People in Your Neighborhood: Social Interactions and Mutual Fund Portfolios. *The Journal of Finance 70*(6), 2679–2732.

Rantala, V. (2019). How do investment ideas spread through social interaction? Evidence from a Ponzi scheme. *The Journal of Finance 74*(5), 2349–2389.

Rényi, A. and G. Szekeres (1967). On the height of trees. *Journal of the Australian Mathematical Society 7*(4), 497–507.

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

Sadler, E. (2025). Seeding a simple contagion. *Econometrica 93*(1), 71–93.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press.

Schelling, T. (1978). *Micromotives and Macrobehavior*. New York: Norton.

Schröder, J. M., E.-M. Merz, B. Suanet, and P. Wiepking (2023). The social contagion of prosocial behaviour: How neighbourhood blood donations influence individual donation behaviour. *Health & Place 83*, 103072.

Shadmehr, M. and D. Bernhardt (2011). Collective action with uncertain payoffs: Coordination, public signals and punishment dilemmas. *American Political Science Review 105*, 829–851.

Stein, J. C. (2008). Conversations among competitors. *American Economic Review 98*(5), 2150–2162.

Traag, V. A. (2016). Complex contagion of campaign donations. *PLOS ONE 11*(4), e0153539.

Tsvetkova, M. and M. Macy (2015a). The Contagion of Prosocial Behavior and the Emergence of Voluntary-Contribution Communities. In B. Gonçalves and N. Perra (Eds.), *Social Phenomena: From Data Analysis to Models*, pp. 117–134. Cham: Springer International Publishing.

Tsvetkova, M. and M. W. Macy (2014). The social contagion of generosity. *PLoS ONE 9*(2), e87275.

Tsvetkova, M. and M. W. Macy (2015b). The social contagion of antisocial behavior. *Sociological Science 2*, 36–49.

Van der Hofstad, R. (2023a). *Random Graphs and Complex Networks. Volume 1.* Cambridge, MA: MIT Press.

Van der Hofstad, R. (2023b). *Random Graphs and Complex Networks. Volume 2.* Cambridge, MA: MIT Press.

Vega-Redondo, F. (2007). *Complex Social Networks*. Cambridge: Cambridge University Press.

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

Watts, D. J. and S. H. Strogatz (1998). Collective dynamics of 'small-world' networks. *Nature 393*(6684), 440–442.

Wormald, N. (1978). *Some problems in the enumeration of labelled graphs*. Ph.d. thesis, Newcastle University.

7 Appendix: Proofs of the results in the main text

7.1 Proof of Proposition 1

We wish to establish the existence of a symmetric perfect Bayesian Equilibrium. We first define a best response correspondence:

$$BR_\beta(\sigma) = \begin{cases} 1.\mathbb{1}_{\beta=+}, & \mathcal{I}_\lambda^{(n)}(\sigma) > \frac{c}{v} - 1 \\ [0,1], & \mathcal{I}_\lambda^{(n)}(\sigma) = \frac{c}{v} - 1 \\ 1.\mathbb{1}_{\beta=-}, & \mathcal{I}_\lambda^{(n)}(\sigma) < \frac{c}{v} - 1. \end{cases} \tag{18}$$

Online Appendix B shows that calculating $\mathcal{I}_\lambda^{(n)}(\sigma)$ via Bayes rule results in function that is a ratio of polynomials with no roots in the denominator and is therefore continuous in σ. Hence, $\sigma \mapsto BR_\beta(\sigma)$ satisfies the conditions of Kakutani's fixed point theorem and there exists a strategy $\sigma^* = BR_\beta(\sigma^*)$. Our equilibrium is then this strategy and the associated beliefs that result in $\mathcal{I}_\lambda^{(n)}(\sigma^*)$. Off-equilibrium beliefs only arise at an information set in an equilibrium where $\sigma^* = 0$. In this case it is trivial to specify any belief where $\mathcal{I}_\lambda^{(n)}(\sigma^*) = 0$ in such an event.

7.2 Proof of Lemma 2

Lemma 2 is a special case of Lemma 11 which is proved in Appendix 7.10. For clarity, we show explicitly here how Lemma 11 implies Lemma 2. In the case where $\lambda\sigma < 1$, observe that for the Poisson distribution we have $G_1'(1) = \lambda$, whence (13) gives

$$\mathcal{I}_\lambda(\sigma) = \frac{\delta\lambda\sigma}{1 - \delta\lambda\sigma}.$$

We now turn to the case where $\lambda\sigma > 1$. The generating function for the Poisson distribution has the property

$$G_0^{(k)}(z) = \lambda^k G_0(z).$$

Substituting this into the expression for the expected number of neighbors conditional on exposure (due to viral inference), we have

$$\begin{aligned} \sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma) &= G_0(1-\sigma)\sum_{k=1}^{\infty}\left(\frac{\sigma^k}{(k-1)!} - \frac{(\sigma\rho_\sigma)^k}{(k-1)!}\right)\lambda^k \\ &= G_0(1-\sigma)\left(\lambda\sigma e^{\lambda\sigma} - \lambda\sigma\rho_\sigma e^{\lambda\sigma\rho_\sigma}\right) \end{aligned}$$

But now recall that $G_0(z) = e^{\lambda(z-1)}$, so

$$\begin{aligned}\lambda\sigma G_0(1-\sigma)\left(e^{\lambda\sigma} - \rho_\sigma e^{\lambda\sigma}\right) &= \lambda\sigma e^{-\lambda(1-\sigma-1)}\left(e^{\lambda\sigma} - \rho_\sigma e^{\lambda\sigma}\right) \\ &= \lambda\sigma\left(1 - \rho_\sigma e^{\lambda\sigma(\rho_\sigma - 1)}\right) \\ &= \lambda\sigma(1-\rho_\sigma^2)\end{aligned}$$

where the last equality comes from the fact that $\rho_\sigma = e^{\lambda\sigma(\rho_\sigma-1)}$. So for the Poisson distribution we arrive at an expected number of neighbors conditional on exposure given by

$$\mathbb{E}[A \mid \mathcal{S}] = \frac{\lambda\sigma\left(1-\rho_\sigma^2\right)}{1 - G_0(1-\sigma+\sigma\rho_\sigma)},$$

and finally since the denominator is equal to $1-\rho_\sigma$ we get $\mathbb{E}[A \mid \mathcal{S}] = \lambda\sigma(1+\rho_\sigma)$, so by (13) we have

$$\mathcal{I}_\lambda(\sigma) = \frac{\delta\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} \times (\lambda\sigma(1+\rho_\sigma) - 1) = \frac{\delta\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} - \delta\rho_\sigma\frac{1-\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma}$$

which proves Lemma 2.

7.3 Proof of Proposition 3

By Lemma 2, expected influence when $\sigma = 1$ is given by

$$\mathcal{I}_\lambda(1) = \begin{cases} \frac{\delta\lambda}{1-\delta\lambda}, & \lambda < 1 \\ \frac{\delta\lambda\rho_1}{1-\delta\lambda\rho_1} - \delta\rho_1\frac{1-\lambda\rho_1}{1-\delta\lambda\rho_1}, & \lambda > 1 \end{cases}$$

In what follows we can always recover the same comparative static result for σ by replacing λ with $\lambda\sigma$ (this is essentially invoking the duality result in Van der Hofstad (2023a, Theorem 3.7)), so it is without loss of generality to focus on $\sigma = 1$. We first consider nonviral strategies ($\lambda < 1$). In this case we have

$$\mathcal{I}_\lambda(1) = \frac{\delta\lambda}{1-\delta\lambda},$$

and

$$\frac{d\mathcal{I}_\lambda(1)}{d\lambda} > 0 \iff (1-\delta\lambda) + \delta\lambda > 0$$

which is always true. We now consider viral strategies. We will show that the result holds for all $\lambda > 2$, and prove the case where $1 < \lambda \leq 2$ in E.1. Both cases rely on the derivative of expected influence with respect to λ, which is given by

$$\frac{d}{d\lambda}\left(\frac{\delta\lambda\rho_1}{1-\delta\lambda\rho_1} - \delta\rho_1\frac{1-\lambda\rho_1}{1-\delta\lambda\rho_1}\right) = \frac{\delta\frac{d(\lambda\rho_1)}{d\lambda}}{(1-\delta\lambda\rho_1)^2} + \frac{(1-\delta)\frac{d(\lambda\rho_1)}{d\lambda}}{\delta\rho_1(1-\delta\lambda\rho_1)^2} - \delta\frac{d\rho_1}{d\lambda}\frac{1-\lambda\rho_1}{1-\delta\lambda\rho_1}, \tag{19}$$

We begin by finding an explicit formula for $\frac{d(\lambda\rho_1)}{d\lambda}$. By (5) and the implicit function theorem we have

$$\frac{d\rho_1}{d\lambda} = \left((\rho_1 - 1) + \lambda\frac{d\rho_1}{d\lambda}\right) e^{\lambda(\rho_1 - 1)} = \frac{\rho_1(\rho_1 - 1)}{1 - \lambda\rho_1}.$$

So,

$$\frac{d(\lambda\rho_1)}{d\lambda} = \rho_1 - \frac{\lambda\rho_1(1-\rho_1)}{1-\lambda\rho_1} = \frac{\rho_1(1-\lambda)}{1-\lambda\rho_1}.$$

It follows that (19) can be written as

$$\frac{dI_\lambda(1)}{d\lambda} = \frac{\delta\rho_1(1-\lambda)(1+\rho_1(1-\delta))}{(1-\lambda\delta\rho_1)^2(1-\lambda\rho_1)} - \frac{\delta\rho_1(\rho_1-1)(1-\lambda\rho_1)(1-\lambda\delta\rho_1)}{(1-\lambda\delta\rho_1)^2(1-\lambda\rho_1)},$$

which is strictly decreasing (for $\lambda > 1$) if and only if

$$(1-\rho_1)(1-\lambda\rho_1)(1-\lambda\delta\rho_1) < (\lambda-1)(1+\rho_1(1-\delta)) \tag{20}$$

Now observe that (20) is true whenever $\lambda > 2$, since in this case we have the chain of inequalities

$$(1-\rho_1)(1-\lambda\rho_1)(1-\lambda\delta\rho_1) < 1 < \lambda - 1 < (\lambda-1)(1+\rho_1(1-\delta)).$$

It remains to show that the inequality holds when $1 < \lambda < 2$. We sketch the proof strategy for this case here but leave the details to the Online Appendix E.1. To prove the claim, we rewrite (20) in a way that gives us a simpler inequality that is sufficient to prove the claim. We show that this simpler inequality holds when $\lambda = 1$, and then show that both sides of the inequality are strictly increasing for all $\lambda > 1$. We conclude the proof by showing that at $\lambda = 2$, the RHS of this inequality is still smaller than the smallest value of the LHS, so the inequality holds everywhere.

7.4 Proof of Proposition 4

First, $\sigma = 0$ is always a best response to itself since

$$v - c + v\mathcal{I}_\lambda(0) = v - c < 0,$$

so agents strictly prefer not to adopt if they expect others to do the same. Hence, no-adoption ($\sigma = 0$) is always an equilibrium. It follows that full-adoption ($\sigma = 1$) is an equilibrium if and only if a mixed strategy equilibrium exists. A nonviral mixed strategy equilibrium exists if and

only if there is a $\sigma \in (0,1)$ satisfying

$$v - c + v\mathcal{I}_\lambda(\sigma) = 0 \iff v - c + \frac{v\delta\lambda\sigma}{1-\delta\lambda\sigma}$$
$$\iff \delta\lambda\sigma = \frac{c-v}{c}.$$

while a viral mixed strategy equilibrium exists if and only if there is a $\sigma \in (0,1)$ with $\rho_\sigma < 1$ satisfying

$$v - c + v\mathcal{I}_\lambda(\sigma) = 0 \iff v - c + \frac{v\delta\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} - v\delta\rho_\sigma\frac{1-\lambda\sigma\rho_\sigma}{1-\delta\lambda\sigma\rho_\sigma} = 0$$
$$\iff \delta\lambda\sigma\rho_\sigma = \frac{c - v(1-\delta\rho_\sigma)}{c + v\rho_\sigma}.$$

It follows that in the subcritical regime (when $\lambda < 1$), we have a mixed strategy equilibrium if and only if

$$1 > \lambda > \frac{c-v}{\delta c} = \underline{\lambda}.$$

Moreover, when $\lambda < \underline{\lambda}$, we have $\lambda\sigma < 1 - \frac{v}{c}$ for all $\sigma \in [0,1]$ and so $\sigma = 0$ is the only equilibrium. This proves the first two cases of Proposition 4. Next, in the supercritical regime ($\lambda > 1$), we know from Proposition 3 that influence is strictly decreasing for any fixed σ. In particular, there exists a smallest λ, which we call $\overline{\lambda}$ such that

$$\overline{\lambda}\delta\rho_1 = \frac{c - v(1-\delta\rho_1)}{c + v\rho_1}.$$

For any $1 < \lambda < \overline{\lambda}$, there is a full-adoption equilibrium, and this is the unique viral equilibrium (the nonviral mixed equilibrium and no-adoption equilibrium are also still present). However, for $\lambda > \overline{\lambda}$, full-adoption is no longer an equilibrium, since $\overline{\lambda}\delta\rho_1 < \frac{c-v(1-\delta\rho_1)}{c+v\rho_1}$, and so by Proposition 3 there is a unique viral equilibrium in mixed strategies determined by $\lambda\sigma = \overline{\lambda}$. This proves the final two cases of Proposition 4.

7.5 Proof of Corollary 5

The first equality holds because the LHS is necessarily 0 for all $\lambda < 1$. The second inequality holds because the fraction of agents in the largest component of $G(n,p)$ can be made arbitrarily small as $\lambda \to 1^+$, and in a full-adoption equilibrium (which exists as $\lambda \to 1^+$ by Proposition 4) the potential adopter network is all of $G(n,p)$.

7.6 Proof of Proposition 6

The details of the proof are virtually identical to the proof of Proposition 4. We provide a sketch for completeness. The key difference is that when λ is small, $\sigma = 1$ is the unique nonviral equilibrium, since

$$c - v - v\mathcal{I}_\lambda(1) > 0.$$

However when λ becomes sufficiently large (exceeds $\underline{\lambda}$), full-adoption is no longer an equilibrium since agents expect to influence too many others. This leads to a mixed strategy nonviral equilibrium and no-adoption equilibrium when $\underline{\lambda} < \lambda < \overline{\lambda}$. Finally, when $\lambda > \overline{\lambda}$, the giant component is large enough such that expected influence under full-adoption is small, and so a viral full-adoption (and viral mixed) equilibrium emerges.

7.7 Proof of Corollary 7

The proof follows from Proposition 6 in the same way that Corollary 5 follows from Proposition 4. As such, we omit the details.

7.8 Proofs of Propositions 8 and 9

We first note that in the subcritical regime ($\lambda < 1$), all equilibria are nonviral and so their size is 0. Consider the diffusion of prosocial behaviors. By Proposition 4, we know that when $1 < \lambda < \overline{\lambda}$, the unique viral equilibrium is the full-adoption equilibrium. It follows that the size of a large cascade in this equilibrium is simply $1 - \rho_1$, which is strictly increasing in λ since extinction becomes less likely as the network gets more connected. However, once $\lambda > \overline{\lambda}$, full-adoption is no longer an equilibrium and the unique equilibrium is the mixed strategy equilibrium σ^* determined by

$$\lambda\sigma^*\delta\rho_{\sigma^*} = \frac{c - v(1 - \delta\rho_{\sigma^*})}{c + v\rho_{\sigma^*}}, \tag{21}$$

where

$$\rho_{\sigma^*} = e^{\lambda\sigma^*(1-\rho_{\sigma^*})}.$$

Since ρ_{σ^*} is a function of $\lambda\sigma^*$, and the equilibrium is unique, it follows that we must have $\lambda\sigma^* = \overline{\lambda}$. In other words, in equilibrium, σ is chosen such that ρ_{σ^*} is held constant. Hence when $\lambda > \overline{\lambda}$, an increase in λ keeps $\lambda\sigma^*$ fixed, and therefore σ^* must be strictly decreasing in λ. Finally, since ρ^* is constant for all $\lambda > \overline{\lambda}$, and diffusion is determined by $\sigma^*(1 - \rho_{\sigma^*})$, we have that the diffusion of prosocial behaviors is strictly decreasing in λ in the largest equilibrium, as claimed.

Now we consider the diffusion of antisocial behaviors. Proposition 6 tells us that when $\lambda < \overline{\lambda}$, there are no viral equilibria, while for $\lambda > \overline{\lambda}$, there is a full-adoption equilibrium. Hence by the same reasoning as in with prosocial behaviors, the size of a large cascade, $1 - \rho_1$, is increasing in λ whenever $\lambda > \overline{\lambda}$.

Proposition 9 is immediate. The diffusion of prosocial behaviors is maximized at $\lambda = \overline{\lambda}$, while the diffusion of antisocial behaviors is 0 if and only if $\lambda < \overline{\lambda}$.

7.9 Proof of Propositions 10

It is convenient to prove the proposition by way of the following Lemma.

Lemma 13. *In the largest equilibrium, a marginal increase (decrease) in $\overline{\lambda}$ has the following effects:*

1. ***Prosocial Behaviors:*** *No effect when $1 < \lambda \leq \overline{\lambda}$, strictly and continuously increases (decreases) the size of large cascades when $\lambda > \overline{\lambda}$.*

2. ***Antisocial Behaviors:*** *No effect when $\lambda \neq \overline{\lambda}$, and a discontinuous decrease in (no effect on) the size of large cascades when $\lambda = \overline{\lambda}$.*

Proof. We begin with the case of prosocial behaviors. If $1 < \lambda < \overline{\lambda}$, then by Proposition 4, the largest equilibrium is the full-adoption equilibrium, so the marginal effect of a change in $\overline{\lambda}$ on the largest equilibrium is null— the full adoption equilibrium remains. On the other hand, if $\lambda > \overline{\lambda}$ then a viral mixed equilibrium exists in which $\lambda\sigma = \overline{\lambda}$, and hence an increase (decrease) in $\overline{\lambda}$ increases (decreases) σ in order to maintain $\lambda\sigma = \overline{\lambda}$, leading to a strict increase (decrease) in the diffusion of prosocial behaviors.

Next we consider antisocial behaviors. In this case, the marginal effect of a change in $\overline{\lambda}$ on the largest equilibrium is only nonzero when $\lambda = \overline{\lambda}$, since the largest equilibrium is no adoption to the left of $\overline{\lambda}$, and full adoption to the right of $\overline{\lambda}$. At the point $\lambda = \overline{\lambda}$, an increase in $\overline{\lambda}$ leads to a discontinuous decrease in the size of equilibrium. However, the marginal effect of a decrease in $\overline{\lambda}$ is always null because the discontinuity is right-continuous rather than left-continuous. □

Since increasing δ increases expected influence for all λ, it shifts $\overline{\lambda}$ upwards. Hence the case of δ in Proposition 10 follows immediately from Lemma 13.

To prove the cases for v and c, observe that these moderate equilibrium spread through $\overline{\lambda}$, so by Lemma 13 it suffices to show that $\overline{\lambda}$ is strictly increasing in v and strictly decreasing in c.

To this end, recall the definition

$$\overline{\lambda}\delta\rho_1 = \frac{c - v(1 - \delta\rho_1)}{c + v\rho_1}.$$

The RHS of the above equation is decreasing in v and increasing in c. In particular, since $\lambda\delta\rho_1$ is decreasing when $\lambda < \overline{\lambda}$, it follows that if the RHS increases, $\overline{\lambda}$ decreases. This implies that $\overline{\lambda}$ is increasing in v and decreasing in c, which completes the proof.

7.10 Proof of Lemma 11

Suppose agents are playing the strategy σ. Let $F_{1,\sigma}(z)$ be the generating function for the distribution over discounted finite "forward component sizes" in the potential adopter network (*not* necessarily the whole network) from following a randomly chosen edge. Then as in Callaway et al. (2000) we have

$$F_{1,\sigma}(z) = \delta z G_1[1 - \sigma + \sigma F_{1,\sigma}(z)], \tag{22}$$

from which it follows that

$$F_{1,\sigma}(1) = \delta G_1[1 - \sigma + \sigma F_{1,\sigma}(1)] = \delta\rho_\sigma.$$

We can find the expected finite forward component size in the potential adopter network by taking the derivative of (22) at $z = 1$ and substituting in the expression for $F_{1,\sigma}(1)$. That is,

$$F'_{1,\sigma}(1) = \delta G_1[1 - \sigma + \sigma F_{1,\sigma}(1)] + \delta\sigma F'_{1,\sigma}(1) G'_1[1 - \sigma + \sigma F_{1,\sigma}(1)] = \frac{\delta\rho_\sigma}{1 - \delta\sigma G'_1(1 - \sigma + \sigma\rho_\sigma)}. \tag{23}$$

Since, in the limit, the graph is locally tree-like, each neighbor of an agent can be considered as an independent cluster of potential influence (we discuss this in more depth in Online Appendix B). If σ is nonviral, then $\rho_\sigma = 1$, and exposure gives no information about the degree distribution. Hence the total expected influence of an agent i is $F'_{1,\sigma}(1)$ multiplied by i's expected discounted number of potential-adopter neighbors: $\mathbb{E}[\delta A_i] = \delta\sigma G'_1(1)$. That is,

$$\mathcal{I}_\theta(\sigma) = \frac{1}{1 - \delta\sigma G'_1(1)} \times \mathbb{E}[A_i] = \frac{\delta\sigma G'_1(1)}{1 - \delta\sigma G'_1(1)}.$$

When σ is viral, there is a giant component of potential adopters. We consider two cases. The first is that the seed is in the giant forward component (which happens with probability $\to 1$ in large networks). In this case, expected influence is the expected finite forward component through each neighbor; this is simply $F'_{1,\sigma}(1)$ for each neighbor. The second case is that the seed is in a finite forward component, in which case an agent is pivotal for the information reaching

the giant component. We return to this second case at the end of the proof, for now we calculate the expected number of neighbors conditional on exposure.

Let $\mathcal{S}_i$ denote the event that at least one neighbor of a randomly chosen vertex i is in the giant component of potential adopters. Let A_i denote the number of neighbors of i who are potential adopters. Let d_i denote i's degree. The quantity we need to find is $\mathbb{E}[A_i \mid \mathcal{S}_i]$, since the total discounted expected influence is $F'_{1,\sigma}(1) \times \mathbb{E}[A_i \mid \mathcal{S}_i]$. By the law of total expectation,

$$\mathbb{E}[A_i \mid \mathcal{S}_i] = \sum_{d_i=1}^{\infty} \mathbb{E}[A_i \mid \mathcal{S}_i, d_i]\mathbb{P}(d_i \mid \mathcal{S}_i).$$

By Sadler (2020, Theorem 3),

$$\mathbb{P}(A_i = k \mid \mathcal{S}_i, d_i) = \frac{1-\rho_\sigma^k}{1-(1-\sigma+\sigma\rho_\sigma)^{d_i}}\mathbb{P}(\text{Bin}(d,\sigma)=k).$$

So we have

$$\mathbb{E}[A_i \mid \mathcal{S}_i, d_i] = \sum_{k=1}^{d_i} k\left(\frac{1-\rho_\sigma^k}{1-(1-\sigma+\sigma\rho_\sigma)^{d_i}}\mathbb{P}(\text{Bin}(d_i,\sigma)=k)\right).$$

Now notice that (dropping the i subscripts for readability)

$$\mathbb{E}[A \mid \mathcal{S}, d]\mathbb{P}(d \mid \mathcal{S}) = \left[\sum_{k=1}^{d} k\left(\frac{1-\rho_\sigma^k}{1-(1-\sigma+\sigma\rho_\sigma)^{d}}\mathbb{P}(\text{Bin}(d,\sigma)=k)\right)\right]\left[\frac{p_d\left(1-(1-\sigma+\sigma\rho_\sigma)^d\right)}{1-\sum_{k\geq 0}p_k(1-\sigma+\sigma\rho_\sigma)^k}.\right]$$

$$= \frac{1}{1-G_0(1-\sigma+\sigma\rho_\sigma)}\sum_{k=1}^{d} kp_d\left(1-\rho_\sigma^k\right)\mathbb{P}(\text{Bin}(d,\sigma)=k).$$

By Tonelli's theorem, the double sum $\sum_{d=1}^{\infty}\sum_{k=1}^{n}$ can instead be computed as

$$\begin{aligned}\sum_{k=1}^{\infty}\sum_{d=k}^{\infty} kp_d\left(1-\rho_\sigma^k\right)\mathbb{P}(\text{Bin}(d,\sigma)=k) &= \sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\sum_{d=k}^{\infty} p_d\mathbb{P}(\text{Bin}(d,\sigma)=k) \\ &= \sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma)\end{aligned}$$

Subtracting 1 from the above expression (for the link along which exposure occurred) gives us the numerator for expected influence under viral strategies. Hence if the seed is in the giant

forward component then expected influence is

$$F'_{1,\sigma}(1)\mathbb{E}[A_i \mid \mathcal{S}_i] = \frac{\delta\rho_\sigma}{1-\delta\sigma G'_1(1-\sigma+\sigma\rho_\sigma)}\left(\frac{\sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma)}{1-G_0(1-\sigma+\sigma\rho_\sigma)}-1\right)$$
$$= \frac{\delta\rho_\sigma(\widehat{G}'_1(1)-1)}{1-\delta\sigma G'_1(1-\sigma+\sigma\rho_\sigma)},$$

where $\widehat{G}'_1(1) \equiv \frac{\sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma)}{1-G_0(1-\sigma+\sigma\rho_\sigma)}$. It remains to prove that when the seed is in a finite forward component (the "second case" mentioned earlier in the proof), the contribution to discounted expected influence is negligible and so the above expression indeed captures the influence in line with Lemma 11. It follows from the above analysis that the seed is in a finite forward component with probability $\frac{F'_{1,\sigma}(1)\mathbb{E}[A_i|\mathcal{S}_i]}{cn}$, where cn is the size of the giant component. We show that in this case the expected influence is $o(n)$. To this effect, notice that the expected size of influence is this case is bounded above by a branching process with expected offspring $\lambda\sigma\delta$, since this would be the expected offspring in the absence of any indirect paths. We decompose this branching process into two parts. The first part counts all agents within $\log\log(n)$ graph distance of the root, and the second part counts all agents greater than distance $\log\log(n)$ of the root. The first part grows as:

$$1+\lambda\sigma\delta+(\lambda\sigma\delta)^2+\cdots+(\lambda\sigma\delta)^{\log\log n} = \frac{(\lambda\sigma\delta)^{\log\log n+1}-1}{\lambda\sigma\delta-1}$$
$$= O\left((\lambda\sigma\delta)^{\log\log n}\right)$$
$$= O\left((\log n)^{\log(\lambda\sigma\delta)}\right) = o(n).$$

where the second last equality follows from the fact that $r^{\log n} = e^{\log r\log n} = n^{\log r}$. Next, letting $\delta = e^{-b}$ for some $b \geq 0$, the second part (agents greater than $\log\log n$ steps from the root) grows at most as fast as

$$\delta^{\log\log n}cn = \frac{cn}{(\log n)^b} = o(n),$$

where c is the fraction of agents in the potential adopter network. This bound holds because there are at most cn agents that could possibly be influenced, but these agents must be discounted by at least $\delta^{\log\log n}$ under our construction— the entire forward component size is "too far way" from any one agent to make a nontrivial difference to the calculation of expected influence. Hence the total expected influence in both the first and second cases is $o(n)$, whence the contribution to expected influence in the case where the seed is in a finite forward component is

$$\frac{F'_{1,\sigma}(1)\mathbb{E}[A_i \mid \mathcal{S}_i]}{cn} \times o(n) = o(1),$$

and therefore contributes negligibly, as claimed.

Online Appendix

A Poisson Model

A.1 Branching Processes and Duality

There are several texts one can consult for a detailed treatment of branching processes, e.g. Van der Hofstad (2023a); Athreya and Ney (1972). The most important results on Galton-Watson branching processes are also outlined in Appendix A of Sadler (2020). As such, we confine ourselves here to presenting results on the duality properties of supercritical branching processes, and for this we follow Van der Hofstad (2023a).

Let $\mathcal{D} \equiv \{p_k\}_{k\geq 0}$ be a probability distribution over the nonnegative integers, and suppose that $\mathcal{D}$ is the offspring distribution of a branching process. The branching process is said to be *subcritical* if $\mathbb{E}[\mathcal{D}] < 1$, and *supercritical* if $\mathbb{E}[\mathcal{D}] > 1$. This is because the *extinction probability,* ρ i.e. the probability that the branching process eventually dies out, is 1 if $\mathbb{E}[\mathcal{D}] < 1$, and strictly less than 1 if $\mathbb{E}[\mathcal{D}] > 1$.

Let $G_{\mathcal{D}}(z) = \sum_k p_k z^k$ denote the generating function for the distribution $\mathcal{D}$. Then the extinction probability ρ is defined by the smallest solution in $[0, 1]$ to the equation $\rho = G_{\mathcal{D}}(\rho)$.

Call the distributions $\{p_k\}_{k\geq 0}$ and $\{p_k^{'}\}_{k\geq 0}$ a *conjugate pair* if

$$p'_k = \rho^{k-1} p_k.$$

It is easy to check that $\{p'_k\}_{k\geq 0}$ is a probability distribution, since

$$\begin{aligned}\sum_{k=0}^{\infty} p'_k &= \rho^{-1} \sum_{k=0}^{\infty} p_k \rho^k \\ &= \rho^{-1} G_{\mathcal{D}}(\rho) \\ &= \rho^{-1} \rho = 1.\end{aligned}$$

It turns out that the distribution for a supercritical branching process conditioned on extinction, is precisely equal to the conjugate distribution defined above. This is stated formally in the following theorem.

Theorem A1 (Van der Hofstad (2023a, Theorem 3.7))**.** *Let $\{p_k\}_{k\geq 0}$ and $\{p'_k\}_{k\geq 0}$ be a conjugate pair of offspring distributions. The branching process with distribution $\{p_k\}_{k\geq 0}$ conditioned on*

extinction has the same distribution as the branching process with offspring distribution $\{p'_k\}_{k\geq 0}$.

The proof follows directly from Bayes' rule. Theorem A1 takes on a particularly nice form for Poisson branching processes. Let $\{p_k\}_{k\geq 0}$ be a Poisson distribution with mean λ. Then the generating function is given by

$$G_0(z) = e^{\lambda(z-1)},$$

and therefore the extinction probability $\rho = \rho_\lambda$ satisfies

$$\rho_\lambda = e^{\lambda(\rho_\lambda - 1)}. \tag{A1}$$

If $\lambda > 1$, then a branching process with offspring distribution $\{p_k\}_{k\geq 0}$ is supercritical, and therefore by Theorem A1 the distribution conditional on extinction is

$$p'_k = \rho_\lambda^{k-1} p_k = \frac{\rho_\lambda^k}{e^{\lambda(\rho_\lambda - 1)}} \cdot \frac{e^{-\lambda}\lambda^k}{k!} = \frac{e^{-\lambda\rho_\lambda}(\lambda\rho_\lambda)^k}{k!},$$

where the second equality follows from Eqn. (A1). But this distribution is again Poisson, with mean

$$\mu_\lambda \equiv \lambda\rho_\lambda < 1.$$

It follows that a branching process with offspring distribution $\text{Po}(\lambda)$ (where $\lambda > 1$) conditioned on extinction, has the same distribution as a branching process with offspring distribution $\text{Po}(\lambda\rho_\lambda)$. We call a branching process with offspring distribution $\text{Po}(\lambda\rho_\lambda)$ the *subcritical dual* of the supercritical branching process with offspring distribution $\text{Po}(\lambda)$.

Once we introduce percolation, the branching process has offspring distribution $\text{Po}(\lambda\sigma)$, and so the subcritical dual has offspring distribution $\text{Po}(\lambda\sigma\rho_\sigma)$, where we omit the dependence of ρ on λ for readability. Finally, although we formally prove Lemma 2 in Appendix 7.2 of the paper, this duality gives us another way of seeing why the influence function takes on the specific form that it does. In particular consider the following theorem.

Theorem A2 (Van der Hofstad (2023a, Theorem 3.5))**.** *Let T denote the total offspring of a branching process with i.i.d. offspring X, having mean offspring $\mu < 1$, then*

$$\mathbb{E}[T] = \frac{1}{1-\mu}.$$

It follows from Theorem A2 that the expected offspring of the subcritical dual branching process in our setting is

$$\mathbb{E}[T] = \frac{1}{1-\lambda\sigma}.$$

This is remarkably close to our expression for the expected influence. The mean of the offspring distribution in our case is $\lambda\sigma$, but we discount each generation by δ, so the effective mean is $\delta\lambda\sigma$, leading to $\mathbb{E}[T] = (1 - \delta\lambda\sigma)^{-1}$ for each of the $\delta\lambda\sigma$ neighbors an individual has (in expectation). On the other hand, in the supercritical regime ($\lambda\sigma > 1$) expected influence only counts agents along paths that go extinct. When we choose an edge at random in the network and follow it to a vertex incident with it, the vertex has, in expectation, $\lambda\sigma\rho_\sigma$, "forward" potential adopting neighbors along paths that go extinct, so the mean of the relevant offspring distribution is $\delta\lambda\sigma\rho_\sigma$. Hence the expected influence is just $\mathbb{E}[T] = (1 - \delta\lambda\sigma\rho_\sigma)^{-1}$ for each neighbor. The exact calculation for the expected number of neighbors in this case is complicated due to viral inference (see Appendix 7.10 for more details).

The reason we can obtain this result via branching processes is that the Erdös-Rényi random graph is intimately related to a branching process with Poisson offspring distribution. We omit the details here, but one can show that $G(n, \lambda/n)$ converges "locally in probability" to a branching process with offspring distribution $\mathrm{Po}(\lambda)$ (see Van der Hofstad, 2023b, Theorem 2.18). In other words, our analysis of $G(n, \lambda/n)$ in terms of Poisson branching processes can be made precise. We take this approach in the following proof, but present it in a slightly more flexible way that also works for graphs with an arbitrary degree distribution.

B Influence

This section provides a rigorous foundation for our definition of influence in finite graphs and its behavior in the limit as $n \to \infty$. The following explanation can be easily adapted to a graph with an arbitrary degree sequence, and even any distribution over values v. In our case, the degree sequence is Poisson and the distribution over values is a point mass.

Fix a graph $G = (V, E)$ on $V = \{1, 2, \ldots, n\}$ vertices with edge set E, and fix a strategy σ which induces a subnetwork of potential adopters $G_a = (V_a, E_a) \subseteq G$. Let $s \in \{1, 2, \ldots, n\}$ be the seed agent, and for each vertex j define $\mathcal{P}_{G_a}(j, s)$ to be the set of all paths in the potential adopter network G_a from j to s. That is,

$$\mathcal{P}_{G_a}(j, s) = \{v_0 v_1, \ldots, v_k \colon k \in \mathbb{N},\ v_i v_{i+1} \in E_a \text{ for all } i, \text{ and } v_i \text{ are distinct}\}.$$

Write $G_a - \{i\}$ for the graph obtained after removing from G_a the vertex i and any edges incident with it. Write $G_a + \{i\}$ for the graph obtained by adding $i \in V$ to the subgraph $G_a \subseteq G$. If

$i \in V_a$, then $G_a + \{i\} = G_a$. The influence set $\text{Inf}(i) = \text{Inf}_{\{G,G_a,s\}}(i)$ of i in G is defined by[1]

$$\text{Inf}(i) = \{j \in V : P_{G_a+\{i\}}(j,s) \neq \emptyset, \text{ and } P_{G_a-\{i\}}(j,s) = \emptyset\}.$$

That is, the set of all agents $j \in V_a$ for whom every path to s contains i, if i were in the potential adopter network. Now, if $\text{Inf}(i) \neq \emptyset$, then at least one of i's neighbors is in $\text{Inf}(i)$. To see this, note that if every $j \in \text{Inf}(i)$ is not a neighbor of i in $G_a + \{i\}$, then any path from j to s must pass through at least one of i's neighbors. If we let $P = v_0 v_1, \ldots, v_k \in \mathcal{P}(j,s)$ with $v_0 = j$ and $v_k = s$, and $k > 2$, (by the assumption that j is not a neighbor of i) then there is some $\ell < k$ with $v_\ell = i$, and where $v_{\ell-1} \neq j$ is a neighbor of i. But if $v_{\ell-1} \notin \text{Inf}(i)$, then there exists some other path from $v_{\ell-1}$ to s not passing through i. Call this path P_ℓ. Then $Pv_{\ell-1}P_\ell$[2] is a path from j to s not containing i, a contradiction of the fact that $j \in \text{Inf}(i)$.

This gives us a useful way to reformulate i's influence set. We have $j \in \text{Inf}(i)$, if and only if i separates a component of agents containing j from a component of agents containing s. This is useful because if i's influence set is nonempty, then as we have seen, i has at least one neighbor over whom he has influence, and every agent over which i has influence is connected to one of i's neighbors. Hence for any $j \in \text{Inf}(i)$, we know that j is in a component with at least one of i's neighbors. If follows that if i is connected to the seed, then i's influence is the sum over the component sizes in $G_a - \{i\}$ for each of his neighbors over whom he has influence (being careful not to double count if two of i's neighbors are in the same component). Formally, letting $C_G(j)$ denote the component in G containing j, and letting $N_G(i)$ denote the set of neighbors if i in G, we have

$$\text{Inf}(i) = \begin{cases} \bigcup_{j:\, s \notin C_{G_a-\{i\}}(j)} C_{G_a-\{i\}}(j)(s), & s \in C_{G_a+\{i\}}(i), \\ \emptyset, & s \notin C_{G_a+\{i\}}(i). \end{cases} \tag{B2}$$

In words, the influence set of i is the union over all components containing i's neighbors who are not on some alternate path back to the seed. Now, in our model i discounts the value of their influence over those in their influence set by a discount factor $\delta < 1$. So given an influence set $\text{Inf}(i)$ in a graph G with potential adopter network G_a and seed s, we define i's *benefit from influence*[3] as the polynomial

$$\mathcal{I}_{\{G,G_a,s,i\}}(\delta) = \sum_{j \in \text{Inf}(i)} \delta^{\ell_{ij}},$$

where ℓ_{ij} is the length of the shortest path between i and j in $G_a + \{i\}$. In our model, G and G_a

[1] Influence is defined for any vertex in G and not just in G_a, because agents outside the potential adopter network still consider what their influence *would* be if they were to join the potential adopter network.

[2] This notation is standard in Diestel (2000).

[3] Importantly, i's benefit from influence is not the same thing as i's influence set. The benefit from influence discounts agents in the influence set by their graph distance to i. We use the term "expected influence" in the paper to mean the "expected benefit from influence".

are both random. This is because nature draws a graph $G = \mathbb{G}(n, p)$, and then G_a is determined by a realization of the strategy σ. We can view $\mathcal{I}_{\{G,G_a,s,i\}}(\delta)$ as a random variable contingent on the realizations of G, G_a, and s. We now show how to calculate the expectation of $\mathcal{I}_{\{G,G_a,s,i\}}(\delta)$.

Fixing a realization of the graph G, the potential adopter network is determined by n independent Bernoulli experiments, each with success probability σ, and so the probability that any particular potential adopter subgraph is realized is determined by a polynomial in σ. Since agents are aware upon exposure that at least one of their neighbors has adopted, they condition the probability of realizing any particular potential adopter network on the knowledge that at least one of their neighbors is a potential adopter. It follows that the conditional probability is a ratio of positive polynomials in the adoption probabilities σ. Finally, the seed agent is chosen uniformly at random.

However, in our model, strategies are conditional upon exposure, i.e. i has the opportunity to act only if they are in the same component as s. Agents don't know the time period t at which they are exposed, nor do they know which of their neighbors exposed them, as such, upon exposure they condition on the fact that they are connected to the seed in the potential adopter network, and on nothing else. So the quantity we are really interested in is i's expected benefit from influence conditional on the event $\mathcal{S}_i \equiv \{s \in C_{G_a+\{i\}}(i)\}$. Note that since $\text{Inf}(i) = \emptyset$ on the complement of $\mathcal{S}_i$, the expected benefit of influence $\mathbb{E}_\sigma[\mathcal{I}_{\{G,G_a,s,i\}}(\delta)]$ is identical to $\mathbb{E}_\sigma[\mathcal{I}_{\{G,G_a,s,i\}}(\delta)\mathbb{1}_{\mathcal{S}_i}]$. Hence for any agent i, ex-ante the expected size of i's influence under the strategy $\sigma \in [0, 1]$ and conditional on $\mathcal{S}_i$ is

$$\begin{aligned}
\mathbb{E}_\sigma[\mathcal{I}_{\{G,G_a,s,i\}}(\delta) \mid \mathcal{S}_i] &= \frac{\mathbb{E}_\sigma[\mathcal{I}_{\{G,G_a,s,i\}}(\delta)\mathbb{1}_{\mathcal{S}_i}]}{\mathbb{P}(\mathcal{S}_i)} \\
&= \frac{\mathbb{E}_\sigma[\mathcal{I}_{\{G,G_a,s,i\}}(\delta)]}{\mathbb{P}(\mathcal{S}_i)} \\
&= \frac{1}{\mathbb{P}(\mathcal{S}_i)} \sum_G \sum_{G_a} \sum_s \mathcal{I}_{\{G,G_a,s,i\}}(\delta)\mathbb{P}(G)\mathbb{P}(G_a \mid G, \sigma)\mathbb{P}(s \mid G_a, G, \sigma), \quad \text{(B3)}
\end{aligned}$$

where $\mathbb{P}(s \mid G_a, G, \sigma) = \frac{1}{n-1}$ is fixed and independent of the choice of G and s, while each $\mathbb{P}(G_a \mid G, \sigma)$ is polynomial in σ, and each $\mathbb{P}(G)$ is a number in $[0, 1]$ determined by the degree distribution.

The most important consequence of this is that the expected benefit from influence upon exposure is a continuous function of σ, since it is a ratio of polynomials in σ with a denominator that has no roots in $[0, 1]$ (except at $\sigma = 0$ which is trivial).

We now take a closer look at what happens as $n \to \infty$. We have proved that i's influence can be obtained by looking at the union over the components containing i's "forward neighbors",

that is, those neighbors who are not in the same forward component as the seed. We claim that as $n \to \infty$, in the absence of a giant component every agent i of degree d has influence over $d-1$ of his neighbors with probability $1-o(1)$.

To prove the above claim, we use a Lemma from from Bollobás & Riordan (2015) which establishes that the configuration model produces a locally tree-like graph. Given a graph G, let $G_{\leq t}(v)$ denote the subgraph of G induced by the vertices within distance t of v; that is, up to the "t-th neighbours" of v. Let $\mathcal{T} = \mathcal{T}_{\mathcal{D}}$ be a branching process (or a "random rooted tree") on $X_1, X_2, \ldots$ with $X_i \sim \mathcal{D}'$, the forward degree distribution (independently for all i). As with G, let $\mathcal{T}_{\mathcal{D}}|_t$ be the subtree of $\mathcal{T}_{\mathcal{D}}$ induced by the vertices within distance t of the root (that is, the first t generations of the process).

Lemma B3 (Bollobás and Riordan (2015))**.** *Let v be a vertex of $G = \mathbb{G}^{\star}_{\mathbf{d}}$ chosen uniformly at random. Then we may couple the random graphs $G_{\leq t}(v)$ and $T_{\mathcal{D}}|_t$ so that they are isomorphic as rooted graphs with probability $1-o(1)$ as $n \to \infty$.*

This gives us the following corollary.

Corollary B4. *Let v be a vertex of $G = \mathbb{G}^{\star}_{\mathbf{d}}$ chosen uniformly at random. If $t \geq 1$ is a constant, then w.h.p. the neighbourhood $G_{\leq t}(v)$ of v in G is a tree.*

Suppose that the graph has no giant component of potential adopters, so that all component sizes are finite. Consider choosing a vertex i and random, and suppose i has degree d. Corollary B4 implies that if i is exposed then with probability $1-o(1)$, i's removal from the potential adopter network creates $d-1$ additional components– one for each of his neighbors– by separating them from the component containing s. This is because by assumption all of i's neighbors are contained in finite components, and so for any $\epsilon > 0$ we can choose a finite size t such that the fraction of vertices in $\text{Inf}(i)$ which are also contained in the tree $G_{\leq t}(v)$ (minus the component containing i and s) is $1-\epsilon$. It follows that i's expected benefit from influence is simply the expected discounted "forward component size" of each of his neighbors who are potential adopters. When agents do not know their own degrees, they compute the expected influence through each of their forward neighbors. In expectation there are $\sigma\mathbb{E}[\mathcal{D}'] = \sigma\lambda$ forward neighbors who are potential adopters. Since each neighbor is (w.h.p.) contained in a separate forward component, i's expected benefit from influence— i.e. total influence discounted by δ— is $\delta\sigma\lambda\mathcal{J}_{\lambda}(\sigma)$ where we define $\mathcal{J}_{\lambda}(\sigma)$ to be the discounted number of agents contained in the component of *one* of i's potential adopting neighbors, under the strategy σ. Equivalently, we can define $\mathcal{J}_{\lambda}(\sigma)$ as the expected influence of an agent with one forward neighbor in the potential adopter network. In the limit as $n \to \infty$ we have an explicit expression for the component sizes

and forward component sizes (See Callaway et al. (2000) and Newman et al. (2001)), namely

$$\mathcal{I}_\lambda(\sigma) = \delta\lambda\sigma\mathcal{J}(\sigma) = \frac{\delta\lambda\sigma}{1-\delta\lambda\sigma}. \tag{B4}$$

It remains to consider what happens when there are components of infinite size in the potential adopter network. We would like do define influence analogously to the above– as the expected discounted forward component size. If the seed is connected to the giant component on paths not passing through i, then i's influence is the discounted *finite* forward component size.[4] This is because all agents not on finite forward components will be exposed independently of the action of i (a different chain of people from the giant component will reach them). Hence when the seed is in the giant forward component, the influence of an agent i of degree d is, in expectation, their discounted number of potential adopting neighbors multiplied by the expected discounted forward component size through each of these neighbors, conditional on those components being finite (i.e. the paths dying out). This turns out to be precisely the expectation of the generating function $F_{1,\sigma}(z)$, which we define in Appendix 7.10 of the paper and we show its expectation to be

$$F'_{1,\sigma}(1) = \frac{\delta\rho_\sigma}{1-\delta\sigma G'_{1,\sigma}(1-\sigma+\sigma\rho_\sigma)}, \tag{B5}$$

where ρ_σ is the forward extinction probability. When $\rho_\sigma = 1$, (B5) coincides with $\mathcal{J}(\sigma)$ from (B4). However, due to viral inference the expected number of potential adopting neighbors is different when strategies are viral. Let $\mathbb{E}[A_i \mid \mathcal{S}_i]$ be the expected degree of a randomly chosen agent in the giant component. Let $\mathcal{F}$ be the event that the seed is in a finite forward component, and let cn be the size of the giant component. Then total expected influence for large n under viral strategies is

$$\mathcal{I}_\theta(\sigma) = F'_{1,\sigma}(1)\mathbb{E}[A_i \mid \mathcal{S}_i]$$

as $n \to \infty$. As we show in Appendix 7.10 of the paper, the correct expression $\mathbb{E}[A_i \mid \mathcal{S}_i]$ is

$$\mathbb{E}[A_i \mid \mathcal{S}_i] = \frac{1}{1-G_0(1-\sigma+\sigma\rho_\sigma)}\sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma) - 1.$$

This concludes our formal discussion of the notion of influence.

B.1 The case of $\delta = 1$.

A special case arises when $\delta = 1$. In this case, the externality is *perfectly global* in the sense that agents receive utility from all others' who adopt, independently of their position in the

[4] It is possible that the seed is connected to the giant component only on paths passing through i, but in this case the contribution to expected influence is negligible, as we show at the end of Appendix 7.10 in the paper.

network. In particular, agents care about the actions of those who are infinitely far away from them.[5] When the network is subcritical, all paths within a single component are finite and so $\ell_{ij} < \infty$ for all j in the same component as i. But when the network is supercritical, the giant component contains infinitely long paths, so there may be agents for whom $\ell_{ij} = \infty$ *and* i can "influence" j in the sense we describe shortly. As such, it is important to distinguish between the case where $\delta < 1$, where agents only care about "local adoption" and $\delta = 1$ where agents care about global adoption.

The key change to expected influence comes when analyzing the case where the seed belongs to one of the finite forward components. Despite the vanishing probability of this occurring, when $\delta = 1$ these events contribute non-negligibly to expected influence because i is essential for connecting the seed to the giant forward component. The contribution to expected influence of this event is given by the product of the probability of the seed being in one of the finite forward components and the expected influence in this event. If the giant component contains a fraction c of the vertices, then the probability that the seed is in one of the finite forward components is $\frac{\frac{\lambda\sigma\rho_\sigma}{1-\lambda\sigma\rho_\sigma}-\rho_\sigma}{cn}$ (in the limit of a large network), which goes to 0, while the expected size of influence diverges to cn. Hence, the contribution to influence is the product of these and is equal to $\frac{\lambda\sigma\rho_\sigma}{1-\lambda\sigma\rho_\sigma} - \rho_\sigma$.

Therefore, when $\delta = 1$, the expected influence function $\mathcal{I}_\lambda \colon [0,1] \to \overline{\mathbb{R}}$ is given by

$$\mathcal{I}_\lambda(\sigma) = \begin{cases} \frac{\lambda\sigma}{1-\lambda\sigma} & \text{if } \lambda\sigma < 1, \\ 2(\frac{\lambda\sigma\rho_\sigma}{1-\lambda\sigma\rho_\sigma} - \rho_\sigma) & \text{if } \lambda\sigma > 1 \end{cases} \tag{B6}$$

where ρ_σ is as in (5). A similar expression can be derived in the case of a graph with an arbitrary degree distribution.

C The Degree Distribution

In order for $\lim_{n\to\infty} \mathbf{d}^{(n)}$ to be well-behaved, we assume—as is standard—that there exists a distribution $\mathcal{D}$ with finite expectation and with p.m.f. $\{p_k\}_{k\in\mathbb{N}}$, such that for each $k \in \mathbb{N}$,

$$\lim_{n\to\infty} \frac{n_k(\mathbf{d}^{(n)})}{n} = p_k, \text{ and} \tag{C7}$$

$$\lim_{n\to\infty} \frac{m(\mathbf{d}^{(n)})}{n} = \frac{\mathbb{E}(\mathcal{D})}{2}. \tag{C8}$$

[5] Here, we are adopting the convention that $1^\infty = 1$. This is essentially the same as the convention in measure theory that $+\infty \cdot 0 = 0$.

The function $n_k(\mathbf{d}^{(n)})$ is the number of vertices of degree k in $\mathbf{d}^{(n)}$, while $m(\mathbf{d}^{(n)})$ is the number of edges in the graph with degree sequence $\mathbf{d}^{(n)}$.

Under these conditions, all of our quantities of interest (the extinction probability, size of the largest component, etc.) are analytic except at the critical threshold (see Janson, 2009, Theorem 3.11 for details). This justifies (among other things) our implicit differentiation of ρ_σ in the proof of Proposition 3.

D Equilibrium

Although we analyze equilibria in the "limit-game", our pure strategy equilibria exist in any game with sufficiently large n, and our mixed strategy equilibria are limits of mixed equilibria in the finite games (as $n \to \infty$). To see this, consider the case where a prosocial behavior diffuses on the network. Note first that for all n, we have $\mathcal{I}_\lambda(0) = 0$, so consider the case where $\sigma = 1$ is a strict pure strategy equilibrium in the limit-game.

Since $\sigma = 1$ is a strict pure strategy equilibrium, it must be that $\mathcal{I}_\lambda(1) > \frac{c}{v} - 1$. But since $\mathcal{I}_\lambda^{(n)}(\sigma)$ converges to $\mathcal{I}_\lambda(\sigma)$ as $n \to \infty$, there must be some N for which $\mathcal{I}_\lambda^{(n)}(1) > \frac{c}{v} - 1$ for all $n \geq N$. Hence $\sigma = 1$ is an equilibrium in all sufficiently large games.

Finally, suppose $\sigma \in (0, 1)$ is a nonviral mixed strategy equilibrium of the limit-game. Then $\mathcal{I}_\lambda(\sigma) = \frac{c}{v} - 1$. Now fix $\epsilon > 0$. As above, for sufficiently large n we must have that $\mathcal{I}_\lambda^{(n)}(\sigma + \epsilon) > \frac{c}{v} - 1$ and $\mathcal{I}_\lambda^{(n)}(\sigma - \epsilon) < \frac{c}{v} - 1$ (alternatively, for viral equilibria the inequalities are reversed). Hence by continuity of $\mathcal{I}_\lambda^{(n)}(\sigma)$ (established in Appendix B), there is a mixed strategy equilibrium in $(\sigma - \epsilon, \sigma + \epsilon)$, i.e. within distance ϵ of σ. Hence we may take a sequence of mixed strategy equilibria within ϵ of σ for each n sufficiently large. Taking $\epsilon \to 0$ implies that the equilibrium σ can be recovered as the limit of mixed strategy equilibria in the finite games.

E Proofs: Additional Details

E.1 Remainder of Proof of Proposition 3

Recall that our aim is to show that the following inequality ((20) in the paper) holds for all $1 < \lambda \leq 2$:

$$(1 - \rho_1)(1 - \lambda\rho_1)(1 - \lambda\delta\rho_1) < (\lambda - 1)(1 + \rho_1(1 - \delta)).$$

To begin, it is convenient to rewrite the inequality as

$$\frac{(1-\lambda\rho_1)(1-\lambda\delta\rho_1)}{1+\rho_1(1-\delta)} < \frac{\lambda-1}{1-\rho_1}.$$

Observe that for any $\delta \in (0,1)$,

$$\begin{aligned}
&\lambda\delta\rho_1^2 + \delta\rho_1(\lambda-1) > 0 \\
\implies\ &\lambda\delta\rho_1^2 + \lambda\delta\rho_1 - \delta\rho_1 > 0 \\
\implies\ &1+\rho_1-\delta\rho_1 > 1+\rho_1-\lambda\delta\rho_1-\lambda\delta\rho_1^2 \\
\implies\ &1+\rho_1(1-\delta) > (1-\lambda\delta\rho_1)(1+\rho_1) \\
\implies\ &(1+\rho_1(1-\delta))(1-\lambda\rho_1) > (1-\lambda\delta\rho_1)(1+\rho_1)(1-\lambda\rho_1) \\
\implies\ &\frac{1-\lambda\rho_1}{1+\rho_1} > \frac{(1-\lambda\rho_1)(1-\lambda\delta\rho_1)}{1+\rho_1(1-\delta)}
\end{aligned}$$

Hence it suffices to prove that

$$\frac{1-\lambda\rho_1}{1+\rho_1} < \frac{\lambda-1}{1-\rho_1}$$

for all $1 < \lambda \leq 2$.

To this effect, consider the limit

$$L \equiv \lim_{\lambda\to 1^+} \frac{\lambda-1}{1-\rho_1}$$

Since the numerator and denominator both approach 0, by L'Hôpital's rule, we have

$$L = \lim_{\lambda\to 1^+} \frac{1}{-\frac{d\rho_1}{d\lambda}} = \frac{1-\lambda\rho_1}{\rho_1(1-\rho_1)}.$$

A second application of L'Hôpital's rule gives

$$\begin{aligned}
L = \lim_{\lambda\to 1^+} \frac{-\rho_1 - \lambda\frac{d\rho_1}{d\lambda}}{\frac{d\rho_1}{d\lambda}(1-2\rho_1)} &= \left(\lim_{\lambda\to 1^+} \frac{1}{-\frac{d\rho_1}{d\lambda}}\right)\left(\lim_{\lambda\to 1^+} \frac{\rho_1}{1-2\rho_1}\right) + \left(\lim_{\lambda\to 1^+} \frac{-\lambda}{1-2\rho_1}\right) \\
L &= L(-1) + 1 \\
\implies\ L &= \frac{1}{2}.
\end{aligned}$$

On the other hand, we have

$$\lim_{\lambda\to 1^+} \lambda\rho_1 = 1 \implies \lim_{\lambda\to 1^+} \frac{1-\lambda\rho_1}{1+\rho_1} = 0 < \frac{1}{2} = L.$$

This shows that (20) holds as $\lambda \to 1^+$. Next we show that both sides of (20) are strictly

increasing. The LHS is immediate, since

$$\frac{d}{d\lambda}\left(\frac{1-\lambda\rho_1}{1+\rho_1}\right) = \frac{\overbrace{-\frac{d(\lambda\rho_1)}{d\lambda}}^{>0}\overbrace{(1+\rho_1)}^{>0} \overbrace{-\frac{d\rho_1}{d\lambda}}^{>0}\overbrace{(1-\lambda\rho_1)}^{>0}}{(1+\rho_1)^2} > 0.$$

For the RHS, we have

$$\frac{d}{d\lambda}\left(\frac{\lambda-1}{1-\rho_1}\right) = \frac{1-\rho_1+\frac{d\rho_1}{d\lambda}(\lambda-1)}{(1-\rho_1)^2},$$

which is positive iff

$$\begin{aligned} & 1-\rho_1+(\lambda-1)\left(\frac{\rho_1(\rho_1-1)}{1-\lambda\rho_1}\right) > 0 \\ \iff & 1 > \frac{(\lambda-1)\rho_1}{1-\lambda\rho_1} \\ \iff & \rho_1 < \frac{1}{2\lambda-1}. \end{aligned}$$

To prove the last inequality above it suffices to show that

$$e^{\lambda\left(\frac{1}{2\lambda-1}-1\right)} \leq \frac{1}{2\lambda-1}, \tag{E9}$$

since this would imply that the smallest solution in $[0,1]$ to the equation $\rho_1 = e^{\lambda(\rho_1-1)}$ must be at some $\rho_1 < \frac{1}{2\lambda-1}$. By using the fact that

$$\begin{aligned} \lambda\left(\frac{1}{2\lambda-1}-1\right) = \frac{2\lambda(1-\lambda)}{2\lambda-1} &= \frac{1}{2}\left[\frac{1-(2\lambda-1)^2}{2\lambda-1}\right] \\ &= \frac{1}{2}\left[\frac{1}{2\lambda-1}-(2\lambda-1)\right] \\ &= \frac{1}{2(2\lambda-1)}+\frac{1}{2}-\lambda, \end{aligned}$$

so we can rewrite (E9) as

$$e^{-\lambda}\left(e^{\frac{1}{2(2\lambda-1)}+\frac{1}{2}}(2\lambda-1)-e^{\lambda}\right) \leq 0.$$

At $\lambda = 1$ the LHS of the above inequality is equal to 0, so it suffices to prove that the function

$$f(\lambda) \equiv e^{\frac{1}{2(2\lambda-1)}+\frac{1}{2}}(2\lambda-1)-e^{\lambda}$$

is strictly increasing. To show this, first note that for $\lambda > 1$,

$$2(\lambda-1)^2 > 0 \implies (2\lambda-1)\lambda-3\lambda+2 > 0 \implies \lambda > \frac{3\lambda+2}{2\lambda-1}.$$

Next recall that $\log(1+x) < x$ for all $x \neq 0$, so we have

$$\frac{\lambda}{2\lambda-1} + \log\left(1 + \left(1 - \frac{1}{2\lambda-1}\right)\right) < \frac{\lambda}{2\lambda-1} + 1 - \frac{1}{2\lambda-1} = \frac{3\lambda-2}{2\lambda-1} < \lambda.$$

Exponentiating both sides gives

$$e^{\frac{\lambda}{2\lambda-1}} \frac{4\lambda-3}{2\lambda-1} < e^{\lambda}.$$

Finally, some algebraic manipulation shows that

$$\frac{df}{d\lambda} = e^{\lambda} - e^{\frac{\lambda}{2\lambda-1}} \frac{4\lambda-3}{2\lambda-1}.$$

So we see that $\frac{df}{d\lambda} > 0$ iff $e^{\lambda} > e^{\frac{\lambda}{2\lambda-1}} \frac{4\lambda-3}{2\lambda-1}$, which we have proved! This proves (E9) which in turn proves that $\rho_1 < \frac{1}{2\lambda-1}$. Hence

$$\frac{d}{d\lambda}\left(\frac{\lambda-1}{1-\rho_1}\right) > 0$$

for all $\lambda > 1$. To summarize, we have thus far shown that

$$\lim_{\lambda \to 1^+} \left(\frac{\lambda-1}{1-\rho_1}\right) = \frac{1}{2} > 0 = \lim_{\lambda \to 1^+} \left(\frac{1-\lambda\rho_1}{1+\rho_1}\right), \tag{E10}$$

$$\frac{d}{d\lambda}\left(\frac{\lambda-1}{1-\rho_1}\right) > 0 \tag{E11}$$

$$\frac{d}{d\lambda}\left(\frac{1-\lambda\rho_1}{1+\rho_1}\right) > 0. \tag{E12}$$

To complete the proof, we show that $\frac{1-\lambda\rho_1}{1+\rho_1}|_{\lambda=2} < \frac{1}{2}$, which shows that (20) holds when $1 < \lambda \leq 2$. It suffices to prove that at $\lambda = 2$,

$$\rho_1 \geq \frac{1}{5},$$

since this would imply

$$\frac{1-2\rho_1}{1+\rho_1} \leq \frac{1-\frac{2}{5}}{1+\frac{1}{5}} = \frac{1}{2}.$$

But for this it suffices to show that

$$e^{2\left(\frac{1}{5}-1\right)} \geq \frac{1}{5},$$

since by definition of ρ_1 this means $\rho_1 \geq \frac{1}{5}$. It is easily verified (numerically or otherwise) that

$$e^{2\left(\frac{1}{5}-1\right)} \approx 0.202 > \frac{1}{5},$$

and this proves that (20) holds for $1 < \lambda \leq 2$, which in turn proves that

$$\frac{d\mathcal{I}_{\lambda}(1)}{d\lambda} < 0$$

for all $\lambda > 1$, as desired.

E.2 Note on Proposition 8

We used the notation $C_{1,\beta}(\lambda)$ in the text to refer to the fraction of agents in the largest component of the potential adopter network in the largest equilibrium. This is different to the ex-ante expected size of diffusion, but the comparative statics are the same for both. In particular,

$$C_{1,\beta}(\lambda) = \sigma^*(1 - \rho_{\sigma^*}),$$

where σ^* is the largest equilibrium under β and given λ, while the ex-ante expected size of diffusion is

$$C_{1,\beta}(\lambda) \times \mathbb{P}(\text{seed triggers a large cascade}) = \sigma^*(1 - \rho_{\sigma^*})^2,$$

so we see that the relevant comparative statics are identical from the ex-ante perspective.

F Extensions

F.1 Discussion of the Assumptions

F.1.1 Nonconstant Values and Transformation of Influence

To isolate the strategic effect of influence, we took the spillover v to be constant for all agents. In some strategic settings, the value from influencing additional agents to take the action depends on how many agents one anticipates are already going to take it. For example, the marginal benefit from participation in a public protest may be smaller when a large number of agents are already attending— in other words, the value of participation exhibits diminishing marginal returns.

To this end, fix λ and suppose that $v = f(\zeta_\sigma) \neq 0$ is a weakly decreasing function of ζ_σ, where $\zeta_\sigma = \sigma(1 - \rho_\sigma)$ is the size of a large cascade under the strategy σ.[6] Then analogous with (8), agents prefer to choose action 1 when $\mathcal{I}_\lambda(\sigma) > \frac{c}{f(\zeta_\sigma)} - 1$. Since $\frac{c}{f(\zeta_\sigma)}$ is weakly increasing, and $\mathcal{I}_\lambda(\sigma)$ is strictly decreasing for $\sigma > \sigma^{\text{crit}}$, it follows that either all equilibria are nonviral, or there is a unique viral equilibrium. Since influence attains its maximum value at the critical threshold, a unique mixed viral equilbrium exists if and only if influence exceeds the threshold at the σ^{crit}:

$$\mathcal{I}_\lambda(\sigma^{\text{crit}}) > \tfrac{c}{f(\sigma^{\text{crit}})} - 1 \iff \frac{1}{1-\delta} > \frac{c}{f(\sigma^{\text{crit}})}.$$

[6] More generally, we may take $v = f(\sigma)$ for a decreasing function of σ and our main results on viral equilibria are preserved.

It follows that whenever a viral equilibrium exists, our characterization in Proposition 4 remains virtually unchanged for $\lambda \geq 1$— a viral full-adoption equilibrium emerges continuously around $\lambda = 1$ and there is a threshold $\overline{\lambda}$ at which the unique viral equilibrium transitions to a mixed strategy. The figure below shows expected influence against $\frac{c}{f(\sigma)} - 1$ for the special case where $f(\zeta_\sigma) = 1 - \zeta_\sigma$. This choice of f means that for any nonviral equilibria, $f(\zeta_\sigma) = 1$, so the threshold at which nonzero equilibria emerge is $c - 1$. In the figure below we take $\delta = .8, c = 2$ and $\lambda = 2.5$.[7]

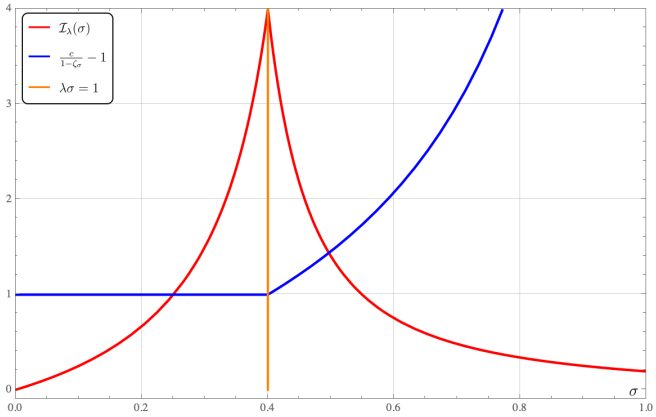

Figure F1: Externalities which decrease in the fraction of the population that adopt have the same qualitative features as our model whenever a viral equilibrium exists.

The figure demonstrates that the characterization in Proposition 4 applies here: importantly at $\lambda = 1$ a viral full-adoption equilibrium emerges. As λ increases there is a threshold $\overline{\lambda}$ at which the unique viral equilibrium transitions to a mixed strategy.

Another natural change to the values in the utility function would be to allow for heterogeneity. If values were unknown but drawn from a common, atomless distribution, say $v_i \sim \mathcal{V}$, then agents' strategies become a function of their privately realized value: $\sigma = \sigma(v_i)$ (Appendix F.1.3 relies on a similar analysis). Fixing any (symmetric) strategy $\sigma(v)$ from other agents, an individual's best response always take on the form of a threshold: if $v > v^*$, take action 1, if $v < v^*$ take action 0. Hence equilibrium takes the form of a threshold rule: all agents with "sufficiently high" values take the action. As such, expected influence can be computed by conditioning the branching process on agents having values at least v^*. As long as this occurs with probability that is independent of agents' degree, it simple scales down the size of the potential adopter network in the same way as λ or σ. For example, if agents play adopt when $v > v^*$ and don't adopt otherwise, then nonviral influence would be given by

$$\frac{\delta\lambda\mathbb{P}(v > v^*)}{1 - \delta\lambda\mathbb{P}(v > v^*)}.$$

[7]These choices of parameter values give us a natural comparison with Figure 4— we could alternatively take $c = 1$ and $f(\zeta_\sigma) = 0.5(1 - \zeta_\sigma)$.

The analysis in this case becomes rather tedious without—as far as we can tell—adding any new qualitative insights, so we view our assumption of homogeneous values as a useful benchmark.

A related but distinct assumption in our model is that influence directly scales v (or more generally, some function of v). If influence were transformed by some function the qualitative features of our model would not change as long as the transformed influence is weakly increasing for $\sigma < \sigma^{\text{crit}}$ and weakly decreasing for $\sigma > \sigma^{\text{crit}}$. So, for example, any positive monotone transformation of influence preserves the qualitative features of our model.

In sum, our model is robust to several alternative utility specifications including diminishing or heterogeneous values and transformations of influence.

F.1.2 No Viral Inference

In our model, agents correctly update their beliefs about their expected degree upon exposure when there is a giant component of potential adopters in the network (this is viral inference). A natural question is whether our results change qualitatively if agents are naive in the sense that they take the degree distribution to the be that of the underlying network regardless of whether they are in the giant component. Since exposure informs an agent that at least one of their links is in the giant component, it stands to reason that viral inference should decrease the expected influence relative to naive beliefs. It is easily shown that expected influence in a naive version of the Poisson model is given by

$$\mathcal{I}_\lambda(\sigma) = \frac{\delta\lambda\sigma\rho_\sigma}{1 - \delta\lambda\sigma\rho_\sigma},$$

so the $-\delta\rho_\sigma$ term in Lemma 2 is the "correction" due to viral inference.

In Figure F2, we plot the expected influence accounting for viral inference, relative to expected influence with naive updating. We plot this only for $\lambda > 1$ since this is the only time viral strategies exist, and we mark $\tilde{\lambda} < \overline{\lambda}$ as the upper thresholds for emergence of viral mixed equilibria under VI and no-VI. We see, as explained in our analysis above, that the expected influence under naive updating lies strictly above the expected influence with viral inference: agents believe they are too influential.

To conclude this section, we emphasize that the introduction of viral inference into the model did not change the qualitative results of our characterization of equilibrium precisely because the expected influence remained qualitatively the same as in the Poisson model. Our analysis of graphs with an arbitrary degree distribution is in many ways analogous (see next section) and we provide further examples illustrating the same point.

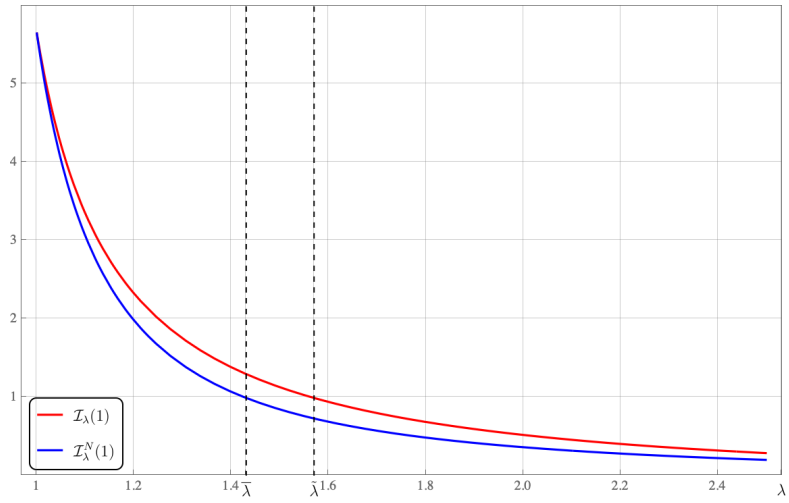

Figure F2: Influence with ($\mathcal{I}_\lambda$) and without ($\mathcal{I}^N_\lambda$) viral inference

F.1.3 Knowledge of Degrees

A critical assumption that simplifies our analysis is that agents do not know their degree. In the settings we are primarily interested in (e.g., participation in a protest), this is a reasonable assumption. In other settings, it may be of interest to know how agents behave when they have knowledge of their degree.

Suppose that agents know their degree. Then, strategies become functions of the degree, i.e., $\sigma\colon \mathbb{N} \to [0,1]$. Since the graph is locally tree-like, higher degrees translate multiplicatively into higher influence. Under degree distribution $\mathcal{D}_\theta$ and strategy *profile* $\boldsymbol{\sigma} \equiv \{\sigma(d)\}_{d\geq 0}$, the expected utility from a prosocial action by an agent of degree d is

$$\mathbb{E}_{\mathcal{D}}[u_\beta(\boldsymbol{\sigma}, d))] = (v + c) + v(d-1)\mathcal{I}_\theta(\boldsymbol{\sigma}).$$

It is straightforward to observe that equilibrium Bayesian strategies will take a threshold form. In the case of a prosocial behavior, agents adopt (not-adopt) with degrees strictly above (below) a threshold degree k^{thres} and may potentially mix at the threshold.[8] Note that agents with $d = 1$ either never (always) adopt in the case of prosocial (antisocial) behaviors. Existence of equilibria can be proved from Brouwer's fixed point theorem, and we can write down an explicit form for the expected influence $I_\theta(\boldsymbol{\sigma})$ in terms of generating functions. The fact that influence is maximized at the critical threshold $\lambda = 1$ suggests that prosocial behaviors emerge continuously and antisocial behaviors emerge discontinuously in this setting as well.

[8] Of course vice-versa in the case of an antisocial behavior.

F.2 A Note on Condition 1

Condition 1 can be weakened to the following.

Condition 2. Let $\sigma \in [0,1]$ be any strategy.

(i) If σ is viral, then $\mathcal{I}_\theta(\sigma)$ is strictly decreasing in σ.

(ii) $\mathcal{I}_\theta(1)$ is strictly decreasing in θ for all $\theta > \theta^{\text{crit}}$.

In Condition 2 we have removed case (i) of Condition 1 since we can show it always holds, and we have replaced the condition that $\mathcal{I}_\theta(\sigma)$ is strictly decreasing in θ for all viral σ, with the condition that this is true at $\sigma = 1$. This amended condition is sufficient for virtually the same equilibrium analysis as we did in the Poisson model, because full-adoption of prosocial behaviors is always an equilibrium around the critical threshold, and (ii) guarantees that as θ get large there is a point where full-adoption is no longer an equilibrium.

F.3 Additional Working for Examples in Section 5.2

F.3.1 Zipf Distribution

It is easily shown that

$$G_1(z) = \frac{1 - e^{-\alpha}}{1 - e^{-\alpha} z}.$$

Thus, the extinction probability $\rho_\sigma = G_1(1 - \sigma + \sigma\rho_\sigma)$ is given by

$$\rho_\sigma = \frac{1}{\theta\sigma}, \tag{F13}$$

where we define $\theta = (e^\alpha - 1)^{-1}$. It follows that the critical threshold θ^{crit} for the emergence of the giant component is simply $\theta = 1$ (or equivalently, $\alpha = \ln(2)$) so that an equilibrium is viral iff $\theta\sigma > 1$.[9]

We will do everything here in terms of α, but it is straightforward to translate all our results in terms of $\theta = (1 - e^\alpha)^{-1}$. In general the generating function for a Zipf distribution is given by

$$G_0(z) = z^{\underline{k}} \Phi(z e^{-\alpha}, 1, \underline{k}).$$

So in the special case where $\underline{k} = 1$, we get

$$G_0(z) = \frac{\ln(1 - e^{-\alpha} z)}{\ln(1 - e^{-\alpha})}.$$

[9] The reason we define $\theta = (e^\alpha - 1)^{-1}$ is so that can write the results here in a similar way to the Poisson model where virality depended on $\lambda\sigma$.

It follows that

$$G_1(z) = \frac{G_0'(z)}{G_0'(1)} = \frac{1-e^{-\alpha}}{1-e^{-\alpha}z},$$

as claimed. The extinction probability ρ_σ under the strategy σ must satisfy

$$\rho_\sigma = \frac{1-e^{-\alpha}}{1-e^{-\alpha}(1-\sigma+\sigma\rho_\sigma)}.$$

We claim that $\rho_\sigma = \frac{e^\alpha-1}{\sigma}$ solves this. To see this, observe that

$$\rho_\sigma\left(1-e^{-\alpha}(1-\sigma+\sigma\rho_\sigma)\right) = -\rho_\sigma^2\sigma e^{-\alpha} + \left(1-e^{-\alpha}(1-\sigma)\right)\rho_\sigma,$$

and so substituting $\rho_\sigma = \frac{e^\alpha-1}{\sigma}$ we get

$$\begin{aligned}
&-\left(\frac{e^\alpha-1}{\sigma}\right)^2(\sigma e^{-\alpha} + \left(1-e^{-\alpha}(1-\sigma)\right)\frac{e^\alpha-1}{\sigma} \\
&\qquad = \frac{1}{\sigma}\left[-(e^{2\alpha}-2e^\alpha+1)e^{-\alpha} + (e^\alpha-1-1+e^{-\alpha}) + \sigma(1-e^{-\alpha})\right] \\
&\qquad = \frac{1}{\sigma}\sigma(1-e^{-\alpha}) \\
&\qquad = 1-e^{-\alpha},
\end{aligned}$$

and so $\frac{e^\alpha-1}{\sigma}$ is a fixed point of the equation $\rho_\sigma = G_1(1-\sigma+\sigma\rho_\sigma)$. It is in fact the smallest solution and therefore the extinction probability.[10]

Importantly, it follows from our analysis that

$$\sigma\rho_\sigma = e^\alpha - 1.$$

To calculate the expected influence, we first calculate

$$G_1'(z) = \frac{e^{-\alpha}(1-e^{-\alpha})}{(1-e^{-\alpha}z)^2} = \frac{e^\alpha-1}{(e^\alpha-z)^2}.$$

[10]We omit the details here but this follows from the fact that it is the smallest root of the quadratic which solves $\rho_\sigma = G_1(1-\sigma+\sigma\rho_\sigma)$.

So at $z = 1 - \sigma + \sigma\rho_\sigma$,

$$\begin{aligned} G_1'(1-\sigma+\sigma\rho_\sigma) &= G_1'(1-\sigma+(e^\alpha-1)) \\ &= G_1'(e^\alpha-\sigma) \\ &= \frac{e^\alpha-1}{(e^\alpha-e^\alpha+\sigma)^2} \\ &= \frac{e^\alpha-1}{\sigma^2} \\ &= \frac{\rho_\sigma}{\sigma}. \end{aligned}$$

Hence $\sigma G_1'(1-\sigma+\sigma\rho_\sigma) = \rho_\sigma$. The last step is to calculate the expected number of neighbors who are potential adopters. We have

$$G_0^{(k)}(1-\sigma) = G_1^{(k-1)}(1-\sigma)G_0'(1) = -\frac{(k-1)!(e^\alpha-1)}{(e^\alpha-(1-\sigma))^k} \times \frac{1}{(e^\alpha-1)\ln(1-e^{-\alpha})},$$

and so

$$\sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma) = -\frac{1}{\ln(1-e^{-\alpha})}\sum_{k=1}^{\infty} k\left[\left(\frac{\sigma}{e^\alpha-(1-\sigma)}\right)^k - \left(\frac{\rho_\sigma\sigma}{e^\alpha-(1-\sigma)}\right)^k\right],$$

Now,

$$\sum_{k=1}^{\infty}\left(\frac{\sigma}{e^\alpha-(1-\sigma)}\right)^k = \frac{\sigma/(e^\alpha-(1-\sigma))}{\left(\frac{\sigma}{e^\alpha-(1-\sigma)}-1\right)^2} = \frac{1}{1-\frac{\sigma}{e^\alpha-(1-\sigma)}} - 1 = \frac{\sigma}{e^\alpha-1}$$

and similarly,

$$\sum_{k=1}^{\infty}\left(\frac{\sigma\rho_\sigma}{e^\alpha-(1-\sigma)}\right)^k = \frac{\sigma/(e^\alpha-(1-\sigma))}{\left(\frac{\sigma\rho_\sigma}{e^\alpha-(1-\sigma)}-1\right)^2} = \frac{1}{1-\frac{\sigma\rho_\sigma}{e^\alpha-(1-\sigma)}} - 1 = \frac{e^\alpha-1}{\sigma}.$$

Finally, we have

$$\frac{1}{1-G_0(1-\sigma+\sigma\rho_\sigma)} = \frac{1}{1-\frac{\ln(e^{-\alpha}\sigma)}{\ln(1-e^{-\alpha})}} = \frac{\ln(1-e^{-\alpha})}{\ln(e^\alpha-1)-\ln(\sigma)}.$$

Making the substitution $\rho_\sigma = \frac{e^\alpha - 1}{\sigma}$ and putting everything together, we have

$$\begin{aligned}\frac{1}{1 - G_0(1-\sigma+\sigma\rho_\sigma)} \sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!} G_0^{(k)}(1-\sigma) &= -\frac{1}{\ln(\rho_\sigma)}\left(\rho_\sigma^{-1} - \rho_\sigma\right) \\ &= \frac{1-\rho_\sigma^2}{-\rho_\sigma \ln(\rho_\sigma)}.\end{aligned}$$

Hence the expected influence when σ is viral (making the substitution $\theta = (e^\alpha - 1)^{-1} = (\sigma\rho_\sigma)^{-1}$) is

$$\begin{aligned}\mathcal{I}_\theta(\sigma) = \frac{\delta\rho_\sigma}{1 - \sigma\delta G_1'(1-\sigma+\sigma\rho_\sigma)}\left[\frac{1-\rho_\sigma^2}{-\rho_\sigma\ln(\rho_\sigma)} - 1\right] &= \frac{\delta\rho_\sigma}{1-\delta\rho_\sigma}\left[\frac{1-\rho_\sigma^2}{-\rho_\sigma\ln(\rho_\sigma)} - 1\right] \\ &= \frac{\delta(1-\rho_\sigma^2)}{-(1-\rho_\sigma)\ln(\rho_\sigma)} - \frac{\delta\rho_\sigma}{1-\delta\rho_\sigma} \\ &= \frac{\delta(\theta\sigma+1)(\theta\sigma-1)}{\theta\sigma(\theta\sigma-\delta)\ln(\theta\sigma)} - \frac{\delta}{\theta\sigma-\delta},\end{aligned}$$

which is precisely the expression given in Section 5.2.1. On the other hand if σ is nonviral, then

$$\mathcal{I}_\theta(\sigma) = \frac{\sigma\delta G_1'(\sigma)}{1-\sigma\delta G_1'(\sigma)} = \frac{\frac{\sigma\delta}{e^\alpha-1}}{1-\frac{\sigma\delta}{e^\alpha-1}} = \frac{1}{\frac{e^\alpha-1}{\sigma\delta} - 1}.$$

We do not attempt to show that Condition 1 holds but it is evidently true from Figure 7.

F.3.2 Exponential Distribution

The generating function for this distribution is given by

$$G_0(z) = \frac{z}{\mu - z(\mu-1)},$$

where $\mu = (1-e^{-\alpha})^{-1}$, and so

$$G_1(z) = \frac{1}{[\mu - z(\mu-1)]^2}.$$

The extinction probability ρ_σ is therefore the solution of a cubic equation and is solvable by radicals. It is easy to show that $G_1'(1) = 2(\mu-1)$, and so subcritical influence (when $\rho_\sigma = 1$) is given by

$$\mathcal{I}_\mu(\sigma) = \frac{4\delta\sigma(\mu-1)}{1-2\delta\sigma(\mu-1)}.$$

It's clear that this expression is increasing in both μ and σ, and moreover that σ is viral if and only if $\sigma > \frac{1}{2(\mu-1)}$. We can also show that, quite remarkably, that

$$\frac{1}{1 - G_0(1-\sigma+\sigma\rho_\sigma)} \sum_{k=1}^{\infty} k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma) = 1 + \frac{2\sigma(\mu-1)}{1+\sigma(\mu-1)(1-\rho_\sigma)}.$$

The extinction probability is the smallest solution in $[0,1]$ to the equation

$$\rho_\sigma = \frac{1}{[\mu - (1-\sigma+\sigma\rho_\sigma)(\mu-1)]^2}.$$

Using computational software (or otherwise), we find that when $\rho_\sigma < 1$, it is given by

$$\rho_\sigma = \frac{2+\sigma(\mu-1)}{2\sigma(\mu-1)} - \frac{\sqrt{4+\sigma(\mu-1)}}{2\sqrt{\sigma(\mu-1)}}.$$

Notice that the key parameter here is $\sigma(\mu-1)$. In fact we can rewrite everything in terms of $\mu = 2\sigma(\mu-1)$ so that

$$\rho_\sigma = \frac{4+\mu}{2\mu} - \frac{\sqrt{8+\mu}}{2\sqrt{\mu}}.$$

The expected forward degree is given by

$$G_1'(1) = 2(\mu-1) > 1 \iff \mu > \frac{3}{2}.$$

Or in terms of α,

$$\frac{1}{1-e^{-\alpha}} > \frac{3}{2} \iff 0 < \alpha < \ln(3).$$

So here, $\alpha = \ln(3)$ is the critical threshold for the emergence of the giant component, which we can more conveniently state as $\mu = 1$. Some work shows

$$\sigma G_1'(1-\sigma+\sigma z) = \frac{2(\mu-1)\sigma}{(\mu-(\mu-1)(1-\sigma+\sigma z))^3}.$$

and substituting ρ_σ into this function gives

$$\sigma G_1'(1-\sigma+\sigma\rho_\sigma) = \frac{16}{\sqrt{(\mu-1)\sigma}(\sqrt{(\mu-1)\sigma}+\sqrt{4+(\mu-1)\sigma})^3}$$

Putting this all together and making the substitution $\theta = 2\sigma(\mu-1)$, we have (after some algebraic manipulation)

$$\frac{\delta\rho_\sigma}{1-\delta\sigma G_1'(1-\sigma+\sigma\rho_\sigma)} = \frac{8\delta}{\sqrt{\theta}\left((4+\theta)\sqrt{\theta}+\theta\sqrt{8+\theta}+4\delta(\sqrt{\theta}-\sqrt{8+\theta})\right)}. \tag{F14}$$

Next we calculate the expected forward degree conditional on exposure under a viral strategy. We have

$$G_0^{(k)}(1-\sigma) = \frac{k!(\mu-1)^{k-1}\mu}{(1+\sigma(\mu-1))^{k+1}}$$

So

$$\begin{aligned}
\sum_k k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}\left[\frac{k!(\mu-1)^{k-1}}{(1+\sigma(\mu-1))^{k+1}}\right] &= \frac{\mu\sigma}{(1+\sigma(\mu-1))^2}\sum_k k(1-\rho_\sigma^k)\left[\frac{\sigma(\mu-1)}{1+\sigma(\mu-1)}\right]^{k-1} \\
&= \frac{\mu\sigma}{(1+\sigma(\mu-1))^2}\left[\frac{1}{\left(1-\frac{\sigma(\mu-1)}{1+\sigma(\mu-1)}\right)^2} - \rho_\sigma\frac{1}{\left(1-\frac{\rho_\sigma\sigma(\mu-1)}{1+\sigma(\mu-1)}\right)^2}\right] \\
&= \mu\sigma\left[1-\frac{\rho_\sigma}{(1+\sigma(\mu-1)(1-\rho_\sigma))^2}\right].
\end{aligned}$$

Next, we have

$$\frac{1}{1-G_0(1-\sigma+\sigma\rho_\sigma)} = \frac{1+\sigma(1-\rho_\sigma)(\mu-1)}{1+\sigma(1-\rho_\sigma)(\mu-1)-(1-\sigma+\sigma\rho_\sigma)} = \frac{1+\sigma(1-\rho_\sigma)(\mu-1)}{\mu\sigma(1-\rho_\sigma)}.$$

Putting everything together, we have

$$\begin{aligned}
\frac{\sum_{k=1}^\infty k(1-\rho_\sigma^k)\frac{\sigma^k}{k!}G_0^{(k)}(1-\sigma)}{1-G_0(1-\sigma+\sigma\rho_\sigma)} &= \frac{1+\sigma(1-\rho_\sigma)(\mu-1)}{\mu\sigma(1-\rho_\sigma)}\cdot\mu\sigma\left[1-\frac{\rho_\sigma}{(1+\sigma(\mu-1)(1-\rho_\sigma))^2}\right] \\
&= \frac{(1-\sigma(1-\rho_\sigma)(\mu-1))^2-\rho_\sigma}{(1-\sigma(1-\rho_\sigma)(\mu-1))(1-\rho_\sigma)} \\
&= \frac{1+2\sigma(\mu-1)+\sigma(\mu-1)(1-\rho_\sigma)}{1+\sigma(\mu-1)(1-\rho_\sigma)} \\
&= 1+\frac{2\sigma(\mu-1)}{1+\sigma(\mu-1)(1-\rho_\sigma)} \\
&= 1+\frac{2\theta}{2+\theta(1-\rho_\sigma)}.
\end{aligned}$$

Hence subtracting 1 for the neighbor in the giant component, the expected number of potential adopting neighbors under an exponential degree distribution is simply $\frac{2\theta}{2+\theta(1-\rho_\sigma)}$. Multiplying (F14) by $\frac{2\theta}{2+\theta(1-\rho_\sigma)}$ gives the desired expression after some further simplification.