

The effects of individual versus community-influenced isolation on SIS epidemic persistence on finite random graphs

Shirshendu Chatterjee*

David Sivakoff†

Matthew Wascher‡

June 1, 2026

Abstract

The contact process, or the SIS epidemic model, is a continuous-time Markov process used to model the spread of a recurring infection on a graph. Each vertex is either healthy or infected, and each infected vertex independently infects each of its healthy neighbors at rate λ and recovers at rate 1. We study the contact process in the presence of additional intervention measures by introducing a third possible state for vertices, which we call isolated. Vertices may enter the isolated state either because of individual decisions or due to community-influenced decisions, which leads to two distinct models that we call the isolation model and the vigilance model, respectively. In the isolation model, infected vertices self-isolate at rate α . In the vigilance model, each healthy vertex causes each of its infected neighbors to isolate at rate α . Unlike the classical contact process, these models lack the key features of attractiveness and existence of a dual, which makes analyzing them more challenging. We study the persistence times of the infection on large, finite, degree-heterogeneous random graphs. We show that the infection in the isolation model persists for at least stretched exponential time in the size of the graph for all values of α and λ . By contrast, in the vigilance model, for every fixed α the persistence time of the infection exhibits a phase transition in λ : for small λ the infection persists for at most a linear time in the size of the graph, while for large λ the infection persists exponentially long. This contrast demonstrates that individual versus community-influenced isolation can substantially affect the persistence of an epidemic.

As a corollary of our main results, we also prove prolonged persistence for the SIRS epidemic model on large finite inhomogeneous random graphs, disproving predictions based on non-rigorous analysis published in the physics literature. In the SIRS model, each vertex is either (i) healthy and susceptible, or (ii) infected, or (iii) removed and temporarily immune. Each infected vertex independently infects each of its healthy neighbors at rate λ and recovers with temporary immunity at rate α , where this immunity is lost at rate 1. We show that for any $\lambda, \alpha > 0$ the persistence time of the SIRS model on the configuration model having n vertices starting from all vertices infected is at least $\exp(n^{1-\eta})$ with high probability for any $\eta > 0$.

1 Introduction

The *contact process*, introduced by [24], is one of the popular models in the context of epidemic spread and it has long been studied as a stochastic epidemic model for recurring infections spreading

*Department of Mathematics, City University of New York, Graduate Center and CCNY

†Departments of Statistics and Mathematics, The Ohio State University

‡Department of Mathematics, Applied Mathematics, and Statistics, Case Western Reserve University

on graphs. Given a graph $G = (V, E)$ with vertices V and edges E , the *classical contact process* allows for two possible states on each vertex, healthy and infected, and has a single parameter λ that determines the rate at which infection spreads across edges from infected to healthy vertices. The contact process has been studied extensively on lattices [30], trees [40], and general random graphs [2].

Recently, there has been increased interest in studying different variants and improvisations of the contact process models that include additional vertex states or edge behaviors as a way to model more complex pathogen dynamics and human behavior. We consider two such models here. The first, the *contact process with isolation*, includes an additional isolated state in which vertices can neither transmit nor receive infection, and has infected vertices self-isolate at rate α . The second, the *contact process with vigilance*, also includes the isolated state but instead of self-isolation, each healthy vertex isolates each of its infected neighbors at rate α . These models are intended to model different ways in which individuals or communities might respond to the presence of infection. An important point of comparison between these models is the survival time of the infection on finite graphs with n vertices. We ask whether there exist distinct parameter regimes (phases) where the survival time is short (at most polynomial in n) versus long (at least super-polynomial in n).

An important difference between our models and the classical contact process is that the classical contact process has a type of monotonicity called *attractiveness* (or equivalently is an attractive particle system) while our models are not known to have this property. This property states that if we have initial infected sets $A \subseteq B$ and ξ_t^A and ξ_t^B are the infected sets at time t beginning from initial infected sets A and B respectively, then there exists a coupling between ξ_t^A and ξ_t^B such that $\xi_t^A \subseteq \xi_t^B$ for all $t \geq 0$. Attractiveness is an important ingredient in the proofs of many results about the classical contact process. Neither the contact process with isolation nor the contact process with vigilance is known to be an attractive particle system, and this presents technical challenges in deriving our results.

We study the contact process with isolation and contact process with vigilance on configuration model random graphs where the degree distribution follows a power law distribution with power law exponent $\gamma > 3$. Our main result on the contact process with isolation shows that this model, like the classical contact process, exhibits at least stretched exponential survival for all choices of λ and α on these graphs. This is somewhat surprising because [6] and [29] previously showed that similar models exhibit polynomial survival on the star graph. By contrast, the classical contact process survives for exponentially long on star graphs, and this property is fundamental to proving that the classical contact process survives exponentially long for all $\lambda > 0$ on high degree random graphs [2, 5, 33].

Our results for the isolation model also apply to another model for epidemic spread and forest fires that incorporates recurring infections and temporary immunity after infections. This model is known as the SIRS epidemic model. Some of its aspects have been studied in the probability literature [31, 32, 9] using rigorous analytic methods and physics literature [38, 43, 17] using non-rigorous approaches. Although the SIRS model is a more natural model for recurring infections compared to the SIS model, only a few sporadic results have been obtained rigorously due to the difficulty posed by lack of monotonicity property described above. To the best of our knowledge, all previously published rigorous results [15, 47] for the SIRS model are for homogeneous graphs and lattices. In this paper, we have been able to make a break through and obtain rigorous result involving the SIRS model on large finite heterogeneous graphs. Our result refutes the claim in the physics literature [17] about having a non-vanishing epidemic threshold for the SIRS model on

heterogeneous random graphs

In contrast with the contact process with isolation and SIRS models, we show that the contact process with vigilance exhibits a phase transition in λ for every fixed α between linear survival time and exponential survival time on graphs satisfying the isoperimetric inequality given in Assumption 2. These graphs include configuration model random graphs where the degree distribution follows a power law distribution with power law exponent $\gamma > 3$ and minimum degree 3. We therefore observe that an individual versus a community-influenced response to an epidemic can have a substantial impact on survival and extinction of the epidemic.

2 Main Results

We define the **contact process with isolation** as follows. Let $G = (V, E)$ be a graph with vertices V and edges E . Let $\xi_t(v)$ denote the state of vertex v at time t , where $\xi_t(v) \in \{0, 1, -1\}$ where state 0 is healthy, state 1 is infectious (or infected), and state -1 is isolating. Let $|\xi_t|$ denote the size of the infected set at time t . Given an initial condition $\xi_0 \in \{0, 1, -1\}^V$ on the states of the vertices in G and parameters λ and α , the process evolves according to the following rules:

1. $\xi_t(v)$ goes from $0 \rightarrow 1$ at rate $\lambda \cdot |\{w \sim v : \xi_t(w) = 1\}|$.
2. $\xi_t(v)$ goes from $1 \rightarrow -1$ at rate α .
3. $\xi_t(v)$ goes from $\{1, -1\} \rightarrow 0$ at rate 1.

Here rate means that the time to event follows an Exponential distribution with the given rate parameter. The contact process with isolation is intended to model an individually-driven response to an epidemic.

We define the **contact process with vigilance** as follows. Let $G = (V, E)$ be a graph with vertices V and edges E . At any given time, each vertex v is in one of three states: 1 (infectious), 0 (healthy), or -1 (isolated). Let $\zeta_t(v)$ be the state of vertex v at time t . Given the initial states of all vertices $\zeta_0 \in \{0, 1, -1\}^V$, an infection rate parameter λ and a vigilance rate parameter α , the process evolves according to the following rules.

1. For v such that $\zeta_t(v) = 0$, $\zeta_t(v) \rightarrow 1$ at rate $\lambda|\{w \sim v : \zeta_t(w) = 1\}|$.
2. For v such that $\zeta_t(v) = 1$, $\zeta_t(v) \rightarrow 0$ at rate 1.
3. For v such that $\zeta_t(v) = 1$, $\zeta_t(v) \rightarrow -1$ at rate $\alpha|\{w \sim v : \zeta_t(w) = 0\}|$.
4. For v such that $\zeta_t(v) = -1$, $\zeta_t(v) \rightarrow 0$ at rate 1.

Again rate means that the time to event follows an Exponential distribution with the given rate parameter. The first rule states that each healthy vertex becomes infected at rate λ times the number of infected neighbors it has. The second rule states each infected vertex recovers and becomes healthy at rate 1. The third rule states that each infectious vertex is identified and isolated at rate α times the number of healthy neighbors it has. The fourth rule states that each isolated vertex becomes healthy and returns from isolation at rate 1. The contact process with vigilance is intended to model a community-driven response to an epidemic.

We now state our main results.

Assumption 1. Suppose $G = (V, E)$ is a configuration model random graph with n vertices where the degree sequence $D_1, \dots, D_n$ is sampled i.i.d. from a power law distribution with

$$\mathbb{P}(D_i \geq k) \asymp \frac{1}{k^{\gamma-1}}$$

with exponent $\gamma > 3$. Moreover, suppose $D_i \geq 3$, so G is connected with high probability.

Assumption 2. Suppose that $G = (V, E)$ is a graph such that for some $0 < \epsilon' < \epsilon < 1/2$ there exists a $\delta > 0$ such that for any set of vertices B satisfying $\epsilon'|V| \leq |B| \leq \epsilon|V|$,

$$|\partial B| \geq \delta|B|,$$

where $\partial B = \{v \in V \setminus B : \exists w \in B \text{ such that } \{w, v\} \in E\}$ denotes the external vertex boundary of B .

Theorem 1. Let $G = (V, E)$ be a graph satisfying Assumption 1 with $|V| = n$, and let $\tau = \inf\{t : |\{\xi_t = 1\}| = 0\}$ be the time to extinction of the contact process with isolation on G . Starting from initial condition $\xi_0(v) = 1$ for all $v \in V$, then for any $\alpha > 0$ and $\lambda > 0$ and $\eta > 0$ we have $\mathbb{P}(\tau < e^{n^{1-\eta}}) \rightarrow 0$ as $n \rightarrow \infty$.

Theorem 2. Let $G = (V, E)$ be any graph with $|V| = n$ and let $\tau = \inf\{t : |\{\zeta_t = 1\}| = 0\}$ be the time to extinction of the contact process with vigilance on G . Fix $\alpha > 0$ and any $\lambda < \alpha$. Starting from initial condition $\zeta_0(v) = 1$ for all $v \in V$, there exists $C > 0$ such that $\mathbb{P}(\tau > Cn) \rightarrow 0$ as $n \rightarrow \infty$.

Theorem 3. Let $G = (V, E)$ be a graph satisfying Assumption 2 for some $\epsilon > 0$ and $\delta > 0$ with $|V| = n$ and let $\tau = \inf\{t : |\{\zeta_t = 1\}| = 0\}$ be the time to extinction of the contact process with vigilance on G . Fix $\alpha > 0$. Starting from initial condition $\zeta_0(v) = 1$ for all $v \in V$, there exists $\lambda_0 = \lambda_0(\alpha, \epsilon, \delta)$ and $C > 0$ such that $\mathbb{P}(\tau < e^{Cn}) \rightarrow 0$ as $n \rightarrow \infty$ for any $\lambda > \lambda_0$.

In Theorem 4, we show that Assumption 2 is satisfied with high probability when G is a graph satisfying Assumption 1 where $\gamma > 3$ and the minimum degree is 3. This leads to the following corollary that gives a direct comparison to Theorem 1 for the contact process with vigilance.

Corollary 1. Let $G = (V, E)$ be a graph satisfying Assumption 1, and let $\tau = \inf\{t : |\{\zeta_t = 1\}| = 0\}$ be the time to extinction of the contact process with vigilance on G . Fix $\alpha > 0$. Starting from initial condition $\zeta_0(v) = 1$ for all $v \in V$, there exists $\lambda_0 = \lambda_0(\alpha, \gamma)$ and $C > 0$ such that $\mathbb{P}(\tau < e^{Cn}) \rightarrow 0$ as $n \rightarrow \infty$ for any $\lambda > \lambda_0$.

The result in Theorem 1 also holds for the SIRS model. We first define the SIRS model as follows. Let $G = (V, E)$ be a graph with vertices V and edges E . Let $\Upsilon_t(v)$ denote the state of vertex v at time t where $\Upsilon_t(v) \in \{0, 1, -1\}$ where state 0 is healthy, state 1 is infections, and state -1 is removed, which is equivalent to isolating in the isolation model. Let $|\Upsilon_t|$ denote the size of the infected set at time t . Given an initial condition $\Upsilon_0 \in \{0, 1, -1\}^V$ on the states of the vertices and parameters λ and α , the process evolves according to the following rules:

1. $\Upsilon_t(v)$ goes from $0 \rightarrow 1$ at rate $\lambda \cdot |\{w \sim v : \Upsilon_t(w) = 1\}|$.
2. $\Upsilon_t(v)$ goes from $1 \rightarrow -1$ at rate α .
3. $\Upsilon_t(v)$ goes from $-1 \rightarrow 0$ at rate 1.

Again, rate means that the time to event follows an Exponential distribution with the given rate parameter. We now formally state the result.

Corollary 2. *Let $G = (V, E)$ be a graph satisfying Assumption 1 with $|V| = n$, and let $\tau = \inf\{t : |\Upsilon_t = 1| = 0\}$ be the time to extinction of SIRS process on G . Starting from initial condition $\Upsilon_0(v) = 1$ for all $v \in V$, then for any $\alpha > 0$ and $\lambda > 0$ and $\eta > 0$ we have $\mathbb{P}(\tau < e^{n^{1-\eta}}) \rightarrow 0$ as $n \rightarrow \infty$.*

To see why this is true, we observe that Lemma 1 also applies to this parameterization of the SIRS model using the same comparison process. That is, given initial condition $\Upsilon_0(V) = \xi_0^R(V)$ and shared parameters λ and α , the analogous coupling demonstrates that

$$\xi_t^R(v) = 1 \implies \Upsilon_t(v) = 1$$

for all $v \in V$ and all $t > 0$. Because our strategy for proving Theorem 1 is to first prove the result for the comparison process and, like the isolation model, the SIRS model stochastically dominates the comparison process, Corollary 2 follows for the SIRS model.

We remark that the key to our proof of Theorem 1 is that the comparison process, which dominates the contact process with isolation and SIRS processes from below, survives for stretched-exponentially long on a “star of stars” subgraph (see Definition 1). As a result, if the power law exponent γ is larger than 2 (rather than 3) we can still show survival for time e^{n^ϵ} for some ϵ depending on λ, α, γ for both the isolation and SIRS models. This follows by identifying a single subgraph isomorphic to a star of stars of order $m = n^\beta$ then following the first part of the proof of Theorem 1. We require $\gamma > 3$ to find $n^{1-\nu}$ such subgraphs (see Lemma 6), though this is likely an artifact of our proof, which can surely be improved to accommodate smaller values of γ .

3 Background and related results

One key question in the study of the contact process concerns the survival time of the infection. Foundational work on this question included studying the contact process on infinite lattices and trees, where it is known that the contact process exhibits a phase transition in λ between almost sure extinction and positive probability of indefinite survival. For a detailed introduction to and summary of results about the contact process, we direct the reader to [30] and [12].

SIS on large finite graphs and extinction-time notions. On any finite graph the SIS/contact process dynamics eventually reaches the absorbing all-healthy state, so the infinite-graph survival notion does not directly carry over. Instead, a central quantity is the *extinction time*, often studied from the worst-case initial condition where all nodes are infected. Early work analyzed extinction-time behavior on homogeneous finite graphs such as d -dimensional tori and grids [14, 35]. More recent results address complex network models including the Barabási–Albert preferential attachment model [1], the configuration model [5, 34, 33, 2], Erdős–Rényi random graphs [4, 37], random regular graphs [36], small-world graphs [13], networks with community structure [44], and general random graph families [46]. In this line of work, a finite-graph system is called *supercritical* (resp. *subcritical*) when the extinction time grows “large” (resp. “small”) with the number of nodes n ; typically “large” means exponential (or stretched-exponential) in n , whereas “small” is closer to logarithmic in n .

A consistent theme is that degree heterogeneity strongly affects persistence. For instance, under extinction-time notions of criticality, it has been shown that $\lambda_c = 0$ for the Barabási–Albert

model [1] and for configuration models with heavy-tailed degree distributions [5]. Moreover, for these heavy-tailed settings the extinction time can be stretched-exponential (and in some cases exponential) in n even for arbitrarily small $\lambda > 0$, with a metastable infected fraction maintained during the long persistence window [34, 33]. In contrast, for degree distributions having exponential moments, recent results show that the corresponding λ_c (in the extinction-time sense) is positive, implying that sufficiently weak infection dies out rapidly regardless of network size [2]; in particular this includes random regular graphs and Erdős–Rényi graphs [36, 4, 37].

Thresholds from small initial prevalence and spectral mean-field baselines. A different and widely used notion of onset asks what happens when the process starts from a small number (or small fraction) of infected nodes. In this setting one again observes two qualitative behaviors: (a) rapid decay to zero (extinction) or (b) growth to a positive metastable prevalence (persistence). A foundational line of work links the onset of endemic SIS behavior to the spectral radius of the adjacency matrix. Early eigenvalue-based arguments and mean-field analyses suggest a threshold of the form $\tau_c \approx 1/\lambda_1(A)$, where $\tau = \beta/\mu$ and $\lambda_1(A)$ denotes the largest eigenvalue (spectral radius) of A [39, 50]. The N-intertwined mean-field approximation (NIMFA) formalizes node-level dynamics that recover the same adjacency-spectral condition in a tractable ODE system [49], and reviews summarize when such approximations are accurate and when they deviate [38]. These static-network baselines are the conditions our multiplex analysis must reduce to when temporal contacts are absent.

For highly heterogeneous networks, however, the effective onset behavior can differ from the simple $\tau_c \approx 1/\lambda_1(A)$ picture, and the probability literature provides complementary perspectives. Using self-duality of the contact process [30], survival from a small seed can be related to metastable prevalence under all-infected initialization. In heavy-tailed random-graph settings, this motivates scaling relations for the probability of long persistence from a small number of initially infected nodes and yields thresholds that may vanish with n (up to logarithmic factors), consistent with the extinction-time viewpoint [5, 34, 33, 2]. When one wishes to make this connection explicit, one may express a survival/metastable-density scaling $\phi(\lambda)$ (for power-law exponent α) and infer how many seeds are needed to trigger persistence; the corresponding network-size scaling is then tied to eigenvalue/maximum-degree scalings in heavy-tailed graphs [18, 8]. (We emphasize that these heavy-tail asymptotics depend on the random-graph model class and exponent regime; we cite them here primarily to highlight that multiple notions of “threshold” coexist in the literature and can behave differently in heterogeneous networks.)

Models with additional vertex and edge behaviors. Recently, there has been increased interest in variant contact process models that include additional vertex and edge behaviors intended to model more complex disease dynamics, behaviors, and population structures. An early example is the adaptive SIS model proposed by [21] in which SI edges randomly rewired. While this model has received considerable attention [20, 22, 41, 11], it has proven difficult to study rigorously, with the aforementioned work focusing on mean field approximations, moment closures, and computational results.

Some variants of the adaptive SIS model have proven more tractable. [6] consider a model in which directed edges temporarily deactivate rather than rewrite and show that this model exhibits a phase transition like that of the classical contact process on $\mathbb{Z}, \mathbb{Z} \bmod n$ and exhibits asymptotic survival time polynomial in the size of the graph on the star graph. [45] study a similar model where bidirectional edges temporarily deactivate and derive an epidemic threshold on the complete graph using differential equation approximations. [26] study a model in which edges rewrite independently

of vertex states and show that this model has a phase transition in λ for certain power law degree distribution graphs and sufficiently fast rewiring.

The evoSIR model, which considers the random rewiring of the adaptive SIS model in the context of an underlying SIR model, has also been studied. In general, SIR models tends to be easier to study than SIS models because in SIR models vertices are removed from the graph upon recovery and so each vertex can become infected at most once. Indeed, more is known about the evoSIR model, including the existence of phase transitions in λ and their continuity behavior on Erdos-Renyi graphs [27] and configuration model graphs [16].

Other variants of the contact process, including the isolation and vigilance models we study in this work, add additional vertex states rather than edge behaviors. [42] studies an ecologically-motivated model that we can translate into the language of our isolation model. In this model all vertices, rather than only infectious vertices, isolate at rate α and return from isolation at rate $\alpha\delta$. He shows this model is monotonic in λ , α , and δ individually, although changing multiple parameters can lead to incomparable processes. He also shows this model has phase transitions on $\mathbf{Z}$ in both λ and δ . For sufficiently small δ , the infection dies out almost surely for all λ , while for larger δ the infection has positive probability of persisting indefinitely only when λ is sufficiently large. We make a more detailed comparison of our isolation model with this model in Section 5.1.

Another related model is the SIRS model, in which after recovering, individuals become temporarily immune to reinfection. Like the isolation model, the SIRS model is not monotonic in λ for reasons analogous to those demonstrated in Figure 4. This has made rigorous study of the SIRS model difficult. However, some results have been derived. [15] show that for sufficiently large λ , the SIRS model has a nontrivial stationary distribution on $\mathbb{Z}^2$. However, their results rely on isoperimetric properties specific to $\mathbb{Z}^2$ and so do not generalize to other graphs.

More recently, [29] showed that the survival time of the SIRS model on the star graph is asymptotically polynomial in the size of the graph. Because the result of [33] that the contact process has survival time $e^{\Theta(n)}$ on power law degree distribution random graphs for all $\lambda > 0$ relies heavily on the fact the contact process has survival time $e^{\Theta(n)}$ on star graphs for all $\lambda > 0$, the results of [29] and [6] pose an interesting question of whether these variant models exhibit qualitatively different behavior from that of the classical contact process on power law degree distribution random graphs. Conversely, [19] show that the expected survival time of the SIRS process is exponential in the size of the graph for certain expander graphs, suggesting one might expect similar behavior on random graphs with power law degree distributions. Theorem 1 answers this question for our isolation model, and by Corollary 2, this also holds for the SIRS model. Very recently, several months after our preprint [7] was posted on arXiv, three independent preprints [28, 25, 23] concerning related aspects of the SIRS model on networks appeared. Some of the results in these works overlap with a corollary of the main results already obtained in [7].

4 Graphical construction

A useful tool for studying the contact process and its variants is a graphical construction sometimes called the Harris construction. For the classical contact process, this construction is defined as follows. For a graph G , consider the spacetime region $G \times [0, \infty)$. At each vertex v and directed edge e of the graph, we will have a Poisson process of temporal marks with appropriate rates, defined as follows.

(C1) For each directed edge $e \in G$, define a Poisson process on $\{e\} \times [0, \infty)$ with intensity λ that




Figure 1: An example of the graphical construction for the classical contact process. Blue vertices are infected, and white vertices are susceptible. We follow the blue paths to determine which vertices are infected at time t .

generates infection arrows.

- (C2) For each vertex $v \in G$, define a Poisson process on $\{v\} \times [0, \infty)$ with intensity 1 that generates recovery dots.

We can then realize the process on $G \times [0, \infty)$ given some initial state $\xi_0 \in \{0, 1\}^V$, where 1 indicates an infected vertex and 0 a healthy vertex, using these marks by doing the following.

1. Label infected vertices 1 (blue in Figure 4) vertically in time until a recovery dot is reached.
2. When an infection arrow is observed, if the source vertex is infected (1, blue), infect the target vertex and repeat step 1. for this newly infected vertex.

The set of vertices in G infected at time t is exactly those vertices reached by a blue path at time t . See Figure 4 for an example. This graphical construction can be used to demonstrate that the classical contact process is attractive. Adding additional vertices to the initial infected set can only increase the number of vertices reached by a blue path at time t .

We can define a similar graphical construction for the contact process with isolation by adding Poisson processes to generate marks that govern isolation transitions.

- (I3) For each vertex $v \in G$, define a Poisson process on $\{v\} \times [0, \infty)$ with intensity α that generates isolation crosses.




Figure 2: In the isolation model adding additional vertices to the initial infected set can lead to a smaller final infected set. Blue vertices are infected, white vertices are susceptible, and red (squiggly, highlighted) vertices are isolated.

We can then realize the process on $G \times [0, \infty)$ given some initial state $\xi_0 \in \{-1, 0, 1\}^V$, where 1 indicates an infected vertex, 0 a healthy vertex, and -1 an isolated vertex using these marks by doing the following.

1. Label infected vertices 1 (blue in Figure 4) vertically in time until a recovery dot or an isolation cross is reached. If a recovery dot is reached first, return the vertex to the health (0) state and remove its color.
2. When an isolation cross is observed on a (1, blue in Figure 4) vertex, change its state to isolated (-1) and color the vertex red vertically in time until a recovery dot is reached. When a recovery dot is reached first, return the vertex to the health (0) state and remove its color.
3. When an infection arrow is observed, if the source vertex is infected (1, blue), infect the target vertex and repeat step 1. and 2. for this newly infected vertex.

Again, the set of vertices in A infected at time t is exactly those vertices reached by a blue path at time t . We can also see from this construction why the argument used to show the classical contact process is attractive fails for the contact process with isolation. Adding additional vertices to the initial infected set may cause additional isolations and larger red regions. As in the example shown in Figure 4, this can lead to a smaller infected set at time t .

5 Isolation model

5.1 Overview

Our goal in the first part of this paper is to prove Theorem 1, and our strategy is to employ an argument used by Chatterjee and Durrett in [5]. They show that the classical contact process

($\alpha = 0$, $\lambda > 0$) survives for time $e^{n^{1-\eta}}$ with high probability by showing that high-degree nodes can sustain the infection for exponentially long in their degree, which is long enough to push the infection to other high-degree nodes. They show that the number of these high-degree nodes whose neighborhoods are highly infected dominate a random walk with positive drift, which leads to stretched-exponential survival on power-law degree distributed random graphs.

There are two main challenges in employing this proof strategy. First, their argument relies crucially on the attractiveness of the classical contact process. The contact process with isolation does not have this property, so our first order of business is to prove that the infected set in the contact process with isolation dominates a comparison process in which all vertices, not just infected vertices, enter isolation at rate α . This process is equivalent to a special case of the model studied by Remenik [42], where his parameter $\delta = 1/\alpha$. Unlike the isolation model, the comparison model is attractive in the same sense as the classical contact process, and Remenik [42] notes this is true of his model generally.

The second main challenge is that the contact process with isolation only survives for a polynomial amount of time on a star graph, which is not necessarily long enough for a high degree node to push the infection to other high degree nodes. This can be proved via the same methods used to study the contact process with avoidance [6] and the SIRS model [29] on stars. To get around this obstacle, we introduce a structure that we call a star of stars of order m , which is essentially a regular tree of depth 2 and degree m . The majority of our proof is then devoted to showing that the comparison process survives for time e^{m^β} with probability at least $1 - e^{-m^\rho}$ on a star of stars of order m . This is a consequence of Lemmas 2 and 5. Once this is established, we show that power-law degree distributed random graphs contain many subgraphs isomorphic to a star of stars of order $m = n^\epsilon$, and our proof is completed along the same lines as in [5].

5.2 Comparison process

We will now formally define the comparison process and establish stochastic domination before briefly discussing why the comparison process is attractive. We define the comparison process as follows. Let $G = (V, E)$ be a graph and denote the state of a vertex v at time t by $\xi_t^R(v) \in \{-1, 0, 1\}$ where state 1 is infectious, state 0 is healthy (and susceptible), and state -1 is isolating. Given an initial state ξ_0^R , this comparison model evolves according to the following rules

1. If $\xi_t^R(v) = 0$, then $\xi_t^R(v) \mapsto 1$ at rate $\lambda \cdot |\{w \sim v : \xi_t(w) = 1\}|$.
2. If $\xi_t^R(v) \in \{0, 1\}$, then $\xi_t^R(v) \mapsto -1$ at rate α .
3. If $\xi_t^R(v) \in \{1, -1\}$, then $\xi_t^R(v) \mapsto 0$ at rate 1.

We now show that the infected set in the comparison process is dominated by the infected set in the contact process with isolation. We note that the analogous statement also holds for the SIRS process.

Lemma 1. *The infected set in the contact process with isolation stochastically dominates the infected set in the comparison process. That is, given initial conditions $\xi_0 = \xi_0^R$ and shared parameters λ and α , there exists a coupling such that*

$$\xi_t^R(v) = 1 \quad \text{implies} \quad \xi_t(v) = 1$$

for all $v \in V$ and all $t > 0$.

Proof. We can demonstrate an appropriate coupling using the graphical construction defined via the Poisson processes of marks (C1), (C2), and (I3) in Section 4.

We can couple the isolation model with the comparison process by generating the same realization of symbols in the graphical constructions for both and considering an arbitrary initial configuration. However, the rules for realizing the process differ for the isolation model and the comparison process. In the isolation model:

1. Color infected (state 1) vertices blue vertically in time until either a recovery dot or isolation cross is reached. If a recovery dot is reached first, return the vertex to the healthy (0) state and remove its color.
2. When an isolation cross is observed on an infected (1, blue) vertex, change its state to isolated (-1) and color the vertex red vertically in time until a recovery dot is reached. When a recovery dot is reached, return the vertex to the healthy (0) state and remove its color.
3. When an infection arrow is observed, if the source vertex is infected (1, blue) and the target vertex is not isolated, infect the target vertex and repeat steps 1.-2. for this newly infected vertex.

We can then follow the blue paths to determine which vertices in G are infected at time t . Note that the path of the infection cannot travel through the red regions indicating isolated vertices.

To construct the comparison process on the same realization of marks in the graphical construction, step 2. above is replaced by 2' below.

- 2'. When an isolation cross is observed on any vertex, color the vertex red vertically in time until a recovery dot is reached, at which time its state reverts to healthy (0) and its color removed.

Thus, any red region that exists in the isolation model will exist in the comparison model, as the isolation cross that initiated the red region in the isolation model will also initiate a red region in the comparison model, and the red regions in both models will continue until they encounter the next recovery dot, which occurs at the same time for both processes. The comparison model may have additional red regions not present in the isolation model.

We now consider that if we were to ignore the red regions, the path of infection in both models will be the same, as they share the same infection arrows and recovery dots. However, we truncate any infection paths that attempt to cross a red region. Thus, if the red regions in the isolation model are contained in the red regions in the comparison model, any infection path in the comparison model that is not truncated by a red region is also not truncated by a red region in the isolation model. This demonstrates a coupling in which under any shared realization of marks in the graphical construction and any initial configuration on $G \times [0, \infty)$, for every vertex v and time t ,

$$\xi_t^R(v) = 1 \implies \xi_t(v) = 1.$$

That is, the set of infected vertices in the comparison model at time t is contained in the set of infected vertices in the isolation model at time t . $\square$

We now observe that the comparison model is attractive. The isolation model is not attractive because additional infections can create additional isolations that might block infection paths that would otherwise successfully persist from time 0 to time t in our graphical construction. However, in

the comparison model, the isolated vertices at any time are determined completely by the isolation crosses and recovery dots. The configuration of infection arrows has no effect on when isolations start and end. Thus, adding more infected vertices to the initially infected set or adding more infection arrows to a realization of the graphical construction can only ever increase the number of infection paths that persist from time 0 to time t . A key consequence that we will use repeatedly is that if the infection persists using only the marks of the graphical construction internal to a subgraph of a graph G , then it persists in that subgraph using the full graphical construction.

5.3 Comparison process on a star of stars

We begin by studying the comparison process on what we call a **star of stars graph of order m** . We define this structure as follows.

Definition 1. *A graph $G = (V, E)$ is called a **star of stars graph of order m** if it is a rooted tree with the following properties.*

- *The root has m children. Call the root the **center** vertex.*
- *Call the m children of the center **hubs**. Each hub has m children, which we call **leaves**.*

We will call the subgraph induced by the center and hubs the **center star** and the subgraph induced by each hub and its leaves a **hub star**. We define the **process internal to a hub star** to mean the realization of the process using only marks in the graphical construction internal to that hub star.

Our goal will be to show that for any $\alpha, \lambda > 0$ there are constants $\beta, \rho > 0$ such that the infection survives on a star of stars of order m for time at least e^{m^β} with probability at least $1 - e^{-m^\rho}$. This is a consequence of Lemmas 2 and 5.

We begin by giving a high-level overview of our proof strategy. If we consider the state of the center, we can classify the process as being in one of three phases at any given time:

- The one-phase, when the center is infected.
- The zero-phase, when the center is healthy.
- The minus one-phase, when the center is isolating.

During the one-phase, the center is spreading the infection to the hubs, potentially reintroducing it to any hub stars where it might have previously gone extinct. During the zero-phase, any hub star can potentially reinfect the center, assuming its own infection has not gone extinct locally. During the minus one-phase, the hub stars behave independently as stars of degree m , and can potentially see their infections go extinct locally.

From this, we can broadly see three reasons the infection might die out on a star of stars.

1. During a one-phase, the center does not infect enough hubs.
2. During a zero-phase, the non-extinct hub stars do not reinfect the center quickly enough.
3. During a minus one-phase, too many hub stars go extinct locally.

Our goal will be to show that each of these is unlikely. However, 1. presents a problem. Namely, the one-phase could be very short, and in this case we would not expect the center to infect many hubs. Thus, we will define long one-phases to be those that last for at least time $1/(1 + \alpha)$ and change item 1. above to

1. During a long one-phase, the center does not infect enough hubs.

Since we are interested in long one-phases, we combine 2. and 3. into

2. Between long one-phases, too many hub stars go extinct locally.

Having set forth our strategy, we now proceed to examine the process during the different center phases.

5.4 The one-phase

Recall that a one-phase occurs when the center becomes infected and ends when the center stops being infected. Let S_i be the time of the i th reinfection of the center, and let T_i be the next time after S_i that the center is not infected. Note that S_i and T_i are stopping times, and on the event $S_i < \infty$, $T_i - S_i$ is exponentially distributed and independent of the σ -field generated by all marks in the graphical construction outside the center during $[S_i, T_i]$. We will call the i th one-phase **long** if $T_i - S_i > 1/(1 + \alpha)$.

For the process to survive, we need the center to spread infection to the hub stars sufficiently so that enough hub stars can survive independently of the center until the next long one-phase. To this end, we will call a hub star **non-extinct** whenever either its hub is infected and or has at least 1 infected leaf.

Our goal in studying the one-phase is to show that there is some $\eta > 0$ such that at the end of a long one-phase, there are at least ηm non-extinct hub stars with high probability irrespective of the configuration at the start of the long one-phase. Since S_i is a stopping time, the strong Markov property allows us to consider the process starting from time 0 with the center infected and show that, conditioned on this initial one-phase being long, the center is likely to spread the infection to enough hub stars. We formalize this in the following lemma.

Lemma 2. *Fix $\alpha, \lambda > 0$ and suppose we start the process at time 0 from an initial configuration ξ_0^R such that the center is infected. Let T_0 be the first time that the center is not infected, let G_0 be the event that $T_0 \geq 1/(1 + \alpha)$, and let N_t be the number of non-extinct hub stars at time $t = 1/(1 + \alpha)$. Then there exists $\eta > 0$ depending on α and λ but not on m such that*

$$\mathbb{P}^{G_0}(N_t < \eta m) \leq \exp\left(\frac{-\eta m}{4}\right)$$

where $\mathbb{P}^{G_0}$ means the probability conditional on the event G_0 .

Proof. Consider a fixed hub. By monotonicity, on the event G_0 we set $T_0 = t$ and observe that the following sequence of events ensures the hub is infected at time T_0 .

- (a) Hub vertex has a recovery dot in $(0, t/3)$
- (b) Hub vertex has no recovery dots in $(t/3, t)$

- (c) Hub vertex has no isolation crosses in $(0, t)$
- (d) Center has an infection arrow to the hub in $(t/3, 2t/3)$

Call the intersection of events (a)–(d) above G_1 . We can use the independence of these events to compute the probability of G_1

$$\mathbb{P}^{G_0}(G_1) = (1 - e^{-t/3})(e^{-2t/3})(e^{-\alpha t})(1 - e^{-\lambda t/3}) =: \gamma_1.$$

Since $t = \frac{1}{1+\alpha}$, we have $\gamma_1 > 0$ depends on α and λ but not on m . This is enough to ensure the hub star is non-extinct at time T_0 . Observe that since we have set $T_0 = t$, N_t is the number of non-extinct hub stars at time T_0 . Then on the event G^0 , N_t stochastically dominates Y where

$$Y \sim \text{Binomial}(m, \gamma_1).$$

Choose $\eta > 0$ depending on λ and α such that $E(Y) \geq 2\eta m$. By a Chernoff bound,

$$\mathbb{P}^{G_0}(N_t < \eta m) \leq \mathbb{P}^{G_0}(Y < \eta m) \leq \exp\left(\frac{-\eta m}{4}\right),$$

which completes the proof. $\square$

5.5 Survival of non-extinct hub stars between consecutive long one-phases

We next consider the survival of the infection in hub stars between consecutive long one-phases. We show an initially non-extinct hub star can survive independently of its surroundings for polynomial time with probability at least $O(1)$. Because the end of the long one-phase is a stopping time, the strong Markov property allows us to consider this starting from time 0.

Lemma 3. *Consider the process internal to a hub star, and let L_t be the number of infected leaves at time t . Suppose $L_0 > 0$, and let $\tau_H = \inf\{t > 0 : L_t = 0\}$. Then there exists $1 > \rho > 0$ and a constant $\kappa > 0$ such that*

$$\mathbb{P}(\tau_H > (2 + 2/\alpha)m^\rho) \geq \kappa.$$

Proof. For our fixed hub vertex, let S_i^H be the time of the i th reinfection of the hub vertex, and let T_i^H be the next time after S_i^H at which the hub is no longer infected. Call the i th time interval $[S_i^H, T_i^H)$ the i th **infectious phase**, provided $S_i^H < \infty$. We say the i th infectious phase is **long** if $T_i^H - S_i^H > 1/(1 + \alpha)$. We first consider the probability that time 1 is the start of a long infectious phase for this hub. We can guarantee this if either the hub is initially infected at time 0 and

- (C1) There are no recovery dots or isolation crosses on the hub in $[0, 1 + 1/(1 + \alpha)]$,

which has probability $e^{-(2+\alpha)}$; or a leaf is infected and

- (C2) There is a recovery dot on the hub in $[0, 1/2]$,
- (C3) There are no recovery or isolation crosses on the leaf in $[0, 1]$,
- (C4) There is an infection arrow from the leaf to the hub in $[1/2, 1]$,

(C5) There are no recovery dots on the hub in $[1/2, 1 + 1/(1 + \alpha)]$, and

(C6) There are no isolation crosses on the hub in $[0, 1 + 1/(1 + \alpha)]$

The probability of satisfying (C2-C6) is

$$K := (1 - e^{-1/2})(e^{-(1+\alpha)})(e^{-\lambda/2})(e^{-(1/2+1/(1+\alpha))})(e^{-\alpha(1+1/(1+\alpha))}) < e^{-(2+\alpha)}.$$

We discount any hubs that fail to satisfy the above conditions and turn our attention to the hubs that have a long infectious phase starting from time 1. Since the comparison process is attractive, we can simplify our argument by henceforth assuming that long infectious phases last for exactly $1/(1 + \alpha)$ time. Broadly, the process internal to the hub star can survive between long one-phases if

1. the hub infects a sufficient number of leaves during a long infectious phase, and
2. a sufficient number of leaves stay infected until the start of the next long infectious phase.

We consider 1. first. Suppose T_i^H is the end of a long infectious phase and let $t = 1/(1 + \alpha)$. By monotonicity we can set $T_i^H = S_i^H + t$ and consider the time interval $[S_i^H, S_i^H + t]$. For each leaf, we can ensure it is infected at time $S_i^H + t$ if the leaf has

- (L1) A recovery dot in $[S_i^H, S_i^H + t/3]$,
- (L2) An infection arrow from the hub in time $[S_i^H + t/3, S_i^H + t]$,
- (L3) No recovery dots in $[S_i^H + t/3, S_i^H + t]$, and
- (L4) No isolation crosses in $[S_i^H, S_i^H + t]$,

The probability of satisfying (L1)-(L4) is

$$p^* = (e^{-t/3})(1 - e^{-2\lambda t/3})(e^{-2t/3})(e^{-\alpha t}).$$

The number of infected leaves at the end of a long infectious phase stochastically dominates X^* , where

$$X^* \sim \text{Binomial}(m, p^*).$$

Let D be the event that there are fewer than $mp^*/2$ infected leaves at time T_i^H . Then using a Chernoff bound,

$$\mathbb{P}(D) \leq \mathbb{P}(X^* < mp^*/2) \leq \exp\left(\frac{-p^*m}{4}\right). \quad (1)$$

We now consider the survival of these infections until the next long one-phase. Let T_ℓ be the time from T_i^H until the next long one-phase. For $C > 0$ to be chosen later, define the following bad events.

- (D1) Consider the first $2C \log m$ time intervals after T_i^H during which the hub is not in state 0. Let T_ℓ^* be the total amount of this time, and let D_1 be the event $T_\ell^* > 4C \log m$. Since $T_\ell^* \sim \text{Gamma}(2C \log m, 1)$, we have

$$\mathbb{P}(D_1) \leq \left(\frac{2}{e}\right)^{C \log m} = m^{-C(1-\log 2)}$$

by a Chernoff bound.

- (D2) Consider the first $2C \log m$ time intervals after T_i^H during which the hub is in state 0. Let T_ℓ^{**} be the total amount of this time, and let D_2 be the event that $T_\ell^{**} > \frac{4C \log m}{\alpha}$. Since T_ℓ^{**} is dominated by a $\text{Gamma}(2C \log m, \alpha)$ random variable, we have

$$\mathbb{P}(D_2) \leq \left(\frac{2}{e}\right)^{C \log m} = m^{-C(1-\log 2)}$$

by a Chernoff bound.

- (D3) There is no leaf that is state 1 at time T_i^H and stays in state 1 until time $T_i^H + 4C \log m + \frac{4C \log m}{\alpha}$. Call this event D_3 .
- (D4) More than $\frac{2\lambda}{\lambda+\alpha} C \log m$ of the next $2C \log m$ time intervals after T_i^H during which the hub is in state 0 are ended by transitions of the hub to state -1 . Call this event D_4 .
- (D5) None of the next $\frac{2\lambda}{\lambda+\alpha} C \log m$ infectious phases after time T_i^H are long or there are fewer than $C \log m$ infectious phases after time T_i^H . Call this event D_5 .

We observe on the event $D^c \cap D_1^c \cap D_2^c \cap D_3^c \cap D_4^c \cap D_5^c$ that $T_\ell \leq 4C \log m + \frac{4C \log m}{\alpha}$. We also observe

$$\begin{aligned} (D^c \cap D_1^c \cap D_2^c \cap D_3^c \cap D_4^c \cap D_5^c)^c &= D \cup D_1 \cup D_2 \cup D_3 \cup D_4 \cup D_5 \\ &\subseteq D \cup D_1 \cup D_2 \cup (D_3 \cap D^c) \cup (D_4 \cap D_3^c \cap D_1^c \cap D_2^c) \\ &\quad \cup (D_5 \cap D_4^c) \end{aligned}$$

First, $(D_3 \cap D^c)$ requires that there are at least $mp^*/2$ infected leaves at time T_i^H . The probability that a leaf that is in state 1 at time T_i^H does not leave state 1 before time $T_i^H + 4C \log m + \frac{4C \log m}{\alpha}$ is $e^{-(4C \log m + \frac{4C \log m}{\alpha})(1+\alpha)}$, and this occurs independently for each leaf that is in state 1 at time T_i^H . Thus,

$$\mathbb{P}(D_3 \cap D^c) \leq \exp \left(- (mp^*/4) (e^{-(4C \log m + \frac{4C \log m}{\alpha})(1+\alpha)}) \right).$$

If we choose C depending on λ and α such that $e^{-(4C \log m + \frac{4C \log m}{\alpha})(1+\alpha)} = m^{-1/2}$, then,

$$\mathbb{P}(D_3 \cap D^c) \leq e^{-p^* \sqrt{m}/4}$$

Next, consider $(D_4 \cap D_3^c \cap D_1^c \cap D_2^c)$. On the event $D_1^c \cap D_2^c$, there are at least $2C \log m$ time intervals during which the hub is in state 0 that occur before time $T_i^H + 4C \log m + \frac{4C \log m}{\alpha}$, and on the event D_3^c there exists some leaf that is in state 1 from time $T_i^H + t$ until time $T_i^H + t +$

$4C \log m + \frac{4C \log m}{\alpha}$. On the event $D_3^c \cap D_1^c \cap D_2^c$, there exists some leaf that is infected continuously through the next $2C \log m$ time intervals during which the hub is in state 0, so during each of these intervals there is probability at least $\frac{\lambda}{\alpha + \lambda}$ that the 0 phase is ended by a transition to state 1. Therefore, the number of 0 phases ended by a transition of the hub to state 1 dominates $Y^* \sim \text{Binomial}(2C \log m, \frac{\lambda}{\alpha + \lambda})$. Thus,

$$\begin{aligned} \mathbb{P}(D_4 \cap D_3^c \cap D_1^c \cap D_2^c) &\leq P(Y^* \leq C \log m) \\ &= P(Y^* \leq E(Y^*)/2) \\ &\leq \exp\left(-\frac{\lambda C \log m}{2(\alpha + \lambda)}\right) \\ &= m^{-\frac{\lambda C}{2(\alpha + \lambda)}}, \end{aligned}$$

by a Chernoff bound.

Finally, consider the event $(D_5 \cap D_4^c)$. On D_4^c there are at least $C \log m$ infectious phases after time T_i^H . Since each infectious phase is long with probability e^{-1} ,

$$\mathbb{P}(D_5 \cap D_4^c) \leq \exp\left(-\frac{e^{-1} C \log m}{4}\right) = m^{-\frac{C}{4e}}$$

We now compute

$$\begin{aligned} &\mathbb{P}((D^c \cap D_1^c \cap D_2^c \cap D_3^c \cap D_4^c \cap D_5^c)^c) \\ &\leq \mathbb{P}(D) + \mathbb{P}(D_1) + \mathbb{P}(D_2) + \mathbb{P}(D_3 \cap D^c) \\ &\quad + \mathbb{P}(D_4 \cap D_3^c \cap D_1^c \cap D_2^c) + \mathbb{P}(D_5 \cap D_4^c) \\ &\leq 5 \max\left(\exp\left(\frac{-p^* m}{4}\right), m^{-C(1-\log 2)}, e^{-p^* \sqrt{m}/4}, m^{-\frac{\lambda C}{2(\alpha + \lambda)}}, m^{-\frac{C}{4e}}\right) \\ &= 5 \max\left(m^{-C(1-\log 2)}, m^{-\frac{\lambda C}{2(\alpha + \lambda)}}, m^{-\frac{C}{4e}}\right) \end{aligned} \tag{2}$$

for large enough m .

If we consider cycles of the process internal to the hub moving between long infectious phases, then equation (2) gives an upper bound on the probability that following a long infectious phase the hub star fails to survive long enough to have a subsequent long infectious phase. Now choose $1 > \rho > 0$ small enough such that

$$(2 + 2/\alpha)m^\rho(1 + \alpha) \leq (1/2) \left(5 \max\left(m^{-C(1-\log 2)}, m^{-C\frac{\lambda}{\alpha + \lambda}}, m^{-\frac{5C}{108}}\right)\right)^{-1}.$$

Then the probability of seeing fewer than $(2 + 2/\alpha)m^\rho(1 + \alpha)$ successful cycles is at most $1/2$ for large enough m . Each cycle must last at least $1/(1 + \alpha)$ time, since the initial long infectious phase lasts at least this much time, and so this guarantees

$$\mathbb{P}(L_{t^*} > 0) \geq K/2 =: \kappa$$

for $t^* = (2 + 2/\alpha)m^\rho$ as desired. □

5.6 Time between long one-phases

We now study the time between consecutive long one-phases. We begin with the following lemma.

Lemma 4. *Suppose at time 0 the process is in the zero-phase and there are h non-extinct hub stars. Then the probability the center is infected at time 1 with no other state changes to the center in the interim is at least $q\mathbf{1}_{h>0}$ where*

$$q = (e^{-3\alpha})(e^{-8/3})(1 - e^{-1/3})(e^{-2\lambda/3}).$$

Proof. We first observe that if $h = 0$, the process is extinct and so $q\mathbf{1}_{h>0} = 0$ as expected. If $h > 0$, then there exists some hub star where either the hub is infected at time 0 or some leaf is infected at time 0. In this case, the following sequence of marks in the graphical construction ensures the center is infected at time 1.

- (E1) None of the center, our chosen hub, or our chosen leaf have an avoidance cross in time $(0, 1)$. This event has probability $e^{-3\alpha}$.
- (E2) Neither the center nor the leaf have a recovery dot in time $(0, 1)$. This event has probability e^{-2} .
- (E3) The hub has a recovery dot in time $(0, 1/3)$. This event has probability $1 - e^{-1/3}$.
- (E4) The hub does not have a recovery dot in $(2/3, 1)$. This event has probability $e^{-2/3}$.
- (E5) The leaf has an infection arrow to the hub in time $(1/3, 2/3)$. This event has probability $1 - e^{\lambda/3}$.
- (E6) The hub has an infection arrow to the center in time $(2/3, 1)$. This event has probability $1 - e^{-\lambda/3}$.

Since events (E1)-(E6) are independent in our graphical construction, we can multiply the probabilities and obtain $q\mathbf{1}_{h>0} = (e^{-3\alpha})(e^{-8/3})(1 - e^{-1/3})(e^{-2\lambda/3})$ as desired. $\square$

Note that the strong Markov property implies that Lemma 4 can also be applied starting from any stopping time T^* .

Lemma 5. *Recall that T_i is the next time after the i th reinfection of the center that the center is no longer infected. Let W^{long} be the amount of time from T_i until the start of the next long one-phase, and let N_t be the number of non-extinct hub stars at time $T_i + t$. Then given the configuration $\xi_{T_i}^R$ is such that $L_0 \geq \eta m$ there exists $\rho > 0$ such that*

$$\mathbb{P}^{\xi_{T_i}^R} \left(W^{\text{long}} \leq (2 + 2/\alpha)m^\rho \right) \geq 1 - 5e^{-Cm^\rho}$$

where κ is the constant from Lemma 3, C is a constant depending on α and λ , and $\mathbb{P}^{\xi_{T_i}^R}$ means probability conditioned on $\xi_{T_i}^R$.

Proof. Let W^{long} be the amount of time from T_i until the next long one-phase. Observe that in the comparison process, every vertex switches infinitely many times between state 0 and states $\{-1, 1\}$; in particular, this is true of the center vertex. For any $\rho > 0$, define the following bad events.

- (B1) Consider the first m^ρ time intervals after time T_i during which the center is in states $\{-1, 1\}$. Let T_g^* be the total amount of this time, and let B_1 be the event that $T_g^* > 2m^\rho$. Since $T_g^* \sim \text{Gamma}(m^\rho, 1)$ for every possible $\xi_{T_i}^R$, we observe

$$\mathbb{P}^{\xi_{T_i}^R}(B_1) \leq \left(\frac{2}{e}\right)^{m^\rho} = e^{-m^{-\rho}(1-\log(2))}$$

by the Chernoff bound.

- (B2) Consider the first m^ρ time intervals after time T_i during which the center is in state 0. Let T_g^{**} be the total amount of this time, and let B_2 be the event that $T_g^{**} > \frac{2m^\rho}{\alpha}$. Since T_g^{**} is dominated by a $\text{Gamma}(m^\rho, \alpha)$ random variable for every possible $\xi_{T_i}^R$, we observe

$$\mathbb{P}^{\xi_{T_i}^R}(B_2) \leq \left(\frac{2}{e}\right)^{m^\rho} = e^{-m^{-\rho}(1-\log(2))}$$

by the Chernoff bound.

- (B3) Let N_t denote the number of hub stars with at least one infected vertex at time t . Let B_3 be the event that $L_t \leq \eta\kappa m/2$ for some $T_i < t \leq T_i + (2 + 2/\alpha)m^\rho$, where κ is the constant from Lemma 3. To compute the probability of B_3 , we observe that by Lemma 3, the number of hub stars that have at least one infected vertex for all times in $[T_i, T_i + (2 + 2/\alpha)m^\rho]$ dominates $X \sim \text{Binomial}(\eta m, \kappa)$, since $L_0 \geq \eta m$ by our assumption on $\xi_{T_i}^R$. Then

$$\mathbb{P}^{\xi_{T_i}^R}(B_3) \leq \exp(-\eta\kappa m/8).$$

- (B4) Let B_4 be the event that fewer than $(q/2)m^\rho$ of the next m^ρ times after T_i at which the center changes from state 0 to a state in $\{-1, 1\}$ result in a one-phase (a change to state 1). Each time the center returns to state 0, it is in state 1 after one unit of time with no other changes of the state of the center in the interim with probability at least $q\mathbf{1}_{\{N_t > 0 \text{ for all } T_i < t < T_i + (2+2/\alpha)m^\rho\}}$ by Lemma 4. On the event $B_1^c \cap B_2^c \cap B_3^c$ we have that $h_t > 0$ for all $T_i < t \leq T_i + (2 + 2/\alpha)m^\rho$ and that the next m^ρ times the center changes from state 0 to states $\{-1, 1\}$ all occur before time $T_i + (2 + 2/\alpha)m^\rho$. Thus, on this event, the number of the next m^ρ time intervals after T_i during which the center is in state $\{-1, 1\}$ that start in state 1 dominates $X \sim \text{Binomial}(m^\rho, q)$ and we have

$$\mathbb{P}^{\xi_{T_i}^R}(B_4 \cap B_1^c \cap B_2^c \cap B_3^c) \leq \exp\left(-\frac{qm^\rho}{8}\right).$$

- (B5) Let B_5 be the event that none of the first $(q/2)m^\rho$ one-phases after time T_i is long.

We observe on the event $B_1^c \cap B_2^c \cap B_3^c \cap B_4^c \cap B_5^c$ that $W^{\text{long}} \leq (2 + 2/\alpha)m^\rho$. We also observe

$$\begin{aligned} (B_1^c \cap B_2^c \cap B_3^c \cap B_4^c \cap B_5^c)^c &= B_1 \cup B_2 \cup B_3 \cup B_4 \cup B_5 \\ &\subseteq B_1 \cup B_2 \cup B_3 \cup (B_4 \cap B_1^c \cap B_2^c \cap B_3^c) \cup B_5. \end{aligned}$$

We can now compute

$$\begin{aligned}
\mathbb{P}^{\xi_{T_i}^R}((B_1^c \cap B_2^c \cap B_3^c \cap B_4^c \cap B_5^c)^c) &\leq \mathbb{P}^{\xi_{T_i}^R}(B_1) + \mathbb{P}^{\xi_{T_i}^R}(B_2) + \mathbb{P}^{\xi_{T_i}^R}(B_3) \\
&\quad + \mathbb{P}^{\xi_{T_i}^R}(B_4 \cap B_1^c \cap B_2^c \cap B_3^c) + \mathbb{P}^{\xi_{T_i}^R}(B_5 \cap B_4^c) \\
&\leq 5 \max \left(e^{-m^\rho(1-\log(2))}, (1/2)^{\eta m}, \exp \left(-\frac{qm^\rho}{8} \right) \right) \\
&\leq 5e^{-Cm^\rho}
\end{aligned}$$

for some constant C , completing the proof. $\square$

5.7 Main result for power law random graphs

To prove Theorem 1, we will first show that the configuration model random graph with power law degree distribution contains many stars of stars of order $m = n^\epsilon$ for some $\epsilon > 0$ with high probability.

Lemma 6. *Let $G = (V, E)$ be a graph with n vertices generated according to a configuration model with power law degree distribution with power law exponent $\gamma > 3$ and minimum degree 3. For any $\nu > 0$ there exists an $\epsilon > 0$ such that with high probability, G contains at least $n^{1-\nu}$ disjoint subgraphs isomorphic to a star of stars of order $m = n^\epsilon$.*

Proof. We will use a first-moment argument to identify a large number of vertices that are centers of stars of stars of order $m = n^\epsilon$ with $\epsilon > 0$ chosen sufficiently small by the end of the proof. Let $\delta > 0$ be small, ultimately depending on ϵ and γ . We first show that the probability that vertex 1 is the center of a star of stars of order m conditional on the event that $D_1 \in [n^{3\gamma\epsilon}, n^{3\gamma\epsilon+\delta}]$ tends to 1 as $n \rightarrow \infty$. Until further notice, we condition on the value of D_1 and assume that it is between $n^{3\gamma\epsilon}$ and $n^{3\gamma\epsilon+\delta}$, and by the law of large number, we also assume the total number of half-edges is $S = \sum_{i \in V} D_i \sim E(D_i)n = \mu n$.

By the law of large numbers, with high probability there are at least $cn^{1-2\epsilon(\gamma-1)}$ vertices among $\{2, \dots, n/2\}$ with degrees $D_i \geq n^{2\epsilon}$. Let V_1 denote the first $cn^{1-2\epsilon(\gamma-1)}$ such vertices. We now consider how many vertices among V_1 are connected to vertex 1, and to do so we attempt to pair only the first n^ϵ half edges from each vertex of V_1 to the half-edges of vertex 1. Let E_1 be the set of edges that get formed by pairing one of the first $n^{3\gamma\epsilon}$ half-edges of vertex 1 to one of the first n^ϵ half-edges of any vertex in V_1 — there are $n^{1-2\gamma\epsilon-\epsilon}$ such half-edges. The number of edges formed in this way, $|E_1|$, dominates a Hypergeometric random variable with a population comprised of $n^{1-2\gamma\epsilon-\epsilon}$ “successes” and $S - n^{1-2\gamma\epsilon-\epsilon} \sim \mu n$ “failures” and sample size $n^{3\gamma\epsilon}$. It follows that $|E_1| \geq \mu^{-1}n^{(\gamma-1)\epsilon}(1 - o(1))$ with high probability.

Now, let N_1 be the set of vertices in V_1 with at least one of its first n^ϵ half-edges paired to one of the first $n^{3\gamma\epsilon}$ half-edges of vertex 1. Given that $|E_1| = \mu^{-1}n^{(\gamma-1)\epsilon}(1 - o(1))$, the paired half-edges are uniformly distributed on all possible pairings of this size. We can therefore select this pairing by first choosing $|E_1|$ half-edges connected to vertex 1, then choosing the $|E_1|$ half-edges connected to vertices of V_1 , then matching them. By selecting the half-edges of E_1 that are connected to V_1 sequentially, we see that until we have selected $|V_1|/2 \gg |E_1|$ half-edges connected to V_1 , at each step we have at least a $1/2$ probability of selecting a half-edge connected to a vertex of V_1

that has not had any half-edges selected previously. It follows that $|N_1| \geq n^{(\gamma-1)\epsilon}/(4\mu)$ with high probability.

For each of the first $n^{2\epsilon} \leq |N_1|$ vertices in N_1 , we now wish to find at least n^ϵ unique neighbors among the vertices $\{n/2 + 1, \dots, n\}$. Let V_2 be the set of vertices among $\{n/2 + 1, \dots, n\}$ that have at least one unpaired half-edge at this point in the process, identify this half-edge with the corresponding vertex in V_2 , and note that $|V_2| \geq n/2 - n^{1-2\gamma\epsilon-\epsilon} - n^{3\gamma\epsilon} \geq n/4$. Each vertex of N_1 has degree at least $n^{2\epsilon}$, and we have observed only n^ϵ of its half edges. Thus, for each vertex of N_1 , we now consider the next $n^{2\epsilon}/2$ of its half-edges that haven't been paired yet. We call a vertex of N_1 successful if at least n^ϵ of its half-edges are paired with half edges of V_2 , and we check sequentially whether each vertex of N_1 is successful. After having paired the first k half-edges attached to vertices of N_1 in this procedure, the next half-edge attached to a vertex in N_1 gets paired to a half-edge of V_2 with probability at least $(n/4 - k)/(\mu n) \geq 1/8\mu$ while $k \leq n/8$. As the total number of half-edges under consideration is at most $n^{4\epsilon} \leq n/8$, we have that each half-edge attached to N_1 has at least probability $1/8\mu$ to connect to V_2 . Therefore, probability that a fixed vertex among the first $n^{2\epsilon}$ of N_1 is successful is at least the probability that a Binomial($n^{2\epsilon}/2, 1/8\mu$) random variable exceeds n^ϵ , which tends to 1. Therefore, the number of successful vertices in N_1 exceeds n^ϵ with high probability. We have thus shown that the conditional probability that vertex 1 is the center of a star of stars of order $m = n^\epsilon$ given that $D_1 \in [n^{3\gamma\epsilon}, n^{3\gamma\epsilon+\delta}]$ is $1 - o(1)$.

Let $H = \{j \in V : D_j \in [n^{3\gamma\epsilon}, n^{3\gamma\epsilon+\delta}]\}$. Then $|H| \asymp n^{1-3\gamma(\gamma-1)\epsilon}$ with high probability. Let $B_4(i)$ denote the set of vertices within distance 4 of i in G . For $j \in H$, let $H_j = H \setminus \{j\}$ and let A_j be the event that j is the center of a star of stars of order $m = n^\epsilon$ and that $j \notin \bigcup_{i \in H_j} B_4(i)$. Then,

$$\begin{aligned} \mathbb{P}(A_j^c) &\leq \mathbb{P}(j \text{ is not a center}) + \mathbb{P}\left(\left|\bigcup_{i \in H_j} B_4(i)\right| \geq n^{3\gamma\epsilon+2\delta}|H|\right) \\ &\quad + \mathbb{P}\left(j \in \bigcup_{i \in H_j} B_4(i), \left|\bigcup_{i \in H_j} B_4(i)\right| < n^{3\gamma\epsilon+2\delta}|H|\right). \end{aligned}$$

We claim that all three terms on the right side tend to 0. The first term tends to 0 by the argument given above. The second term tends to 0 by the law of large numbers, as the size-biased degree distribution has finite mean. The third term tends to 0 because if $\left|\bigcup_{i \in H_j} B_4(i)\right| < n^{3\gamma\epsilon+2\delta}|H| \leq Cn^{1+3\gamma\epsilon+2\delta-3\gamma(\gamma-1)\epsilon}$, the probability that a half-edge connected to j was paired with one of the half-edges within $\bigcup_{i \in H_j} B_4(i)$ is at most $Cn^{1+3\gamma\epsilon+2\delta-3\gamma(\gamma-1)\epsilon} \cdot n^{3\gamma\epsilon+\delta}/n = Cn^{6\gamma\epsilon+3\delta-3\gamma(\gamma-1)\epsilon} \rightarrow 0$ as $n \rightarrow \infty$ for sufficiently small δ , since $\gamma > 3$. Then by Markov's inequality,

$$\mathbb{P}\left(\sum_{j \in H} 1_{A_j^c} \geq |H|/2\right) = o(1), \tag{3}$$

so with high probability the number of disjoint stars of stars of order $m = n^\epsilon$ with centers in H is at least $|H|/2 \asymp n^{1-3\gamma(\gamma-1)\epsilon}$. Choosing ϵ sufficiently small completes the proof. $\square$

We are now ready to prove Theorem 1. Armed with the key ingredients of Lemmas 2, 5 and 6, the proof of the analogue of Theorem 1 for the comparison process follows by a small modification to the proof that the usual contact process survives for (stretched) exponentially long on power-law

random graphs, due to Chatterjee and Durrett [5]. Thus, we sketch the overview of the proof and omit some of the details.

Proof of Theorem 1. By Lemma 6, for any $\nu > 0$ there is an $\epsilon > 0$ such that there are (at least) $N = n^{1-\nu}$ disjoint stars of stars of order $m = n^\epsilon$; denote these subgraphs by $S_1, \dots, S_N$. Also note that the diameter of G is at most $C \log n$ with high probability by Theorem 7.19 of [48]. By Lemma 1 it suffices to prove that the comparison process survives for time $e^{n^{1-2\nu}}$ with high probability.

Let $\beta > 0$ and let $t_k = ke^{m^\beta} = k \exp(n^{\epsilon\beta})$ for $k = 0, 1, \dots$. We say the star of stars S_i begins a **cycle** at time t if at this time its center ends a long one-phase (has been infected for time $1/(1+\alpha)$) and has at least ηn non-extinct hub stars, where η is the constant from Lemma 2. We say that this cycle is **successful** if another cycle begins within time $(2 + 2/\alpha)m^\rho$, where ρ is the constant from Lemma 5.

For $k \geq 1$, we call S_i **hot** at time t_k if it has at least one cycle that begins in $[t_{k-1}, t_k]$ and every one of its cycles that begin in $[t_{k-1}, t_k]$ is successful. Since each cycle must include a long one-phase of the center, the maximum number of cycles during $[t_k, t_{k+1}]$ is less than e^{m^β} . Therefore, if S_i is hot at time t_k , then by Lemmas 2 and 5, the probability that S_i will not be hot again at time t_{k+1} is at most $e^{m^\beta}(e^{-\eta m/4} + 5e^{-Cm^\rho}) \leq e^{-m^\beta}$ for small enough β . Also, observe that since these lemmas concern the dynamics internal to S_i , if S_i is hot at time t_k , then it will (with the same probability) contain an infected vertex throughout the interval $[t_k, t_{k+1}]$.

If S_i is hot at time t_k and S_j is another star of stars, then as above S_i sustains the infection internally through time t_{k+1} and is hot again at time t_{k+1} with probability at least $1 - e^{-m^\beta}$. Now subdivide the time interval $[t_k, t_{k+1}]$ into subintervals of length $C \log n$, and during each subinterval of time, we attempt to “push” the infection from one of the infected vertices in S_i along a fixed shortest-path (of length at most $C \log n$) to the center of S_j . This can be done in a greedy fashion: the infection is pushed from vertex v_ℓ to $v_{\ell+1}$ along this path if in a unit time interval, v_ℓ does not recover or isolate during the entire interval, $v_{\ell+1}$ recovers and/or returns from isolation exactly once, and v_ℓ infects $v_{\ell+1}$ after it is susceptible. Moreover, for a push to be considered a success, we require that the center of S_j enters a long one-phase when it becomes infected at the end of this path. These events each occur with positive probability, and so the probability that one such push results in a long one-phase of S_j is at least $e^{-C \log n} = n^{-C}$, for some $C > 0$. As we make at least $e^{m^\beta/2}$ attempts to push the infection from S_i to S_j , S_j will complete a long one-phase during $[t_k, t_{k+1}]$, and, by Lemmas 2 and 5, will therefore be hot at time t_{k+1} with probability at least $1 - e^{-m^\beta/4}$. By a similar argument, since all vertices are initially infected, the center of S_i has a positive probability of having a long one-phase starting at time 0, so with high probability there are at least βN hot stars of stars at time t_1 .

We define X_k to be the number of stars of stars, $S_1, \dots, S_N$, that are hot at time t_k , and we claim that while $X_k \in [1, N-1]$, the process X_k dominates a random walk process with positive drift. Indeed, since $S_1, \dots, S_N$ are disjoint and the comparison process is attractive, if $X_k \geq 1$, then the number of stars of stars that are hot at time t_k and are not hot at time t_{k+1} is dominated by $Z \sim \text{Binomial}(N, e^{-m^\beta})$. If $X_k \in [1, N-1]$ we can identify an i and j such that S_i is hot at time t_k but S_j is not hot at time t_k , and the probability that all hot stars of stars at time t_k remain hot at time t_{k+1} and that pushing the infection from S_i to S_j makes S_j hot at time t_{k+1} is at least

2/3. Summarizing, the transition probabilities for $X_k \in [1, N-1]$ are bounded as

$$\begin{aligned}\mathbb{P}(X_{k+1} - X_k \geq 1 \mid X_k \in [1, N-1]) &\geq 2/3 \\ \mathbb{P}(X_{k+1} - X_k \leq -z \mid X_k \in [1, N-1]) &\leq \mathbb{P}(\text{Binomial}(N, e^{-m^\beta}) \geq z) \quad \text{for } z \geq 1.\end{aligned}$$

Therefore, X_k dominates a random walk on $\mathbb{Z}$ with increments distributed as $Y - Z$ where $Y \sim \text{Bernoulli}(2/3)$ is independent of Z , until this random walk exits $[1, N-1]$. The moment generating function for an increment of this random walk is

$$\psi(s) = \frac{2e^s - 1}{3} [1 + (e^{-s} - 1)e^{-m^\epsilon}]^N.$$

Observe that $Ne^{-m^\epsilon} \rightarrow 0$ as $n \rightarrow \infty$, so $\psi(s) \rightarrow (2e^s - 1)/3$, and for large n we have $\psi(-1) < 1$. Denoting this random walk by $\hat{X}_k$, it follows that $M_k = \exp(-\hat{X}_k)$ is a nonnegative supermartingale. Therefore, if $X_1 \geq \beta N$ and $K = \inf\{k : \hat{X}_k \leq 0\}$, then $\mathbb{P}(K < \infty) \leq \mathbb{E}(M_K) \leq EM_0 \leq e^{-\beta N} = e^{-\beta n^{1-\eta}}$.

Now we compare X_k to $\hat{X}_k$ each time X_k hits N then jumps to some value less than N . The probability that $X_{k+1} < N/2$ given that $X_k = N$ is at most $e^{-N/3}$, so with high probability no such jumps are made during the first $e^{N/4}$ steps of X_k . Thus, if X_k hits N , we wait for it to reenter $[N/2, N-1]$. Once it does, it dominates $\hat{X}_k$, starting from the same value $\geq N/2$, until the first time $\hat{X}_k$ exits $[1, N-1]$. The probability that $\hat{X}_k$ exits this interval on the left is at most $e^{-N/2}$, and so it will eventually exit on the right, implying X_k returns to N before hitting 0. This implies $X_k \geq 1$ for $k \leq e^{N/4}$ with high probability, and the infection survives for at least time $t_{e^{N/4}} \geq e^{n^{1-2\nu}}$. $\square$

6 Vigilance model

6.1 Overview

Our goal for the vigilance model is to show that for graphs satisfying Assumption 2, for every fixed $\alpha > 0$ there is a phase transition in λ between linear and exponential asymptotic survival regimes. We will study the size of the infected set over time, and so introduce the following notation. Let S_t , I_t , and A_t denote the sets of healthy, infected, and isolated vertices at time t respectively, and let $|S_t|$, $|I_t|$, and $|A_t|$ denote the sizes of these sets.

It is straightforward to show that the infection dies out quickly when $\alpha > \lambda$. If we consider a single SI pair, then the I will become isolated before infecting the S with probability $\frac{\alpha}{\alpha + \lambda} > 1/2$. This suggests that when $\alpha > \lambda$, $|I_t|$ will have negative drift, and we can use a standard random walk comparison to show the process dies out in linear time.

When λ is much larger than $\max(\alpha, 1)$, a single SI pair is more likely to result in an infection of the S vertex rather than an isolation or recovery of the I vertex. This suggests it is possible that $|I_t|$ has positive drift. However, there are two complications. First, we must consider that $|A_t|$ could be large. Second, $|\partial I_t|$ could be small, and combined with large $|A_t|$, lead to a configuration in which the number of SI pairs is smaller than $\lambda|I_t|$. If this is the case, then $|I_t|$ will have negative drift, as the next recovery occurs at rate $|I_t|$ while the next infection occurs at rate λ times the number of SI pairs at time t .

To show such configurations are not typical, we prove Lemma 9 and use Assumption 2. Assumption 2 guarantees that when $|I_t|$ is not too large compared to the size of the graph $|V|$, $|\partial I_t|$

is not too small compared to $|I_t|$. Lemma 9 shows that if we choose λ sufficiently large compared to α , then we can obtain a uniform bound on $|A_t|$ in terms of $|I_t|$ that holds for all $t \leq e^{Cn}$ with high probability, where the constant C depends on α and λ as well as the constants ϵ and δ in Assumption 2.

We can then construct a renewal argument for the survival of the process. We start from all vertices infected and show in Lemma 10 that there is some small $\epsilon^* > 0$ such that whenever $|I_t|$ drops below $(2/3)\epsilon^*n$, so long as our uniform bound on $|A_t|$ holds, the probability that $|I_t|$ hits $(1/3)\epsilon^*n$ before it hits ϵ^*n is e^{-Cn} for some constant C that depends on α and λ as well as the constants ϵ and δ in Assumption 2. Thus, we expect to observe $O(e^{-Cn})$ successful cycles of this renewal, which ensures the infection survives so long as no cycle has failed.

6.2 Isoperimetric assumption

We study the vigilance model on graphs satisfying isoperimetric condition in Assumption 2. Before studying the vigilance model such graphs, we show that a configuration model power law random graph with minimum degree 3 satisfies this assumption with high probability.

Theorem 4. *Let $G = (V, E)$ be a configuration model graph with n vertices and degree sequence $D_1, \dots, D_n$ sampled iid such that*

$$\mathbb{P}(D_i \geq k) \asymp \frac{1}{k^{\gamma-1}}, \quad k = 3, 4, \dots$$

where $\gamma > 3$. Then G satisfies Assumption 2 with probability $1 - o(1)$ as $n \rightarrow \infty$.

Proof. Observe that $\mathbb{E}(D_i) < \infty$ and $\text{Var}(D_i) < \infty$ by our choice of $\gamma > 3$. Define $\bar{D} = \frac{1}{n} \sum_{i=1}^n D_i$. By the Weak Law of Large Numbers, there exists a constant b_0 such that

$$\mathbb{P}(3 < \bar{D} < b_0) \rightarrow 1 \quad \text{as } n \rightarrow \infty$$

and call this event A_0 .

Suppose $1/2 > \epsilon > 0$ and suppose we arbitrarily choose a set U of ϵn vertices. Then if $S(U)$ is the sum of the number of half edges of these vertices,

$$3\epsilon n \leq S(U) \leq \sum_{i=n-\epsilon n}^n D_{(i)}$$

where $D_{(i)}$ is the i th order statistic of $\{D_1, \dots, D_n\}$. To bound $\sum_{i=n-\epsilon n}^n D_{(i)}$, we use the results in [10], which provide a central limit theorem for trimmed means that can be applied to the special case of the upper tail sum. Following their notation, we let F be the cdf of the degree distribution, $q_\epsilon = \sup\{x : F(x) \leq 1 - \epsilon\}$, and define

$$F_{\text{trimmed}}(k) = \begin{cases} 0, & \text{if } k \leq q_\epsilon \\ \frac{F(k) - (1 - \epsilon)}{\epsilon}, & \text{if } k > q_\epsilon. \end{cases}$$

When the variance of F_{trimmed} is finite, which our choice of $\gamma > 3$ guarantees, the results of [10] imply that $S(U)/n \xrightarrow{P} \mu_t$ where μ_t is the mean of F_{trimmed} . Observe $\mu_t = E(D|D \geq q_\epsilon)$. We see

$$\mathbb{P}(D > k) = \sum_{\ell=k+1}^{\infty} \frac{C_1}{\ell^\gamma} \sim \int_k^\infty \frac{C_1}{x^\gamma} dx = \frac{C_1 k^{1-\gamma}}{(\gamma-1)},$$

and setting $\mathbb{P}(D > k) = \epsilon$, we get that $q_\epsilon \propto \epsilon^{-1/(\gamma-1)}$. Thus

$$\mathbb{E}(D|D \geq q_\epsilon) = \frac{\sum_{k=q_\epsilon}^{\infty} k^{-\gamma+1}}{\sum_{k=q_\epsilon}^{\infty} k^{-\gamma}} \sim \frac{\int_{q_\epsilon}^{\infty} k^{-\gamma+1} dk}{\int_{q_\epsilon}^{\infty} k^{-\gamma} dk} = \frac{\gamma-1}{\gamma-2} q_\epsilon.$$

Thus we have

$$\epsilon n \mathbb{E}(D|D \geq q_\epsilon) = C_2 n \epsilon^{(\gamma-2)/(\gamma-1)} (1 + o(1))$$

and so we can choose $k_\epsilon = C \epsilon^{(\gamma-2)/(\gamma-1)}$ for some constant $C > C_2$ such that

$$\mathbb{P}(S(U) \leq k_\epsilon n) \rightarrow 1 \text{ as } n \rightarrow \infty.$$

Let A_1 be the event that $S(U) \leq k_\epsilon n$. Then $\mathbb{P}(A_1) = 1 - o(1)$.

We now count the external vertex boundary of U . Order the half edges in U , $u_1, \dots, u_{S(U)}$. Beginning from u_1 , we pair each half edge of U to another half edge of V . If a half edge u_j has already been matched with another half edge u_i for $i < j$, then we skip the half edge u_j . We can skip over at most $S(U)/2$ half edges, and so we are guaranteed at least $S(U)/2$ matchings. When we match each non-skipped half edge u_i , there are three possible outcomes:

1. u_i matches to u_j for some $j > i$,
2. u_i matches to the half edge of some vertex $v \notin U$ that has a half edge which is already matched to another half edge u_ℓ with $\ell < i$, or
3. u_i matches to the half edge at some vertex $v \notin U$ such that no half edge u_ℓ with $\ell < i$ has previously matched with a half edge at v .

Whenever we observe outcome 3., we increase our count of the external vertex boundary of U by 1. The probability of observing outcomes 1. or 2. is at most $2S(U)/(\bar{D}n - 2S(U))$. If we choose ϵ sufficiently small, then on the event $A_0 \cap A_1$ we have $\bar{D} - 2S(U) \geq 3 - 2k_\epsilon > 2$. Thus, if we let X count the number of non-skipped half edges of vertices in U that result in outcomes 1. or 2., then given $A_0 \cap A_1$, if $S(U) = sn$ with $s \leq k_\epsilon$, we have $X \preceq Y + sn/2$ where

$$Y \sim \text{Binomial}(sn/2, s),$$

and $\mu = \mathbb{E}(Y) = \frac{s^2 n}{2s}$.

Let $k > 1$ (to be chosen later) and let A_2 be the event $Y \leq \frac{sn}{2k}$. Using a Chernoff bound [3]

$$\mathbb{P}\left(Y \leq \frac{sn}{2k} \middle| A_0 \cap A_1\right) \geq 1 - \exp\left(-\frac{sn}{2k} \left(\log \frac{.9\bar{D}}{ks} - 1\right) + \frac{s^2 n}{.9\bar{D}}\right).$$

Let $k' > k$, to be chosen later. Then for sufficiently small $\epsilon > 0$ and $s \leq k_\epsilon$,

$$1 - \exp\left(-\frac{sn}{2k} \left(\log \frac{.9\bar{D}}{k} + \log(1/s) - 1\right) + \frac{s^2 n}{.9\bar{D}}\right) \geq 1 - \exp\left(-\frac{sn}{2k'} \log(1/s)\right). \quad (4)$$

Since we chose $k > 1$, on the event $Y \leq \frac{sn}{2k}$, the external vertex boundary of U is at least size

$$sn - \left(sn/2 + \frac{sn}{2k}\right) = \frac{(k-1)sn}{2k}.$$

Thus, on the event $A_0 \cap A_1 \cap A_2$, the set of vertices U has external vertex boundary at least $\frac{(k-1)sn}{2k} \geq \frac{3(k-1)}{2k}|U|$, since the minimum degree is 3 so that $S(U) \geq 3|U|$.

Now consider that the number of possible sets of vertices of size ϵn is $\binom{n}{\epsilon n}$. Using Stirling's approximation

$$\binom{n}{\epsilon n} \leq \exp \left(1.1\epsilon \log(1/\epsilon)n \right) \quad (5)$$

for all sufficiently large n when ϵ is sufficiently small. On the event A_1 , $S(U) = s$ for some $3\epsilon n \leq s \leq k_\epsilon n$, and the preceding estimates hold uniformly for all sets U and numbers of half edges $S(U)$ in this range. Let A_2^i be the event that A_2 holds for set U^i where $i = 1, \dots, \binom{n}{\epsilon n}$ enumerates all possible sets of size ϵn . Using a union bound and the fact that $s \log(1/s)$ is increasing for $s \in (0, 1/e)$, by choosing ϵ so that $k_\epsilon < 1/e$ we have

$$\mathbb{P}(A_0 \cap A_1 \cap A_2^1 \cap \dots \cap A_2^{\binom{n}{\epsilon n}}) \geq 1 - o(1) - \exp \left(-\frac{k_\epsilon n}{2k'} \log(1/k_\epsilon) + 1.1\epsilon \log(1/\epsilon)n \right).$$

Noting that $k_\epsilon = C\epsilon^{(\gamma-2)/(\gamma-1)} \gg \epsilon$, we can choose ϵ small enough such that $\frac{k_\epsilon}{2k'} \log(1/k_\epsilon) - 1.1\epsilon \log(1/\epsilon) > \epsilon' > 0$, so that

$$1 - \exp \left(-\frac{k_\epsilon n}{2k'} \log(1/k_\epsilon) + 1.1\epsilon \log(1/\epsilon)n \right) \geq 1 - e^{-\epsilon' n}.$$

□

6.3 The subcritical regime

The following lemma establishes the existence of a subcritical regime when $\alpha < \lambda$.

Lemma 7. *Suppose we have any graph $G = (V, E)$ and rates $\lambda, \alpha > 0$ such that $\alpha > \lambda$. Then the epidemic dies out quickly.*

Proof. Recall $|I_t|$ is the count of infected vertices at time t . We will show that $|I_t|$ is stochastically dominated by a random walk that is biased toward 0.

Starting at any time t there are four possibilities for the next event: infection, recovery, isolation, and return from isolation. Infection increases $|I_t|$ by 1, recovery and isolation decrease $|I_t|$ by 1, and return from isolation does not change $|I_t|$ but increases the number of SI edges in the system. We can observe

1. The rate of the next infection is $\lambda(\#SI \text{ edges})$
2. The rate of the next recovery event is $|I_t|$
3. The rate of the next isolation event is $\alpha(\#SI \text{ edges})$

The the ratio of the rate of events that increase $|I_t|$ by 1 to the rate of all events that change $|I_t|$ is

$$\frac{\lambda(\#SI \text{ edges})}{\lambda(\#SI \text{ edges}) + |I_t| + \alpha(\#SI \text{ edges})} \leq \frac{\lambda}{\lambda + \alpha}$$

where the right-hand side does not depend on the number of SI edges. Thus, the probability that the next event that changes $|I_t|$ increases $|I_t|$ by 1 is at most $\frac{\lambda}{\lambda+\alpha}$ and the probability that the next event that changes $|I_t|$ decreases $|I_t|$ by 1 is at least $\frac{\alpha}{\lambda+\alpha}$. Thus, we can stochastically dominate $|I_t|$ via coupling with a continuous time process W_t such that $W_0 = |I(0)|$ and W_0 evolves according to the following rules:

1. For all t such that $|I_t| > 0$, W_t transitions when $|I_t|$ transitions.
2. For all t such that $|I_t| = 0$, W_t transitions at rate 1.
3. When W_t transitions, it transitions to $W_t + 1$ with probability $\frac{\lambda}{\lambda+\alpha}$ and transitions to $W_t - 1$ with probability $\frac{\alpha}{\lambda+\alpha}$.

Let $\tau_0 = \inf\{t : W_t = 0\}$ and observe that $|I_t| \leq W_t$ for all $t \leq \tau_0$. In addition, observe that W_t is a random walk biased toward 0, and so if $\tau_0 = \inf\{t : W_t = 0\}$, there exists $C > 0$ such that $P(\tau_0 > Cn) \rightarrow 0$ as $n \rightarrow \infty$. $\square$

6.4 The supercritical regime

Consider the vigilance model with infection rate λ and vigilance rate α . Define the continuous time process $V_t = |A_t| - \eta|I_t|$. We study the behavior of this process in Lemma 8 and use the results to prove Lemma 9.

Lemma 8. *Choose $\eta < \epsilon\delta/24$ where ϵ and δ satisfy Assumption 2 for G and $\delta \leq 12$. Choose $\lambda > \frac{2\alpha}{\eta}$ and let $V_t = |A_t| - \eta|I_t|$. Suppose V_0 is in the interval $(2\eta n, 2.1\eta n)$. Then $|V_t|$ hits $[0, \eta n]$ before it hits $(3\eta n, \infty)$ with probability at least $1 - \exp\left(\frac{-\eta^2 n}{43}\right)$.*

Proof. Define the jump process $\{(\hat{S}_k, \hat{I}_k, \hat{A}_k) : k \in \mathbb{Z}_{\geq 0}\}$ to be the embedded discrete time Markov chain of (S_t, I_t, A_t) , which increments whenever an event (infection, isolation or recovery) occurs. Let p_k be the probability that the next event in this process is either an infection or an isolation. Then

$$p_k = \frac{(\alpha + \lambda)(\#SI \text{ edges in } (\hat{S}_k, \hat{I}_k, \hat{A}_k))}{(\alpha + \lambda)(\#SI \text{ edges in } (\hat{S}_k, \hat{I}_k, \hat{A}_k)) + |\hat{A}_k| + |\hat{I}_k|}.$$

Define

$$U_k = |\hat{A}_k| - \eta|\hat{I}_k|,$$

and observe that since $\eta < 1/4$ and $\lambda > \frac{2\alpha}{\eta} > 8\alpha$, while $U_k > \eta n$

$$\begin{aligned} \mathbb{E}(U_{k+1} - U_k | (\hat{S}_k, \hat{I}_k, \hat{A}_k)) &= p_k \left(\frac{\alpha}{\alpha + \lambda} (1 + \eta) - \eta \frac{\lambda}{\lambda + \alpha} \right) + (1 - p_k) \left(\frac{-(|\hat{A}_k| - \eta|\hat{I}_k|)}{|\hat{A}_k| + |\hat{I}_k|} \right) \\ &\leq p_k \left(\frac{\alpha}{\lambda} (1 + \eta) - \eta \frac{8\alpha}{\alpha + 8\alpha} \right) + (1 - p_k) \left(\frac{-U_k}{n} \right) \\ &\leq p_k \left((1/2)\eta(1 + \eta) - (8/9)\eta \right) + (1 - p_k) \left(-\eta \right) \\ &\leq -\frac{1}{4}\eta. \end{aligned}$$

Let $\mathcal{K} = \inf\{k : U_k \leq \eta n\}$ be the first time that U_k hits $[0, \eta n]$. Now define the discrete time process W_k by

$$W_k = U_{k \wedge \mathcal{K}} + \frac{1}{4}\eta(k \wedge \mathcal{K})$$

and note that W_k is a supermartingale. Let time $t = 0$ correspond to increment $k = 0$ in the embedded discrete process $(\hat{S}_k, \hat{I}_k, \hat{A}_k)$, so that $W_0 = U_0 = V_0$. Note that for all k , $|W_{k+1} - W_k| \leq 2$, so we can apply Azuma's inequality to see that

$$\mathbb{P}(W_k - W_0 > .9\eta n) \leq \exp\left(\frac{-(.81)\eta^2 n^2}{8k}\right).$$

Observe that when $W_k - W_0 \leq .9\eta n$ holds, we have

$$\begin{aligned} U_k + \frac{1}{4}\eta k - U_0 &\leq .9\eta n \\ U_k &\leq U_0 + .9\eta n - \frac{1}{4}\eta k \end{aligned}$$

Suppose $W_k - W_0 \leq .9\eta n$ holds for $k = 1, \dots, \mathcal{K}$. This ensures that U_k hits the interval $[0, \eta n]$ before it hits the interval $(3\eta n, \infty)$. It also ensures we must have $\mathcal{K} - 1 \leq 8n$, otherwise

$$U_{\mathcal{K}-1} \leq U_0 - 2\eta n + .9\eta n < \eta n$$

which is a contradiction since $\mathcal{K} = \inf\{k : U_k \leq \eta n\}$. Thus we can observe

$$\begin{aligned} \mathbb{P}(\mathcal{K} > 8n + 1) &\leq P(W_k - W_0 > .9\eta n \text{ for some } k = 1, \dots, m) \\ &\leq \sum_{k=1}^{\mathcal{K}} \exp\left(\frac{-(.81)\eta^2 n^2}{8k}\right) \\ &\leq \sum_{k=1}^{8n+1} \exp\left(\frac{-(.81)\eta^2 n^2}{8k}\right) \\ &\leq (8n + 1) \exp\left(\frac{-\eta^2 n}{40}\right) \\ &\leq \exp\left(\frac{-\eta^2 n}{41}\right) \text{ for large enough } n. \end{aligned}$$

So long as U_k does not hit the interval $(3\eta n, \infty)$, V_t does not hit the interval $(3\eta n, \infty)$, and when U_k hits the interval $(0, \eta n]$, V_τ hits the interval $(0, \eta n]$. Thus for large enough n , with probability at least $1 - \exp\left(\frac{-\eta^2 n}{41}\right)$, after V_t hits the interval $(2\eta n, 2.1\eta n)$, it enters the interval $(0, \eta n]$ and does so before visiting $(3\eta n, \infty)$. □

Lemma 9. *Chose $\eta < \epsilon\delta/24$ where ϵ and δ satisfy Assumption 2 for G and $\delta \leq 12$. Choose $\lambda > \frac{2\alpha}{\eta}$ and let $V_t = |A_t| - \eta|I_t|$. Suppose V_0 is in the interval $(2\eta n, 2.1\eta n)$. Let $\tau = \inf\{t > 0 : V_t > 3\eta n\}$. Then there is $c > 0$ (depending on λ, α, η) such that $\mathbb{P}(\tau > e^{cn}) \geq 1 - e^{-cn}$.*

Proof. Recall that V_0 is in the interval $(2\eta n, 2.1\eta n)$ and let τ_1 be the first time V_t leaves $(\eta n, 3\eta n)$. We can observe τ_1 is bounded below by the time it takes for $(.9\eta n)/(1 + \eta) \geq \eta n/2$ transitions that change V_t to occur, since any transition can change the value of V_t by at most $1 + \eta$. The possible transitions are infections, isolations, recoveries of infected vertices, and recoveries of isolated vertices, and these happen at combined rate

$$\begin{aligned} & (\lambda + \alpha)(\# \text{number SI edges}_t) + |A_t| + |I_t| \\ & \leq (\lambda + \alpha) \binom{n}{2} + 2n \\ & \leq (\lambda + \alpha)n^2 \text{ for large enough } n \end{aligned} \tag{6}$$

Therefore the amount of time for the continuous time process V_t to hit $(0, \eta n]$ starting from $(2\eta n, 2.1\eta n)$ stochastically dominates the random variable $X \sim \text{Gamma}(\eta n/2, (\lambda + \alpha)n^2)$. Thus using a Chernoff bound

$$\mathbb{P}\left(\tau_1 < \frac{\eta/2}{(\lambda + \alpha)n^2}\right) \leq \mathbb{P}\left(X < \frac{\eta/2}{(\lambda + \alpha)n^2}\right) \leq \exp(-n).$$

Now consider cycles of V_t hitting the interval $(2\eta n, 2.1\eta n)$ and then hitting either $(0, \eta n]$ or $(3\eta n, \infty)$. Say that a cycle fails if either after hitting the interval $(2\eta n, 2.1\eta n)$ V_t hits $(3\eta n, \infty)$ before it hits $(0, \eta n]$ or V_t leaves $(\eta n, 3\eta n)$ in less time than $\eta n/2$. Using a union bound, the probability of failure on a cycle is at most

$$\exp\left(\frac{-\eta^2 n}{41}\right) + \exp(-n)$$

and thus the probability of any failures in $\exp\left(\frac{\eta^2 n}{84}\right)$ cycles is at most

$$\begin{aligned} & \exp\left(\frac{\eta^2 n}{84}\right) \left(\exp\left(\frac{-\eta^2 n}{41}\right) + \exp(-n) \right) \\ & \leq \exp\left(\frac{-\eta^2 n}{84}\right) \text{ for large enough } n \end{aligned} \tag{7}$$

Therefore, if we observe no failures in $\exp\left(\frac{\eta^2 n}{84}\right)$ cycles we will have that if V_t begins below $2.1\eta n$, it will remain below $3\eta n$ for at least time

$$\begin{aligned} & \frac{\eta/2}{(\lambda + \alpha)n^2} \exp\left(\frac{\eta^2 n}{84}\right) \\ & \geq \exp\left(\frac{\eta^2 n}{85}\right) \text{ for large enough } n \end{aligned} \tag{8}$$

and this happens with probability at least $1 - \exp\left(\frac{-\eta^2 n}{84}\right)$. $\square$

We now prove our main result.

Lemma 10. Suppose we consider the vigilance model on a graph $G = (V, E)$ where $|V| = n$ satisfying Assumption 2. Fix $\alpha > 0$ and let $T_k = \inf\{t : |I_t| = k\}$. Then there exist ϵ and η such that if $|I_0| = (2/3)\epsilon n$, $|A_0| < 2\eta n$, and $\lambda > \lambda_0$ where λ_0 depends on α , ϵ , and η , then $\mathbb{P}(T_{\epsilon n} < T_{(1/3)\epsilon n}) \geq 1 - \exp\left(\frac{-\epsilon\eta^2 n}{85}\right)$

Proof. Start by fixing $\alpha > 0$, choosing $\eta < \epsilon\delta/24$ where ϵ and δ satisfy Assumption 2 for G , $\delta \leq 12$, and $\lambda_0 = \max\{\frac{2\alpha}{\eta}, \alpha + 9/\delta\}$.

From Lemma 9 we have that with probability at least $1 - \exp\left(\frac{-\eta^2 n}{84}\right)$, $V_t = |A_t| - \eta|I_t| \leq 3\eta n$ for all $t \leq \exp\left(\frac{\eta^2 n}{85}\right)$. Thus for $t < \min\{\exp\left(\frac{\eta^2 n}{85}\right), T_{(1/3)\epsilon n}, T_{\epsilon n}\}$, we have $|I_t| \in ((1/3)\epsilon n, \epsilon n)$ and $V_t \leq 3\eta n$, which together imply

$$\begin{aligned} |A_t| &\leq (3 + \epsilon)\eta n \\ &\leq 4\eta n. \end{aligned} \tag{9}$$

From our initial condition on $|I_0|$, the definitions of $T_{(1/3)\epsilon n}$ and $T_{\epsilon n}$, and Assumption 2, we can observe that $|\partial I_t| \geq \delta(\epsilon/3)n$ for all $t < \min\{\exp\left(\frac{\eta^2 n}{85}\right), T_{(1/3)\epsilon n}, T_{\epsilon n}\}$. We can further observe that since $|A_t| \leq 4\eta n$,

$$\begin{aligned} |\partial I_t \cap S_t| &\geq \delta(\epsilon/3)n - 4\eta n \\ &\geq \delta\epsilon n/6 \end{aligned} \tag{10}$$

by our choice of η .

Now consider events that change $|I_t|$. $|I_t|$ decreases by 1 when either an infected vertex recovers or becomes isolated, and I_t increases by 1 when a susceptible vertex becomes infected. Thus, for $t < \min\{\exp\left(\frac{\eta^2 n}{85}\right), T_{(1/3)\epsilon n}, T_{\epsilon n}\}$ the probability that the next event that changes $|I_t|$ increases it by 1 rather than decreases it by 1 is

$$\begin{aligned} &\frac{(\lambda - \alpha)|\partial I_t \cap S_t|}{(\lambda - \alpha)|\partial I_t \cap S_t| + |I_t|} \\ &\geq \frac{(\lambda - \alpha)\epsilon\delta n/6}{(\lambda - \alpha)\epsilon\delta n/6 + \epsilon n} \\ &\geq 3/5 \end{aligned} \tag{11}$$

by our choice of $\lambda > \lambda_0 \geq \alpha + 9/\delta$.

Recall the embedded discrete time process $(\hat{S}_k, \hat{I}_k, \hat{A}_k)$. For all increments k that occur before time $\min\{\exp\left(\frac{\eta^2 n}{85}\right), T_{(1/3)\epsilon n}, T_{\epsilon n}\}$, $|\hat{I}_k|$ stochastically dominates a random walk R_k that starts at $R_0 = (2/3)\epsilon n$, increases by 1 with probability 3/5 and decreases by 1 with probability 2/5 each increment. We continue this coupling until one of the following occurs:

1. $|\hat{I}_k|$ hits ϵn
2. $|\hat{I}_k|$ hits $(1/3)\epsilon n$
3. k hits K such that increment K occurs after time $\frac{18n}{(\alpha + \lambda)\delta}$

after which we let R_k continue independently of $|\hat{I}_k|$. We can observe that if all of the following hold

1. $R_{2\epsilon n} \geq \epsilon n$
2. R_k hits ϵn before it hits $(1/3)\epsilon n$
3. $K > 2\epsilon n$

then our coupling ensures that $|\hat{I}_k|$ hits ϵn before it hits $(1/3)\epsilon$ and this happens for some increment $k < \exp\left(\frac{\eta^2 n}{85}\right)$. This in turn implies that $T_{\epsilon n} \leq \min\{\exp\left(\frac{\eta^2 n}{85}\right), T_{(1/3)\epsilon n}\} \leq T_{(1/3)\epsilon n}$. We will bound each of failure of each of these events.

1. Applying a Binomial Chernoff bound, we observe that

$$\mathbb{P}(R_{2\epsilon n} < \epsilon n) < \exp\left(\frac{-\epsilon^2 n^2}{225}\right).$$

2. The probability that R_k hits $(1/3)\epsilon n$ before it hits ϵn is given by a standard Gambler's ruin analysis:

$$\mathbb{P}(R_k \text{ hits } (1/3)\epsilon n \text{ before it hits } \epsilon n) = \frac{1 - (3/2)^{(1/3)\epsilon n}}{1 - (3/2)^{(2/3)\epsilon n}} \leq \exp\left(\frac{-\epsilon n}{9}\right).$$

3. Observe that $|\hat{I}_k|$ changes when either an infection occurs, an isolation occurs, or an infected vertex recovers. This happens at combined rate

$$\begin{aligned} &(\alpha + \lambda)(\# \text{ SI edges}) + |\hat{I}_k| \\ &\geq (\alpha + \lambda)(\delta \epsilon n / 6). \end{aligned} \tag{12}$$

Thus the number of increments k that occur in time $\frac{18n}{(\alpha + \lambda)\delta}$ stochastically dominates a random variable $Y \sim \text{Poisson}(3\epsilon n)$. Applying a Chernoff bound for Y

$$\mathbb{P}(K < 2\epsilon n) \leq \mathbb{P}(Y < 2\epsilon n) \leq \exp\left(\frac{-\epsilon n}{4}\right) \tag{13}$$

Therefore we can bound the probability of failure by

$$\begin{aligned} \mathbb{P}(T_{(1/3)\epsilon n} < T_{\epsilon n}) &\leq \mathbb{P}(R_{2\epsilon n} < \epsilon n) + \mathbb{P}(R_k \text{ hits } (1/3)\epsilon n \text{ before it hits } \epsilon n) \\ &\quad + \mathbb{P}(K < 2\epsilon n) + \mathbb{P}(\text{Lemma 9 fails}) \end{aligned} \tag{14}$$

We have already established $\mathbb{P}(R_{2\epsilon n} < \epsilon n) < \exp\left(\frac{-\epsilon^2 n^2}{225}\right)$ and $\mathbb{P}(R_k \text{ hits } (1/3)\epsilon n \text{ before it hits } \epsilon n) \leq \exp\left(\frac{-\epsilon n}{9}\right)$ and $\mathbb{P}(K < 2\epsilon n) < \frac{-\epsilon n}{4}$ and $\mathbb{P}(\text{Lemma 9 fails}) < \exp\left(\frac{-\eta^2 n}{84}\right)$.

Substituting these into equation 14, we have

$$\begin{aligned}\mathbb{P}(T_{\epsilon n} < T_{(1/3)\epsilon n}) &\geq 1 - \left(\exp\left(\frac{-\epsilon^2 n^2}{225}\right) + \exp\left(\frac{-\epsilon n}{9}\right) + \exp\left(\frac{-\epsilon n}{4}\right) + \exp\left(\frac{-\eta^2 n}{84}\right) \right) \\ &\geq 1 - \exp\left(\frac{-\epsilon \eta^2 n}{85}\right)\end{aligned}\tag{15}$$

for large enough n . $\square$

We are now ready to prove Theorem 3.

Proof. Choose some $\epsilon > 0$ and $\delta \in (0, 12)$ for which G satisfies Assumption 2, and fix $\alpha > 0$. Start from all vertices infected so $|I_0| = n$, let $\tau_1 = \inf\{t : |I_t| = (2/3)\epsilon n\}$, and let $T_k^1 = \inf\{t \geq \tau_1 : |I_t| = k\}$. For $i = 2, 3, \dots$, let $\tau_i = \inf\{t > \tau_{i-1} : |I_t| = (2/3)\epsilon n\}$ and let $T_k^i = \inf\{t \geq \tau_i : |I_t| = k\}$. By Lemma 10, if we choose $\eta < \epsilon\delta/24$ and $\lambda > \lambda_0 = \max(\frac{2\alpha}{\eta}, \alpha + 9/\delta)$ then

$$\mathbb{P}(T_{\epsilon n}^1 < T_{(1/3)\epsilon n}^1) \geq 1 - \exp\left(\frac{-\epsilon \eta^2 n}{86}\right).$$

We can also observe for each $i = 1, 2, \dots$ that if $T_{\epsilon n}^i > T_{(1/3)\epsilon n}^i$ this guarantees $\tau_{i+1} < \infty$ because $\lim_{t \rightarrow \infty} |I_t| = 0$ and to reach 0 from ϵn , $|I_t|$ must pass through $(2/3)\epsilon n$. We can thus conclude

$$\{T_{\epsilon n}^j < \infty\} \supseteq \{T_{\epsilon n}^j < T_{(1/3)\epsilon n}^j\} \supseteq \{T_{\epsilon n}^j < T_{(1/3)\epsilon n}^j | T_{\epsilon n}^{j-1} < T_{(1/3)\epsilon n}^{j-1}\} \cap \dots \cap \{T_{\epsilon n}^1 < T_{(1/3)\epsilon n}^1\}$$

and use the strong Markov property to apply Lemma 10 to see that for each $i = 1, 2, \dots$

$$\mathbb{P}(T_{\epsilon n}^i > T_{(1/3)\epsilon n}^i | T_{\epsilon n}^{i-1} > T_{(1/3)\epsilon n}^{i-1}) \geq 1 - \exp\left(\frac{-\epsilon \eta^2 n}{86}\right).$$

Taking a union bound on the probability of failure of Lemma 10 for $i = 1, \dots, j$, we see that

$$\mathbb{P}(T_{\epsilon n}^j < \infty) \geq \mathbb{P}(T_{\epsilon n}^j > T_{(1/3)\epsilon n}^j) \geq 1 - j \exp\left(\frac{-\epsilon \eta^2 n}{86}\right).$$

If we choose $j = \exp\left(\frac{\epsilon \eta^2 n}{172}\right)$ then this probability is $1 - \exp\left(\frac{-\epsilon \eta^2 n}{172}\right)$, so with high probability we will have $T_{\epsilon n}^j < \infty$ for $j = \exp\left(\frac{\epsilon \eta^2 n}{172}\right)$.

To obtain a lower bound on $T_{\epsilon n}^j$, consider that for each $i = 1, \dots, j$ there must be at least one event that involves a vertex changing states between $T_{\epsilon n}^{i-1}$ and $T_{\epsilon n}^i$. Given our graph $G = (V, E)$, the total rate at which any such event occurs is bounded above by $(\lambda + \alpha)|E| + |V|$. Since $|V| = n$ and $|E| < n^2$, the total rate is bounded above by $(\lambda + \alpha + 1)n^2$. Thus, if for notational convenience we denote $T_{\epsilon n}^0 = 0$ and consider the sum

$$S_j = \sum_{i=1}^j T_{\epsilon n}^i - T_{\epsilon n}^{i-1},$$

S_j dominates the random variable $X \sim \text{Gamma}(j, (\lambda + \alpha + 1)n^2)$. We can apply Chebyshev's inequality to see

$$\begin{aligned}
\mathbb{P}\left(S_j < \frac{j}{2(\lambda + \alpha + 1)n^2}\right) &\leq \mathbb{P}\left(X < \frac{j}{2(\lambda + \alpha + 1)n^2}\right) \\
&\leq \mathbb{P}\left(|X - \frac{j}{(\lambda + \alpha + 1)n^2}| > \frac{j}{2(\lambda + \alpha + 1)n^2}\right) \\
&\leq \frac{j \cdot 4(\lambda + \alpha + 1)^2 n^4}{j^2(\lambda + \alpha + 1)^2 n^4} \\
&= \frac{4}{j}.
\end{aligned} \tag{16}$$

On the event $\{T_{en}^j < \infty\} \cap \{S_j \geq \frac{j}{2(\lambda + \alpha + 1)n^2}\}$ we have that $\tau \geq \frac{j}{2(\lambda + \alpha + 1)n^2}$. Setting $j = \exp\left(\frac{\epsilon \eta^2 n}{172}\right)$ and taking a union bound on the probability of failure for large n we observe

$$\begin{aligned}
\mathbb{P}\left(\tau < \exp\left(\frac{\epsilon \eta^2 n}{344}\right)\right) &\leq \mathbb{P}\left(\tau < \frac{\exp\left(\frac{\epsilon \eta^2 n}{172}\right)}{2(\lambda + \alpha + 1)n^2}\right) \\
&\leq \mathbb{P}(T_{en}^j = \infty) + \mathbb{P}\left(S_j < \frac{\exp\left(\frac{\epsilon \eta^2 n}{172}\right)}{2(\lambda + \alpha + 1)n^2}\right) \\
&\leq \mathbb{P}(T_{en}^j = \infty) + \mathbb{P}\left(X < \frac{\exp\left(\frac{\epsilon \eta^2 n}{172}\right)}{2(\lambda + \alpha + 1)n^2}\right) \\
&\leq \exp\left(\frac{-\epsilon \eta^2 n}{172}\right) + 4 \exp\left(\frac{-\epsilon \eta^2 n}{172}\right) \\
&\rightarrow 0 \text{ as } n \rightarrow \infty,
\end{aligned} \tag{17}$$

and this completes the proof. $\square$

References

- [1] N. Berger, C. Borgs, J. T. Chayes, and A. Saberi. On the spread of viruses on the internet. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, page 301–310, USA, 2005. Society for Industrial and Applied Mathematics.
- [2] S. Bhamidi, D. Nam, O. Nguyen, and A. Sly. Survival and extinction of epidemics on random graphs with general degree. *Annals of Probability*, 49(1):244–286, 2021.
- [3] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 02 2013.
- [4] E. A. Cator and H. Don. Explicit bounds for critical infection rates and expected extinction times of the contact process on finite random graphs. *Bernoulli*, 27(3):1556–1582, 2021.
- [5] S. Chatterjee and R. Durrett. Contact processes on random graphs with power law degree distributions have critical value 0. *The Annals of Probability*, 37(6), Nov. 2009.

- [6] S. Chatterjee, D. Sivakoff, and M. Wascher. The effect of avoiding known infected neighbors on the persistence of a recurring infection process. *Electronic Journal of Probability*, 27:1 – 40, 2022.
- [7] S. Chatterjee, D. Sivakoff, and M. Wascher. The effects of individual versus community-influenced isolation on sis epidemic persistence on finite random graphs, 2025.
- [8] F. Chung, L. Lu, and V. Vu. Eigenvalues of random power law graphs. *Annals of Combinatorics*, 7(1):21–33, June 2003.
- [9] J. Cox and R. Durrett. Limit theorems for the spread of epidemics and forest fires. *Stochastic Processes and their Applications*, 30(2):171–191, 1988.
- [10] H. David and H. Nagaraja. *Order Statistics*. John Wiley & Sons, 2003.
- [11] G. Demirel, E. Barter, and T. Gross. Dynamics of epidemic diseases on a growing adaptive network. *Scientific Reports*, 10(7):42352, 2017.
- [12] R. Durrett. *Dynamics on Graphs*. 2024.
- [13] R. Durrett and P. Jung. Two phase transitions for the contact process on small worlds. *Stochastic Processes and their Applications*, 117(12):1910–1927, 2007.
- [14] R. Durrett and X.-F. Liu. The Contact Process on a Finite Set. *The Annals of Probability*, 16(3):1158 – 1173, 1988.
- [15] R. Durrett and C. Neuhauser. Epidemics with Recovery in $D = 2$. *The Annals of Applied Probability*, 1(2):189 – 206, 1991.
- [16] R. Durrett and D. Yao. Susceptible–infected epidemics on evolving graphs. *Electronic Journal of Probability*, 27(none):1 – 66, 2022.
- [17] S. C. Ferreira, R. S. Sander, and R. Pastor-Satorras. Collective versus hub activation of epidemic phases on networks. *Phys. Rev. E*, 93:032314, Mar 2016.
- [18] A. Flaxman, A. Frieze, and T. Fenner. High degree vertices and eigenvalues in the preferential attachment graph. *Internet Mathematics*, 2(1):1–19, 2005.
- [19] T. Friedrich, A. Göbel, N. Klodt, M. Krejca, and M. Pappik. Analysis of the survival time of the sirs process via expansion. *Electronic Journal of Probability*, 29, 01 2024.
- [20] T. Gross and B. Blasius. Adaptive coevolutionary networks: a review. *Journal of the Royal Society, Interface*, 5(20):259–271, 2007.
- [21] T. Gross, C. J. D. D’Lima, and B. Blasius. Epidemic dynamics on an adaptive network. *Phys. Rev. Lett.*, 96:208701, May 2006.
- [22] D. Guo, S. Trajanovski, R. van de Bovenkamp, H. Wang, and P. Van Mieghem. Epidemic threshold and topological structure of susceptible-infectious-susceptible epidemics in adaptive networks. *Phys. Rev. E*, 88:042802, Oct 2013.

- [23] A. Göbel, N. Klodt, and M. S. Krejca. Reemergence of the epidemic threshold in sirs infections on connected stars, 2026.
- [24] T. E. Harris. Contact Interactions on a Lattice. *The Annals of Probability*, 2(6):969 – 988, 1974.
- [25] Z. He, S. Dhara, and D. Mukherjee. Waning immunity fails to restore a positive epidemic threshold on power-law networks, 2026.
- [26] E. Jacob and P. Mörters. The contact process on scale-free networks evolving by vertex updating. *Royal Society open science*, 4(5):170081, 2017.
- [27] Y. Jiang, R. Kassem, G. York, M. Junge, and R. Durrett. Sir epidemics on evolving graphs, 2019.
- [28] P. Lam and O. Nguyen. Supercriticality of the sirs on random networks, 2026.
- [29] P. Lam, O. Nguyen, and I. Yang. Optimal bound for survival time of the sirs process on star graphs, 2024.
- [30] T. M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, volume 324. Springer Berlin, Heidelberg, 2013.
- [31] D. Mollison. Modelling biological invasions: chance, explanation, prediction. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 314(1167):675–693, 12 1986.
- [32] D. Mollison and K. Kuulasmaa. *Spatial epidemic models: theory and simulations*, pages 291–309. 1985.
- [33] T. Mountford, J.-C. Mourrat, D. Valesin, and Q. Yao. Exponential extinction time of the contact process on finite graphs. *Stochastic Processes and their Applications*, 126(7):1974–2013, 2016.
- [34] T. Mountford, D. Valesin, and Q. Yao. Metastable densities for the contact process on power law random graphs. *Electronic Journal of Probability*, 18(none), Jan. 2013.
- [35] T. S. Mountford. A metastable result for the finite multidimensional contact process. *Canadian Mathematical Bulletin*, 36(2):216–226, 1993.
- [36] J.-C. Mourrat and D. Valesin. Phase transition of the contact process on random regular graphs. *Electronic Journal of Probability*, 21:1–17, 2016.
- [37] D. Nam, O. Nguyen, and A. Sly. Critical value asymptotics for the contact process on random graphs. *Transactions of the American Mathematical Society*, 375(6):3899–3967, June 2022. Publisher Copyright: © 2022 American Mathematical Society.
- [38] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani. Epidemic processes in complex networks. *Rev. Mod. Phys.*, 87:925–979, Aug 2015.
- [39] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86(14):3200–3203, Apr. 2001.

- [40] R. Pemantle. The contact process on trees. *The Annals of Probability*, pages 2089–2116, 1992.
- [41] N. Pipatsart, W. Triampo, and C. Modchang. Stochastic models of emerging infectious disease transmission on adaptive random networks. *Comput Math Methods Med*, 2017(2403851), 2017.
- [42] D. Remenik. The contact process in a dynamic random environment. *The Annals of Applied Probability*, 18(6):2392–2420, 2008.
- [43] J. C. M. Silva, D. H. Silva, F. A. Rodrigues, and S. C. Ferreira. Comparison of theoretical approaches for epidemic processes with waning immunity in complex networks. *Phys. Rev. E*, 106:034317, Sep 2022.
- [44] D. Sivakoff. Contact process on a graph with communities. *ALEA, Latin American Journal of Probability and Mathematical Statistics*, 14:9–31, 2017.
- [45] I. Tunc, M. S. Shkarayev, and L. B. Shaw. Epidemics in adaptive social networks with temporary link deactivation. *Journal of statistical physics*, 151(1–2), 2013.
- [46] D. Valesin. The contact process on random graphs. *Ensaaios Matemáticos*, 40(1):1–115, December 2024.
- [47] J. van den Berg, G. R. Grimmett, and R. B. Schinazi. Dependent random graphs and spatial epidemics. *The Annals of Applied Probability*, 8(2):317–336, 1998.
- [48] R. van der Hofstad. *Random graphs and complex networks*, volume 2. Cambridge university press, 2024.
- [49] P. Van Mieghem, J. Omic, and R. Kooij. Virus spread in networks. *IEEE/ACM Transactions on Networking*, 17(1):1–14, 2009.
- [50] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos. Epidemic spreading in real networks: an eigenvalue viewpoint. In *22nd International Symposium on Reliable Distributed Systems, 2003. Proceedings.*, pages 25–34, 2003.