

Finding Super-spreaders in Network Cascades

Elchanan Mossel

Department of Mathematics, MIT

ELMOS@MIT.EDU

Anirudh Sridhar

Department of Mathematics, MIT

ANISRI@MIT.EDU

Editors: Shipra Agrawal and Aaron Roth

Abstract

Suppose that a cascade (e.g., an epidemic) spreads on an unknown graph, and only the infection times of vertices are observed. What can be learned about the graph from the infection times caused by multiple distinct cascades? Most of the literature on this topic focuses on the task of recovering the *entire* graph, which requires $\Omega(\log n)$ cascades for an n -vertex bounded degree graph. Here we ask a different question: can the important parts of the graph be estimated from just a few (i.e., constant number) of cascades, even as n grows large?

In this work, we focus on identifying super-spreaders (i.e., high-degree vertices) from infection times caused by a Susceptible-Infected process on a graph. Our first main result shows that vertices of degree greater than $n^{3/4}$ can indeed be estimated from a constant number of cascades. Our algorithm for doing so leverages a novel connection between vertex degrees and the second derivative of the cumulative infection curve. Conversely, we show that estimating vertices of degree smaller than $n^{1/2}$ requires at least $\log(n)/\log \log(n)$ cascades. Surprisingly, this matches (up to $\log \log n$ factors) the number of cascades needed to learn the *entire* graph if it is a tree.

Keywords: Network cascades, Susceptible-Infected process, structure learning, sample complexity

1. Introduction

Cascading behaviors in networks are ubiquitous. In human interaction networks, the spread of diseases through local interactions can quickly escalate, causing population-level pandemics (Brauer and Castillo-Chavez, 2012). Similarly, the structure of social networks plays a fundamental role in the mass dissemination of ideas, such as product adoption or misinformation (Leskovec et al., 2007; Strang and Soule, 1998; Rogers, 2003; Friggeri et al., 2014). In all of these examples, understanding the key factors which facilitate cascades is crucial for the analysis and control of these processes.

The structure of the underlying network plays a fundamental role in the evolution of a cascading process. Consider, for instance, the impact of a high-degree vertex. Once such a vertex is affected by the cascade, they may, in turn, disseminate the cascade’s effects on a massive scale to their neighbors. Due to their potential to accelerate the course of a cascade in this manner, high-degree vertices are of special interest in the analysis of spreading processes. In the context of epidemics, quarantining or immunizing high-degree individuals, also known as *super-spreaders*, can greatly mitigate the damage caused by a pandemic (Pastor-Satorras and Vespignani, 2002; Dezső and Barabási, 2002). In social networks, firms may want to leverage the impact of high-degree vertices, also known as *influencers*, to accelerate the speed of product adoption (Hinz et al., 2011).

A fundamental hurdle in carrying out these ideas in practice is that the underlying network may be partially or fully unknown. In epidemiology, disease-spreading contacts between individuals are often not measured. In marketing, even if retailers have access to certain social networks, only

some of the observed links may be effective in facilitating a cascade, and many links may also be unobserved (e.g., if two individuals communicate through other means). While it is possible to learn aspects of the network through contact tracing or interviewing, such approaches are quite intensive and time-consuming (Spencer et al., 2020). To mitigate this issue, we aim to understand the extent to which more efficient *data-driven* approaches can be used to learn about the network. Specifically, we study how network structure can be estimated from the set of “infection times” caused by a cascade, which we call the *cascade trace*. Typically, a single cascade trace may not be enough to learn meaningful information, but much more can be said by combining information from multiple different cascades on the same network.

The task of estimating a network from multiple cascade traces, known as structure learning, has received considerable attention in the past decade (Gomez-Rodriguez et al., 2012; Du et al., 2012; Netrapalli and Sanghavi, 2012; Abrahao et al., 2013; Daneshmand et al., 2014; Khim and Loh, 2018; Hoffmann and Caramanis, 2019; He et al., 2020; Hoffmann et al., 2020; Wilinski and Lokhov, 2021). Almost all the literature on this topic focuses on exactly recovering the *entire* network under sparsity constraints (i.e., bounds on the maximum degree). However, the number of cascade traces required to do so grows as a function of the network size (Netrapalli and Sanghavi, 2012; Abrahao et al., 2013; Daneshmand et al., 2014). This can be prohibitive in applications such as epidemiology where it is costly to allow many cascades to spread. Motivated by this limitation, we ask a different question:

Can the important parts of the network, such as the location of high-degree vertices, be learned from just a few cascade traces?

We find that the answer is quite subtle. In an n -vertex graph, we show that vertices of degree at least $n^{3/4}$ can indeed be correctly estimated from a constant number of cascade traces. Conversely, estimating vertices of degree smaller than \sqrt{n} can be almost as hard as learning the *entire* network, in terms of the number of cascade traces required.

1.1. Summary of contributions

In a bit more detail, we assume that each cascade spreads on an n -vertex graph G with the following structure: a constant number of (high-degree) vertices have degree at least n^α for some $\alpha \in (0, 1)$, and all other vertices have degree at most $n^{o(1)}$. The goal is to understand how many cascade traces are needed to estimate the set of high-degree vertices with probability $1 - o(1)$.

Our first main contribution is an estimator for the set of high-degree vertices which requires only a constant number of cascade traces. Our estimator is quite simple, and relies only on properties of the *infection curve* $I(t)$, which counts the number of infections that occur before time t . The key idea behind our method is a novel connection between vertex degrees and the second derivative of the infection curve. Specifically, we show that if v becomes infected at time $T(v)$, then a discretized version of the second derivative of $I(t)$ evaluated at time $T(v)$ is a (nearly) unbiased estimator for the degree of v . Motivated by this insight, our estimator computes the (appropriately discretized) second derivative of $I(t)$ at $T(v)$ for each v , and compares it to a threshold. The set of estimated high-degree vertices is then given by all vertices which pass the threshold in all of the observed cascade traces (see Algorithm 1 for more details). Our analysis shows that this estimator succeeds provided $\alpha > 3/4$. Notably, our algorithm is a significant departure from existing methods in the literature on learning in graphical models. Indeed, prior algorithms typically focus on learning the




Figure 1: Phase diagram for the possibility and impossibility of estimating high-degree vertices in a graph G . *Blue region*: High-degree vertices can be estimated from more than $1/(\alpha - 3/4)$ traces. *Red region*: Estimating high-degree vertices is impossible, even when G is known to be a tree. *Green region*: Full recovery of G is possible if G is a tree, hence estimation of high-degree vertices is also possible. *White region*: Regimes with small gap between the bounds provided by our analysis. *Yellow region*: The main open problem left given our work is to determine the sample complexity in this region.

precise edges and neighborhoods in the graph (Netrapalli and Sanghavi, 2012; Abrahao et al., 2013; Daneshmand et al., 2014; Gomez-Rodriguez et al., 2012; Du et al., 2012; He et al., 2020; Wilinski and Likhov, 2021). In contrast, we demonstrate how degree information can be *directly* estimated from just a few traces without needing to infer neighborhoods.

A natural follow-up question is whether one can identify high-degree vertices of degree much smaller than $n^{3/4}$, while still using only a constant number of traces. We provide a negative answer, showing that when $\alpha \in (0, 1/2)$, there exist “hard” instances for which $\log(n)/\log \log(n)$ traces are needed to estimate high-degree vertices, no matter the estimator. In other words, it is *information-theoretically impossible* to correctly estimate high-degree vertices in this case. In our proof, we show that estimating high-degree vertices in a particular ensemble of graphs is equivalent to a variant of the *sparse mixture detection problem* (Dobrusin, 1958; Burnashev and Begmatov, 1991; Ingster, 1996; Donoho and Jin, 2004; Cai and Wu, 2014). Following the approach of Cai and Wu (2014), we solve our version of the sparse mixture detection problem, which leads to a lower bound of $\log(n)/\log \log(n)$ traces for the estimation of high-degree vertices. A surprising implication of our impossibility result is that for $\alpha \in (0, 1/2)$, estimating high-degree vertices can be nearly as hard as learning the entire network. Indeed, using $\log(n)$ traces, the Tree Reconstruction algorithm of Abrahao et al. (2013) is able to exactly recover G , provided it is a tree. On the other hand, the graphs in our ensemble of hard instances are all trees, hence $\log(n)/\log \log(n)$ traces are required to estimate high-degree vertices even if G is known to be a tree. In summary, if G is known to be a tree, the sample complexity (with respect to the number of cascade traces) of estimating high-degree vertices matches that of learning G *exactly*, up to at most a $\log \log(n)$ factor.

Together, our two results highlight a curious phase transition in α (see Figure 1). When $\alpha \in (3/4, 1)$, high-degree vertices can be identified using essentially minimal data about the cascade.

When $\alpha \in (0, 1/2)$, we have the other extreme: learning high-degree vertices is nearly as challenging as learning the entire network, when G is a tree. Our work leaves open the intriguing regime of $\alpha \in [1/2, 3/4]$.

1.2. Related Work

The problem of inferring a network from cascade traces was first studied in the context of epidemiology (Wallinga and Teunis, 2004) and information flow in blogs (Adar and Adamic, 2005). Since then, there has been a large body of theoretical and empirical work on the subject. On the empirical side, Gomez-Rodriguez et al. (2012) initiated a study of scalable and mathematically principled methods with their NetInf algorithm. Since then, a number of follow-up works have tackled the exact network inference problem using a combination of likelihood-based approaches, convex optimization, message passing and machine learning (Daneshmand et al., 2014; Gomez-Rodriguez et al., 2012; Du et al., 2012; He et al., 2020; Wilinski and Lokhov, 2021). In a related vein, there has been significant interest from the machine learning community in the task of learning Hawkes processes from vertex activation times; see Lima (2023) and references therein. We expect that our methods may be useful in this setting as well.

The first results on the sample complexity (i.e., number of traces needed) of exactly recovering the underlying graph were by Netrapalli and Sanghavi (2012), which applied to discrete-time cascades (i.e., the independent cascade model) spreading on bounded-degree graphs. This work was soon followed by Abrahao et al. (2013), who studied the sample complexity of exact network inference using a variant of the SI process to model the cascade. Since their work, a number of extensions have been explored, including inference from noisy infection times (Hoffmann and Caramanis, 2019), correlated cascades (Khim and Loh, 2018) and learning mixtures of graphs (Hoffmann et al., 2020). A common technique underlying these is an analysis of maximum likelihood estimator for the underlying network. We note that our algorithm is a significant departure from these methods, as it is able to directly estimate degree information *without* knowledge of the precise topology.

Our work falls under the broader umbrella of literature studying the sample complexity of learning structures in graphical models. Our work deals with observations from *transient* models, though learning from *stationary* models such as Gaussian Graphical Models and Ising models have also received considerable attention (Ravikumar et al., 2011; Meinshausen and Bühlmann, 2006; Misra et al., 2020; Wang et al., 2010; Cai et al., 2016; Kelner et al., 2020; Klivans and Meka, 2017; Wu et al., 2019; Vuffray et al., 2020; Santhanam and Wainwright, 2012; Bresler, 2015; Bresler et al., 2013). While such frameworks are naturally quite different from ours, there are also many qualitative similarities. For the task of recovering the precision matrix of a n -dimensional Gaussian graphical model, it is known that $\Theta(\log n)$ independent samples are necessary and sufficient under some sparsity conditions (Misra et al., 2020; Wang et al., 2010; Cai et al., 2016; Kelner et al., 2020). Similarly, $\Theta(\log n)$ independent samples are necessary and sufficient for learning bounded-width Ising models (Bresler et al., 2008; Wu et al., 2019; Vuffray et al., 2020; Santhanam and Wainwright, 2012; Bresler, 2015) as well as general classes of Markov Random Fields (Bresler et al., 2013). These results parallel the sample complexity of exactly recovering the graph from cascade traces, and many algorithmic approaches are similar in flavor.

While the work above largely focuses on recovering the underlying network, we mention a few recent papers which also aim to learn the “important” parts of the relevant graphical model, though with significantly different interpretations of importance. Boix-Adserà et al. (2022) as well

as [Bresler and Karzand \(2020\)](#) focus on the task of learning tree-structured Ising models with a good prediction accuracy with respect to the ground-truth model. [Eckles et al. \(2022\)](#) leverage cascade traces from a discrete-time cascade model to perform influence maximization. These works, in conjunction with ours, demonstrate the potential to learn key information about distributions and systems *without* fully learning the underlying model.

1.3. Organization

The rest of the paper is organized as follows. In Section 2, we rigorously formulate the problem of estimating high-degree vertices and state our main results. In Section 3, we present our method for estimating high-degree vertices (Algorithm 1) and provide an overview of its proof of correctness. Section 4 contains a proof sketch of our impossibility result. We conclude and discuss future directions in Section 5. All proof details can be found in the appendices.

2. Definitions and Main Results

Notation. For a set S , we let $|S|$ denote its cardinality. For a graph $G = (V, E)$, we denote the number of vertices by $n := |V|$ and let $\mathcal{N}(v)$ denote the neighborhood of a vertex v . For a subset of vertices $S \subset V$, we denote $\text{cut}(S) := |\{(u, v) \in E : u \in S, v \notin S\}|$. Throughout the paper, we use standard asymptotic notation and all limits are as $n \rightarrow \infty$ unless otherwise specified.

The Susceptible-Infected process. We model the cascade using the well-known Susceptible-Infected (SI) process, which is a continuous-time Markov process in which susceptible vertices in a graph become infected upon interacting with infected neighbors. In more detail, we let G be the graph upon which the cascade spreads. We denote by $\mathcal{I}(t)$ the set of vertices which are infected at a time index $t \geq 0$, and we let $\lambda > 0$ represent the *interaction rate* between individuals. Initially, we assume that only a single vertex v_0 is infected; that is, $\mathcal{I}(0) = \{v_0\}$. Given $\mathcal{I}(t)$, the cascade evolves as follows: for $v \in V \setminus \mathcal{I}(t)$, it holds that as $\epsilon \rightarrow 0$,

$$\mathbb{P}(v \in \mathcal{I}(t + \epsilon) | \mathcal{F}_t) = \epsilon \lambda |\mathcal{N}(v) \cap \mathcal{I}(t)| + o(\epsilon), \quad (1)$$

where $\{\mathcal{F}_t\}_{t \geq 0}$ is the natural filtration corresponding to the stochastic evolution of the cascade and $o(\epsilon) \rightarrow 0$ faster than $\epsilon \rightarrow 0$. In words, the rate at which a susceptible vertex becomes infected is proportional to the number of its infected neighbors and the interaction rate λ . Given the evolution of the cascade $\mathcal{I}(t)$, a few important quantities include the *infection curve* $I(t) := |\mathcal{I}(t)|$, the set $\mathcal{I}[a, b]$ for $0 \leq a \leq b$, which denotes the set of infections occurring in the interval $[a, b]$, and $I[a, b] := |\mathcal{I}[a, b]|$.

We remark that the SI process (1) is closely related to a number of other well-studied models. It is known to be equivalent to First Passage Percolation with $\text{Exp}(\lambda)$ edge weights, a popular model in mathematical physics ([Auffinger et al., 2017](#); [Richardson, 1973](#)). It is also equivalent to the Contact Process without recovery ([Liggett, 1999](#); [Mollison, 1978](#)), and has been used as a rigorous foundation for standard population-based epidemiological models ([Brauer and Castillo-Chavez, 2012](#), Chapter 9). In recent years, a large body of work has also investigated generalizations of (1) with edge-dependent interaction rates (e.g., due to mask-wearing tendencies, multiple viral strains); see, e.g., ([Allard et al., 2009](#); [Alexander and Day, 2010](#); [Eletreby et al., 2020](#)).

For a given vertex $v \in V$, we formally define its *infection time* in the SI process to be

$$T(v) := \inf\{t \geq 0 : v \in \mathcal{I}(t)\}.$$

We call the collection of infection times and associated vertices $\mathbf{T} := \{(v, T(v))\}_{v \in V}$ the *cascade trace*.

Learning from cascade traces. Suppose that the underlying graph G is unknown, but we observe multiple independent cascade traces on G , given formally by

$$(\mathbf{T}_1, \dots, \mathbf{T}_K) \sim \mathcal{SI}(G, v_{0,1}) \otimes \dots \otimes \mathcal{SI}(G, v_{0,K}), \quad (2)$$

where $v_{0,i}$ is the source vertex for the i th cascade. For brevity, for a graph G and collection of source vertices $v_0 := (v_{0,1}, \dots, v_{0,K})$, we denote \mathbb{P}_{G, v_0} to be the probability measure corresponding to (2). Broadly, we are interested in the following question of structural inference: *What can be learned about G from K cascade traces?*

While this has been a question of significant interest in the past decade, the existing literature on the topic almost exclusively focuses on *exactly* recovering G . In the context of cascade traces, this task was first studied by [Netrapalli and Sanghavi \(2012\)](#) for a discrete-time cascade model. In particular, they proved that for n -vertex graphs G with maximum degree at most Δ , $K = \text{poly}(\Delta) \log n$ is necessary and sufficient to exactly recover G . [Abraham et al. \(2013\)](#) studied the exact recovery problem for a slightly more general version of the SI process (1), showing that $K = \text{poly}(\Delta) \log(n)$ suffices for general graphs (provided Δ does not scale as a function of n), and that $K = \Theta(\log n)$ traces suffices if G is a tree.

A key takeaway from this literature on exact recovery of G is that the number of cascade traces K must grow as a function of the network size n . However, this requirement may be prohibitive in settings such as epidemiology where one must make decisions based on a relatively small amount of data. In light of this issue, we study whether it is possible to learn the *important* parts of a network – such as the location of high-degree vertices – from just a few cascade traces. To the best of our knowledge, this angle has not been previously investigated.

Estimating high-degree vertices. To formalize the task of estimating high-degree vertices, we assume that most of the vertices in G have a relatively small degree, except for a constant number of vertices with substantially larger degree. The latter type is the set of high-degree vertices we wish to estimate. We mathematically capture this structural property with the class of graphs $\mathcal{G}(n, m, d, D)$, described below.

Definition 1 We say that $G \in \mathcal{G}(n, m, d, D)$ if and only if (1) G is a connected graph on n vertices, (2) there are at most m vertices of degree at least D , and (3) all other vertices have degree at most d . We denote the set of high-degree vertices in $G \in \mathcal{G}(n, m, d, D)$ to be $\text{highdeg}(G) := \{v \in V : \deg(v) \geq D\}$.

We remark that the assumption of connectivity ensures that all vertices are eventually infected, which is in general necessary for the estimation of high-degree vertices. Indeed, if a high-degree vertex never becomes infected, it is impossible to estimate it. Next, we assume the following about the scaling of the parameters m, d, D with respect to n .

Assumption 2 We assume that m is constant, $d = n^{o(1)}$ and $D = n^\alpha$ for some $\alpha \in (0, 1)$.

We remark that this assumption is quite different from the usual assumption made in the literature on structure learning. Indeed, in much of the work on exact recovery of G , it is assumed that

its maximum degree is bounded as $n \rightarrow \infty$ (Netrapalli and Sanghavi, 2012; Abrahao et al., 2013; Daneshmand et al., 2014). A notable exception is the Tree Reconstruction algorithm of Abrahao et al. (2013), which has provable guarantees for *any* tree with a sufficiently large number of vertices. However, as we are only interested in estimating high-degree vertices rather than specific connections in the network, we may take significantly more relaxed assumptions compared to the bulk of the structure learning literature. In the remainder of the paper, when referencing any graph G we will assume that $G \in \mathcal{G}(n, m, d, D)$ and that Assumption 2 holds.

Our first main result shows that we can correctly identify all vertices in G of degree larger than $n^{3/4}$, using a *constant* number of cascade traces – far less than what is needed to learn edges in the graph.

Theorem 3 *Suppose that $\alpha \in (3/4, 1)$ and that $D = n^\alpha$. If K satisfies*

$$K > \frac{1}{\alpha - 3/4},$$

then there is an estimator $\widehat{\text{HD}}$ such that for any $G \in \mathcal{G}(n, m, d, D)$ and any collection of source vertices $v_0 = (v_{0,1}, \dots, v_{0,K})$,

$$\mathbb{P}_{G, v_0} \left(\widehat{\text{HD}}(\mathbf{T}_1, \dots, \mathbf{T}_K) = \text{highdeg}(G) \right) = 1 - o(1),$$

where $o(1) \rightarrow 0$ as $n \rightarrow \infty$.

Though Theorem 3 shows that a constant number of traces suffices for any $\alpha > 3/4$, K grows as α approaches $3/4$. While this condition is needed in our analysis to handle worst-case examples of G , we remark that in practice K could be smaller and may not explode as α approaches $3/4$. To prove the theorem, we develop a novel algorithm for estimating high-degree vertices from a constant number of cascade traces (Algorithm 1). A key supporting result underlying the theorem and Algorithm 1 is that the time at which a high-degree vertex is infected can be accurately estimated from just a *single* cascade. See Section 3 for more details.

Our second main contribution is a negative result, showing that *any* estimator requires at least $\log(n)/\log \log(n)$ traces to successfully identify high-degree vertices of degree smaller than \sqrt{n} .

Theorem 4 *Fix $\epsilon > 0$, let $\alpha \in (0, 1/2)$. Suppose that $D = n^\alpha$, $d \geq \log^2 n$ and $m = 1$. Additionally, assume that*

$$K \leq \left(\frac{1 - 2\alpha - \epsilon}{5} \right) \frac{\log n}{\log \log n}. \quad (3)$$

Then there exists an ensemble $\mathcal{H} \subset \mathcal{G}(n, m, d, D)$ and a probability distribution μ over graphs in \mathcal{H} and source vertices $v_0 = (v_{0,1}, \dots, v_{0,K})$ such that, for any estimator HD ,

$$\mathbb{P}_{G, v_0 \sim \mu} (\text{HD}(\mathbf{T}_1, \dots, \mathbf{T}_K) = \text{highdeg}(G)) = o(1).$$

An important implication of the theorem is that if K is smaller than $\log(n)/\log \log(n)$, no estimator can successfully identify high-degree vertices for every graph in the class $\mathcal{G}(n, m, d, D)$. The ensemble \mathcal{H} used in the theorem is a collection of trees with a particular structure; see Sections 4 and D for more details. This observation, in the context of prior work on the sample complexity of exact recovery of G (Abrahao et al., 2013), shows that learning high-degree vertices when $\alpha \in (0, 1/2)$ is *nearly* as hard as learning the full graph in the case of trees. To emphasize this point, we rephrase the result of Abrahao et al. (2013) below.

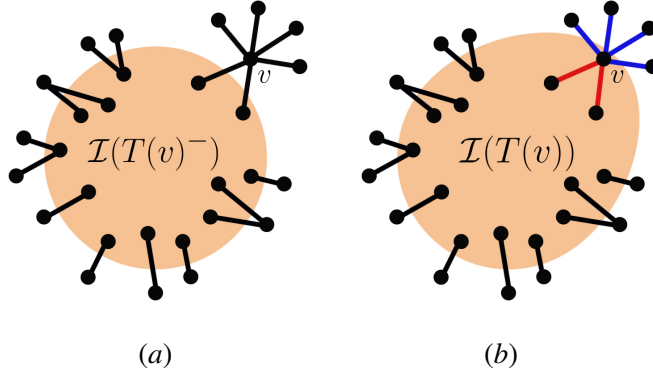


Figure 2: Visualization of the edges contributing to $\text{cut}(\mathcal{I}(t))$ before (a) and after (b) a vertex v becomes infected. In Figure 2(b), the number of blue and red edges denote the positive and negative change to $\text{cut}(\mathcal{I}(t))$, respectively, upon v being infected.

Theorem 5 *Suppose that $K \geq C \log n$, for some universal constant $C > 0$. Then there exists an estimator \widehat{G} such that for any collection of source vertices $v_0 = (v_{0,1}, \dots, v_{0,K})$ and any tree G , it holds that $\mathbb{P}_{G, v_0}(\widehat{G} = G) = 1 - o(1)$.*

Together, Theorems 4 and 5 show that there is at most a $\log \log n$ multiplicative gap between the sample complexity of estimating high-degree vertices for $\alpha \in (0, 1/2)$ and that of fully recovering G , in the case where G is a tree.

3. Estimating high-degree vertices when $\alpha > 3/4$

The starting point for the design of our estimator is the following simple observation: when a high-degree vertex becomes infected, many of its neighbors will become infected shortly after. In a bit more detail, we could assess whether a vertex v is high-degree by examining the number of infections occurring in a δ -size interval after v is infected, given formally by $I[T(v), T(v) + \delta]$. Since we can write $I[T(v), T(v) + \delta] = I(T(v) + \delta) - I(T(v))$, the quantity $I[T(v), T(v) + \delta]$ can be thought of as the discrete first derivative of the infection curve $I(t)$. This suggests that high-degree vertices could be estimated by examining the points at which the first derivative of $I(t)$ is large. Unfortunately, such an approach fails in general. Indeed, if we observe many infections in the interval $[T(v), T(v) + \delta]$, then it is challenging to discern whether the infections were mainly caused by a *single*, common infection (namely, v), or by many other infected vertices which just happened to be infected around the same time as v . However, it turns out that if one considers *higher-order* information – specifically, the *second* derivative of $I(t)$ – then high-degree vertices can be clearly identified.

To see this concretely, we begin by examining the first derivative of $I(t)$ in more detail. By (1),

$$J(t) := \mathbb{E} \left[\left. \frac{d}{dt} I(t) \right| \mathcal{F}_t \right] = \lambda \sum_{v \in V \setminus \mathcal{I}(t)} |\mathcal{N}(v) \cap \mathcal{I}(t)| = \lambda \text{cut}(\mathcal{I}(t)). \quad (4)$$




Figure 3: Plots of the discrete first (a) and second derivative (b) of the infection curve with $\delta = 0.075$ generated from a graph G with approximately 500,000 vertices and one high-degree vertex. The infection time of the high-degree vertex can be identified by the red dotted line in (a) and the large peak in the second derivative plot (b), highlighted by the red rectangle. G was chosen to be a balanced, 5-regular tree of height 8, where one vertex in the 6th layer of the tree has degree 7500.

Now suppose that a vertex v joins the cascade at time t (that is, $T(v) = t$). Since no other vertices can be infected at the exact same time due to the continuous-time nature of the cascade, we have

$$J(t) - J(t^-) = \lambda (|\mathcal{N}(v) \setminus \mathcal{I}(t)| - |\mathcal{N}(v) \cap \mathcal{I}(t)|) = \lambda \deg(v) - 2\lambda |\mathcal{N}(v) \cap \mathcal{I}(t)|. \quad (5)$$

The first equality in (5) can be explained as follows. When v joins the cascade, all of its susceptible neighbors, i.e., those in $\mathcal{N}(v) \setminus \mathcal{I}(t)$, become exposed. On the other hand, the edges which connect v to $\mathcal{I}(t)$ – namely, edges which connect v to $\mathcal{N}(v) \cap \mathcal{I}(t)$ – can no longer lead to infection events, hence the contribution for these edges is subtracted. See Figure 2 for a visualization of this idea.

Crucially, (5) shows that as long as a vertex is infected before most of its neighborhood (i.e., $|\mathcal{N}(v) \cap \mathcal{I}(t)|$ is much smaller than $\deg(v)$), the second-order changes in the infection curve, captured by $J(t) - J(t^-)$, can reveal the degree of a vertex. However, since $J(t)$ is not a directly observable quantity, we must approximate it from the cascade traces. A natural way to do so is by choosing $\delta > 0$ and using the discrete derivative $I[t, t + \delta]/\delta$ as a proxy for $J(t)$. Similarly, we can use $I[t - \delta, t]/\delta$ as a proxy for $J(t^-)$. We can therefore approximate the degree of v using

$$\widehat{\deg}_\delta(v) := \frac{I[T(v), T(v) + \delta] - I[T(v) - \delta, T(v)]}{\delta}.$$

In words, we estimate the degree of v by computing the discrete second derivative of the infection curve when v gets infected. As shown through the empirical example in Figure 3, this statistic is quite effective in identifying the infection time of a high-degree vertex from just a single cascade (this is also shown formally in Section B).

In light of (5), our estimator for $\text{highdeg}(G)$ returns the set of all vertices v for which $\widehat{\deg}_\delta(v)$ is large in all cascade traces; see Algorithm 1. As Theorem 6 shows, this method succeeds with high probability provided $\alpha > 3/4$.

Algorithm 1 Finding high-degree vertices via second derivative thresholding

Input: Cascade traces $\mathbf{T}_1, \dots, \mathbf{T}_K$ and parameters $\delta, \tau > 0$

Output: A set $\widehat{\text{HD}} \subset V$

1: For each $i \in [K]$ and $v \in V$ compute

$$\widehat{\text{deg}}_{\delta,i}(v) := \frac{I_i[T_i(v), T_i(v) + \delta] - I_i[T_i(v) - \delta, T_i(v)]}{\delta}.$$

2: Return $\widehat{\text{HD}} := \left\{ v \in V : \widehat{\text{deg}}_{\delta,i}(v) \geq \tau, \text{ for all } i \in [K] \right\}.$

Theorem 6 Suppose that $\alpha > 3/4$ and that

$$\delta = \frac{1}{n^{\alpha-1/4}}, \quad \tau = \frac{n^\alpha}{\log n}, \quad K > \frac{1}{\alpha - 3/4}.$$

Additionally assume that $d = n^{o(1)}$ and $D = n^\alpha$. Then for any $G \in \mathcal{G}(n, m, d, D)$, the output of Algorithm 1 is equal to $\text{highdeg}(G)$, with probability $1 - o(1)$.

Remark 7 A useful feature of Algorithm 1 is that it is adaptive to the unknown parameters m (the number of high-degree vertices) and λ (the spreading rate of the cascade). In contrast, the prior work on structure learning leverages likelihood-based methods which require precise knowledge of the cascade model (Netrapalli and Sanghavi, 2012; Abrahao et al., 2013).

Remark 8 Though our main result concerns the number of cascades K needed to learn high degree vertices, our techniques show that the infection times of high-degree vertices can be estimated from a single cascade; see Lemma 20 in Section B.

We note that Theorem 3 follows immediately from Theorem 6 as a corollary. We provide a sketch of the proof of Theorem 6 below, and defer the full details to Sections B and C.

Proof [Proof sketch of Theorem 6] Let $\gamma < \alpha$ be sufficiently close to α . The strategy of our proof is to show, for all $v \in V$, that

$$\left| \widehat{\text{deg}}_{\delta,i}(v) - \lambda \text{deg}(v) \right| \leq n^\gamma, \tag{6}$$

in all cascades. If (6) is true for a particular $i \in [K]$ and v is a low-degree vertex, then $\widehat{\text{deg}}_{\delta,i}(v) \leq \lambda \text{deg}(v) + n^\gamma \leq 2n^\gamma$. Else if v is high-degree, then $\widehat{\text{deg}}_{\delta,i}(v) \geq \lambda \text{deg}(v) - n^\gamma \geq \frac{\lambda}{2} n^\alpha$. If λ is not known a priori, any choice of threshold satisfying $\tau = n^{\alpha-o(1)}$ works in distinguishing the two cases. In particular, $\tau = n^\alpha / \log(n)$ suffices.

Due to the way $\widehat{\text{deg}}_{\delta,i}(v)$ is computed, establishing (6) reduces to showing that $I[t, t + \delta]/\delta$ concentrates around $J(t)$ and $I[t, t - \delta]/\delta$ concentrates around $J(t^-)$ for the values of t corresponding to vertices' infection times. For this task, choosing the right value of δ is crucial. If δ is too small, then $I[t, t + \delta]$ may also be quite small, making meaningful concentration around its mean impossible. On the other hand, if δ is too large, the mean of $I[t, t + \delta]/\delta$ may deviate significantly from $J(t)$ due to the potentially large changes to the set of infections in the interval $(t, t + \delta]$. An additional complication arising if δ is too large is that there may be a large number of low-degree vertices consistently infected around the same time as a high-degree vertex. If u is one such low-degree vertex, then $\widehat{\text{deg}}_{\delta}(u)$ can be quite large across multiple traces, leading to false positives.

Examining these conditions in more detail exposes why we must assume $\alpha > 3/4$ for Algorithm 1 to succeed. Let us focus on the behavior of $I[t, t + \delta]/\delta$ (the analysis of $I[t - \delta, t]/\delta$ is similar). We show that

$$\frac{I[t, t + \delta]}{\delta} = J(t) \pm O(n\delta) \pm O\left(\sqrt{\frac{n}{\delta}}\right). \quad (7)$$

Above, the $O(n\delta)$ fluctuation represents the deviation between $J(t)$ and the (conditional) mean of $I[t, t + \delta]/\delta$. This is essentially caused by the additional $O(n\delta)$ additional infections which occur in the interval $[t, t + \delta]$. The $O(\sqrt{n/\delta})$ fluctuation comes from the deviation of $I[t, t + \delta]$ from its mean, which can be derived through subgaussian-type concentration inequalities. To correctly estimate high-degree vertices through (5), both fluctuations need to be smaller than $D = n^\alpha$. We therefore require that $\max\{n\delta, \sqrt{n/\delta}\} = o(n^\alpha)$, which is equivalent to

$$\delta = o\left(n^{-(1-\alpha)}\right) \text{ and } \delta = \omega\left(n^{-(2\alpha-1)}\right). \quad (8)$$

We remark that (8) already requires $\alpha > 2/3$. Next, we need to ensure that there is no low-degree vertex u that is infected around the same time as a high-degree vertex in all cascade traces. Otherwise, as we mentioned earlier, $\widehat{\deg}_{\delta,i}(u)$ could be quite large, which would imply that the output of Algorithm 1 contains false positives. With this scenario in mind, suppose that a low-degree u becomes infected first at time t , and at that point in time, a high-degree v has M infected neighbors. By (1), it holds for small δ that

$$\mathbb{P}(|T(u) - T(v)| \leq \delta | \mathcal{F}_t) = \mathbb{P}(v \in \mathcal{I}(t + \delta) | \mathcal{F}_t) = \delta M + o(\delta).$$

The probability that u and v are close in all cascade traces is therefore at most $O((\delta M)^K)$. A union bound over pairs of low-degree u and high-degree v shows that

$$\mathbb{P}(\exists u \in V, v \in \text{highdeg}(G) : |T_i(u) - T_i(v)| \leq \delta \text{ in all cascades}) \leq O(nm(\delta M)^K). \quad (9)$$

To show that the right hand side of (9) is $o(1)$, we first need to choose δ such that $\delta M = o(1)$. Although M can in principle be as large as $\deg(v)$ (i.e., if all of v 's neighbors are infected before v), we show that with high probability $M \leq \sqrt{n}$ (see Lemma 15 in Section A). Hence we may choose δ satisfying

$$\delta = o\left(n^{-1/2}\right). \quad (10)$$

Writing $\delta = n^{-\beta}$ for $\beta > 0$, (8) and (10) imply that $\beta \in (1/2, 2\alpha - 1)$. This interval is nonempty when $\alpha > 3/4$, and we may choose $\beta := \alpha - 1/4$ to satisfy (8) and (10). Finally, since $\delta = n^{-(\alpha-1/4)}$ and $M \leq \sqrt{n}$, $K > (\alpha - 3/4)^{-1}$ is needed for the right hand side of (9) to be $o(1)$. ■

4. Impossibility of estimation when $\alpha \in (0, 1/2)$

Connections to sparse mixture detection. The cornerstone of our proof is a correspondence between high-degree estimation and the problem of sparse mixture detection. In a bit more detail, let G' be a graph on $n/2$ vertices of maximum degree $n^{o(1)}$ (i.e., all its vertices are low-degree). The graph G will be formed by adding an additional $n/2$ vertices labelled $1, \dots, n/2$ to G' , along with a single edge per added vertex which connects to a vertex in G' . A natural way of generating a

graph G without any high-degree vertices is to let each of the new $n/2$ vertices connect to a uniform random parent in G' . Through standard probabilistic arguments, it can be seen that no more than $O(\log n)$ neighbors are added to any particular vertex in G' , so G contains no high-degree vertices.

A simple modification of this procedure can also generate graphs with a high-degree vertex. Suppose we fix a vertex v in G' . For each of the additional $n/2$ vertices, with probability $2D/n$, it connects to v . Otherwise, with probability $1 - 2D/n$, it connects to a uniform random vertex in G' . In expectation, v receives D additional connections (so, in particular, $\deg(v) \geq D$ in expectation). As before, all of the other vertices receive at most $O(\log n)$ additional neighbors, with high probability. Hence we obtain a graph G in which v is the only high-degree vertex.

Observe that in the second procedure, the distribution of the infection times of each of the added $n/2$ vertices can be viewed as a mixture of two distributions. To be precise, let P be the distribution of a vertex's infection times in the observed cascade traces if they connect to a uniform random neighbor in G' . Similarly, let Q be the distribution of a vertex's infection times if they connect to v . Then the task of differentiating between the two generative models for G – or equivalently, of testing whether v is high-degree – can be phrased as the following hypothesis testing problem:

$$H_0 : \mathbf{T}(1), \dots, \mathbf{T}(n/2) \sim P \quad H_1 : \mathbf{T}(1), \dots, \mathbf{T}(n/2) \sim ((1 - \epsilon_n) P + \epsilon_n Q), \quad (11)$$

where we set $\epsilon_n := 2D/n$ above, and we denote $\mathbf{T}(j) := (T_1(j), \dots, T_K(j))$ where $T_i(j)$ is the infection time of vertex j in the i th cascade.

It turns out that (11) is a special case of a well-studied problem known as *sparse mixture detection*, which dates back to the work of Dobrusin (1958). Burnashev and Begmatov (1991) studied a variant related to Gaussian processes, and Ingster (1996) identified a subtle phase transition between the possibility and impossibility of detection for mixtures of Gaussian distributions. The problem was popularized by the work of Donoho and Jin (2004), who formulated an adaptive method in place of the likelihood ratio called *higher criticism*, again for mixtures of Gaussians. Later, Cai and Wu (2014) developed an analysis of the sparse mixture problem for generic distributions. We remark that Mossel and Roch (2017) used sparse mixture detection in the inference of combinatorial structure, though in a very different context.

Impossibility of detecting high-degree vertices. We solve our instance of the sparse mixture detection problem (11) using the techniques developed by Cai and Wu (2014). That is, we bound the total variation distance between the two distributions in (11) by the χ^2 divergence since it tensorizes and because it is easier to study for problems involving mixtures. Our analysis shows that if K satisfies (3), then the χ^2 divergence tends to zero. As a result, it is impossible to distinguish between scenarios where there are no high-degree vertices (H_0) and where v is a high-degree vertex (H_1).

Putting everything together. Finally, we show that the impossibility of detection implies that no estimator can successfully output the correct high-degree vertex with probability greater than $o(1)$. We remark that the ensemble \mathcal{H} and distribution μ stated in the theorem essentially corresponds to the random constructions for G described in this subsection. For more details, see Section D.

5. Discussion and future work

In this work, we study the sample complexity of estimating super-spreaders (i.e., high-degree vertices) in network cascades, using only the infection times of vertices across multiple cascades. This is the first work to consider this task, to the best of our knowledge. Our results show that estimating vertices with degree at least $n^{3/4}$ can be done using a constant number of cascades, while estimating vertices with degree at most $n^{1/2}$ requires almost as many cascades as learning the entire graph (if the graph is a tree). We discuss several important avenues for future work below.

- **The regime** $\alpha \in [1/2, 3/4]$. In terms of sample complexity, our results show that the regime $\alpha \in (3/4, 1)$ is “easy”, i.e., the number of traces needed is minimal. On the other hand, $\alpha \in (0, 1/2)$ captures a “hard” regime, where the sample complexity nearly matches that of learning the entire graph exactly. What is the right sample complexity for $\alpha \in [1/2, 3/4]$? Is it in the easy regime, the hard regime, or something in between?
- **Estimation in general graphs when** $\alpha \in (0, 1/2)$. The Tree Reconstruction algorithm of [Abraham et al. \(2013\)](#) provides a means of estimating high-degree vertices when the graph is a tree. What about estimating high-degree vertices for general graphs in $\mathcal{G}(n, m, d, D)$?
- **Noisy observations.** In practice, the infection times may be noisy due to delays in reporting or due to quantization (e.g., positive cases being consolidated on a day-by-day basis). To what extent can our methods be adapted to handle these more realistic scenarios?
- **Other motifs.** A high-degree vertex can be viewed as a particular type of motif. What is the sample complexity of detecting and estimating other types of motifs (e.g., cliques)?
- **Other types of graphical models.** Our methods fundamentally leverage the temporal nature of the observed data. How can high-degree vertices be estimated in stationary graphical models, such as Gaussian graphical models or Ising models?

Acknowledgments

E.M. and A.S. acknowledge support from a Vannevar Bush Faculty Fellowship ONR-N00014-20-1-2826, the National Science Foundation under award CCF 1918421 and a Simons Investigator Award (622132). A.S. thanks Jorge Barreras Cortes for insightful discussions.

References

- Bruno Abrahamo, Flavio Chierichetti, Robert Kleinberg, and Alessandro Panconesi. Trace complexity of network inference. In *the Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 491–499, New York, NY, USA, 2013. Association for Computing Machinery.
- Eytan Adar and Lada A. Adamic. Tracking information epidemics in blogspace. In *the Proceedings of the IEEE/WIC/ACM Conference on Web Intelligence*, page 207–214, USA, 2005. IEEE Computer Society.
- HK Alexander and T Day. Risk factors for the evolutionary emergence of pathogens. *Journal of The Royal Society Interface*, 7(51):1455–1474, 2010.

- Antoine Allard, Pierre-André Noël, Louis J. Dubé, and Babak Pourbohloul. Heterogeneous bond percolation on multitype networks with an application to epidemic dynamics. *Physical Review E*, 79:036113, 2009.
- A. Auffinger, M. Damron, and J. Hanson. *50 Years of First Passage Percolation*. American Mathematical Society, 2017.
- Enric Boix-Adserà, Guy Bresler, and Frederic Koehler. Chow-liu++: Optimal prediction-centric learning of tree ising models. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 417–426, 2022.
- F. Brauer and C. Castillo-Chavez. *Mathematical Models in Population Biology and Epidemiology*. Springer, 2012.
- G. Bresler, E. Mossel, and A. Sly. Reconstruction of markov random fields from samples: Some easy observations and algorithms. In *the Proceedings of the 11th International Workshop, APPROX 2008, and 12th International Workshop, RANDOM 2008, LNCS 5171*, pages 343–356. Springer, 2008.
- Guy Bresler. Efficiently learning ising models on arbitrary graphs. In *the Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 771–782. ACM, 2015.
- Guy Bresler and Mina Karzand. Learning a tree-structured ising model in order to make predictions. *Annals of Statistics*, 48(2), 2020.
- Guy Bresler, Elchanan Mossel, and Allan Sly. Reconstruction of markov random fields from samples: Some observations and algorithms. *SIAM Journal on Computing*, 42(2):563–578, 2013.
- M. V. Burnashev and I. A. Begmatov. On a problem of signal detection leading to stable distributions. *Theory of Probability & Its Applications*, 35(3):556–560, 1991.
- Tony T. Cai and Yihong Wu. Optimal detection of sparse mixtures against a given null distribution. *IEEE Transactions on Information Theory*, 60(4):2217–2232, 2014. doi: 10.1109/TIT.2014.2304295.
- Tony T. Cai, Weidong Liu, and Harrison H. Zhou. Estimating sparse precision matrix: Optimal rates of convergence and adaptive estimation. *Annals of Statistics*, 44(2):455–488, 2016.
- Hadi Daneshmand, Manuel Gomez-Rodriguez, Le Song, and Bernhard Schoelkopf. Estimating diffusion network structures: Recovery conditions, sample complexity and soft-thresholding algorithm. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 793–801, Beijing, China, 22–24 Jun 2014. Proceedings of Machine Learning Research.
- Zoltán Dezső and Albert-László Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65:055103, May 2002.
- R. L. Dobrusin. A statistical problem arising in the theory of detection of signals in the presence of noise in a multi-channel system and leading to stable distribution laws. *Theory of Probability & Its Applications*, 3(2):161–173, 1958.

- David Donoho and Jiashun Jin. Higher criticism and for detecting sparse heterogeneous mixtures. *Annals of Statistics*, 32(3), 2004.
- Nan Du, Le Song, Ming Yuan, and Alex Smola. Learning networks of heterogeneous influence. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- Dean Eckles, Hossein Esfandiari, Elchanan Mossel, and M. Amin Rahimian. Seeding with costly network information. *Operations Research*, 70(4), 2022.
- Rashad Eletreby, Yong Zhuang, Kathleen M Carley, Osman Yağın, and H Vincent Poor. The effects of evolutionary adaptations on spreading processes in complex networks. *Proceedings of the National Academy of Sciences of the U.S.A.*, 117(11):5664–5670, 2020.
- Adrien Friggeri, Lada Adamic, Dean Eckles, and Justin Cheng. Rumor cascades. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1):101–110, May 2014.
- Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. Inferring networks of diffusion and influence. *ACM Trans. Knowl. Discov. Data*, 5(4), feb 2012. ISSN 1556-4681.
- Shushan He, Hongyuan Zha, and Xiaojing Ye. Network diffusions via neural mean-field dynamics. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2171–2183. Curran Associates, Inc., 2020.
- Oliver Hinz, Bernd Skiera, Christian Barrot, and Jan U. Becker. Seeding strategies for viral marketing: An empirical comparison. *Journal of Marketing*, 75(6):55 – 71, 2011.
- Jessica Hoffmann and Constantine Caramanis. Learning graphs from noisy epidemic cascades. *Proceedings of the ACM Measurement and Analysis of Computer Systems*, 3(2), jun 2019.
- Jessica Hoffmann, Soumya Basu, Surbhi Goel, and Constantine Caramanis. Learning mixtures of graphs from epidemic cascades. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4342–4352. PMLR, 13–18 Jul 2020.
- J.I. Ingster. *On Some Problems of Hypothesis Testing Leading to Infinitely Divisible Distributions*. Weierstraß-Inst. für Angewandte Analysis und Stochastik, 1996.
- Jonathan Kelner, Frederic Koehler, Raghu Meka, and Ankur Moitra. Learning some popular gaussian graphical models without condition number bounds. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10986–10998. Curran Associates, Inc., 2020.
- Justin Khim and Po-Ling Loh. A theory of maximum likelihood for weighted infection graphs. 2018. Preprint available at <https://arxiv.org/abs/1806.05273>.
- Adam R. Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 343–354. IEEE Computer Society, 2017.

- Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1(1):5–es, may 2007. ISSN 1559-1131.
- T.M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, volume 324. Springer, Berlin, Heidelberg, 1999.
- Rafael Lima. Hawkes processes modeling, inference, and control: An overview. *SIAM Review*, 65(2):331–374, 2023.
- Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34(3):1436 – 1462, 2006.
- Sidhant Misra, Marc Vuffray, and Andrey Y. Lokhov. Information theoretic optimal learning of gaussian graphical models. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2888–2909. PMLR, 09–12 Jul 2020.
- D. Mollison. Markovian Contact Processes. *Advances in Applied Probability*, 10(1):88–108, 1978.
- Elchanan Mossel and Sebastien Roch. Distance-based species tree estimation under the coalescent: Information-theoretic trade-off between number of loci and sequence length. *Annals of Applied Probability*, 27(5), October 2017.
- Praneeth Netrapalli and Sujay Sanghavi. Learning the graph of epidemic cascades. *SIGMETRICS Perform. Eval. Rev.*, 40(1):211–222, jun 2012. ISSN 0163-5999.
- Romualdo Pastor-Satorras and Alessandro Vespignani. Immunization of complex networks. *Phys. Rev. E*, 65:036104, Feb 2002.
- Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge University Press, 2024+. Preprint available at <http://www.stat.yale.edu/~yw562/teaching/itbook-export.pdf>.
- Pradeep Ravikumar, Martin J. Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935 – 980, 2011.
- D. Richardson. Random growth in a tessellation. *Proceedings of the Cambridge Philosophical Society*, 74:515–528, 1973.
- Everett M. Rogers. *Diffusion of Innovations*. Simon and Schuster, 5 edition, 2003.
- Narayana P. Santhanam and Martin J. Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.
- Galen R. Shorack and Jon A. Wellner. *Empirical Processes with Applications to Statistics*. Society for Industrial and Applied Mathematics, 1986.

- KD Spencer, CL Chung, and et al. Stargel A. COVID-19 case investigation and contact tracing efforts from health departments – united states, june 25 – july 24. *Morbidity and Mortality Weekly Report*, 2020.
- David Strang and Sarah A. Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24:265 – 290, 1998.
- Marc Vuffray, Sidhant Misra, and Andrey Likhov. Efficient learning of discrete graphical models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 13575–13585. Curran Associates, Inc., 2020.
- Jacco Wallinga and Peter Teunis. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, 6(160):509–516, 2004.
- Wei Wang, Martin J. Wainwright, and Kannan Ramchandran. Information-theoretic bounds on model selection for gaussian markov random fields. In *2010 IEEE International Symposium on Information Theory*, pages 1373–1377, 2010.
- Mateusz Wilinski and Andrey Likhov. Prediction-centric learning of independent cascade dynamics from partial observations. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11182–11192. PMLR, 18–24 Jul 2021.
- Shanshan Wu, Sujay Sanghavi, and Alexandros G Dimakis. Sparse logistic regression learns all discrete pairwise graphical models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Appendix A. Preliminaries: Cascade dynamics and infection times

In this section we derive several useful results on the behavior of the SI process (1) and the distribution of infection times. Throughout, we will assume that $\mathcal{I}(t)$ is a realization of $\mathcal{SI}(G, v_0)$ for a particular graph $G \in \mathcal{G}(n, m, d, D)$ and vertex $v_0 \in V$.

We begin by discussing an important equivalence between (1) and First Passage Percolation on G with $\text{Exp}(\lambda)$ edge weights. In a bit more detail, let $\{\mathcal{F}_t\}_{t \geq 0}$ denote the natural filtration corresponding to the SI process. Fix $t \geq 0$, and let $\{F(e)\}_{e \in E}$ be a collection of i.i.d. $\text{Exp}(\lambda)$ random variables that are independent of \mathcal{F}_t . For a path¹ P in the graph, we define the *weight* of the path to be $\text{weight}(P) := \sum_{e \in P} F(e)$. Then we have the following result.

Proposition 9 *Let $t \geq 0$, and condition on \mathcal{F}_t . Suppose that $v \in V \setminus \mathcal{I}(t)$. Let $\mathcal{P}(v, t)$ be the collection of paths from $\mathcal{I}(t)$ to v where only the first edge in the path has an endpoint in $\mathcal{I}(t)$. Then the conditional distribution of $T(v)$ with respect to \mathcal{F}_t is given by*

$$T(v) = t + \min_{P \in \mathcal{P}(v, t)} \text{weight}(P).$$

1. Recall that a path is a finite sequence of edges e_1, \dots, e_m such that for each $i \geq 1$, e_i and e_{i+1} share exactly one endpoint.

This result is standard, and essentially follows from the memoryless property of exponential distributions. We defer the reader to (Auffinger et al., 2017, Chapter 6), or to Section E for full details of the proof.

Using the representation in Proposition 9, we can derive a number of useful results on the infection times of vertices. We start by establishing probabilistic upper and lower bounds on conditional infection times of vertices.

Lemma 10 *Let $t \geq 0$, and condition on \mathcal{F}_t . Suppose that $v \in V \setminus \mathcal{I}(t)$. Then $T(v) - t$ is stochastically dominated by a $\text{Exp}(\lambda|\mathcal{N}(v) \cap \mathcal{I}(t)|)$ random variable.*

Proof Conditioned on \mathcal{F}_t , Proposition 9 implies that

$$T(v) - t = \min_{P \in \mathcal{P}(v,t)} \text{weight}(P) \leq \min_{P \in \mathcal{P}(v,t): |P|=1} \text{weight}(P) = \min_{u \in \mathcal{N}(v) \cap \mathcal{I}(t)} F(u, v).$$

Due to properties of exponential random variables, we have that

$$\min_{u \in \mathcal{N}(v) \cap \mathcal{I}(t)} F(u, v) \sim \text{Exp}(\lambda|\mathcal{N}(v) \cap \mathcal{I}(t)|),$$

which proves the claimed result. ■

Lemma 11 *Let $t \geq 0$, condition on \mathcal{F}_t , and suppose that $v \in V \setminus \mathcal{I}(t)$. Then $T(v) - t$ stochastically dominates a $\text{Exp}(\lambda|\mathcal{N}(v)|)$ random variable.*

Proof Conditioned on \mathcal{F}_t , we have that

$$T(v) - t = \min_{P \in \mathcal{P}(v,t)} \text{weight}(P) \geq \min_{u \in \mathcal{N}(v)} F(u, v),$$

where the inequality follows since any path ending in v must include one of the edges incident to v . The claim follows since $\min_{u \in \mathcal{N}(v)} F(u, v) \sim \text{Exp}(\lambda|\mathcal{N}(v)|)$. ■

We next define a useful event concerning the number of infected neighbors of high-degree vertices. We remark that the control of this quantity is especially important for us, as it is equal to the bias of our estimate of vertex degrees (see (5)).

Definition 12 *The event \mathcal{A} holds if and only if*

$$|\mathcal{N}(v) \cap \mathcal{I}(T(v))| \leq 2d\sqrt{n} \log^2 n,$$

for all $v \in V$.

Remark 13 *As we assume throughout this paper that $d = n^{o(1)}$, we will often use that $|\mathcal{N}(v) \cap \mathcal{I}(T(v))| \leq n^{1/2+o(1)}$ when using the event \mathcal{A} .*

Remark 14 *Assuming that $G \in \mathcal{G}(n, m, d, D)$, the event \mathcal{A} is only relevant for high-degree vertices, since*

$$|\mathcal{N}(v) \cap \mathcal{I}(T(v))| \leq \deg(v) = n^{o(1)} \leq 2d\sqrt{n} \log^2 n,$$

otherwise.

Lemma 15 *Suppose that $m = o(\log^2 n)$. Then $\mathbb{P}(\mathcal{A}) = 1 - o(1)$.*

Proof If $\deg(v) \leq 2d\sqrt{n}\log^2 n$, then

$$|\mathcal{N}(v) \cap \mathcal{I}(T(v))| \leq |\mathcal{N}(v)| \leq 2d\sqrt{n}\log^2 n,$$

which satisfies the condition in \mathcal{A} corresponding to v . It therefore suffices to consider the case $\deg(v) \geq 2d\sqrt{n}\log^2 n$. Define the event

$$\mathcal{E}_v := \{|\mathcal{N}(v) \cap \mathcal{I}(T(v))| \geq 2d\sqrt{n}\log^2 n\},$$

and let T be the stopping time at which $\sqrt{n}\log n$ neighbors of v are first infected.

Let $t \geq 0$, and condition on \mathcal{F}_t . Let us also assume that $T = t$ (an event that is \mathcal{F}_t -measurable), so that $|\mathcal{N}(v) \cap \mathcal{I}(t)| = \sqrt{n}\log n$. If $T(v) \leq t$, then we have that $\mathbb{P}(\mathcal{E}_v | \mathcal{F}_t) = 0$. We therefore proceed by considering the case where $T(v) > t$. Notice that if $T = t$,

$$|\mathcal{N}(v) \cap \mathcal{I}(T(v))| = |\mathcal{N}(v) \cap \mathcal{I}(t)| + |\mathcal{N}(v) \cap \mathcal{I}[t, T(v)]| = \sqrt{n}\log n + |\mathcal{N}(v) \cap \mathcal{I}[t, T(v)]|.$$

As a result, we have that

$$\mathbb{P}(\mathcal{E}_v | \mathcal{F}_t) \mathbf{1}(T = t) \leq \mathbb{P}(|\mathcal{N}(v) \cap \mathcal{I}[t, T(v)]| \geq d\sqrt{n}\log^2 n | \mathcal{F}_t) \mathbf{1}(T = t). \quad (12)$$

To bound the probability on the right hand side, we start by bounding $T(v)$. Lemma 10 implies that $T(v) - t$ is stochastically dominated by a $\text{Exp}(\lambda\sqrt{n}\log n)$ random variable. Hence

$$\mathbb{P}\left(T(v) \geq t + \frac{1}{\lambda\sqrt{n}} \middle| \mathcal{F}_t\right) \leq \frac{1}{n}. \quad (13)$$

We can now bound the probability of interest as follows.

$$\begin{aligned} \mathbb{P}(|\mathcal{N}(v) \cap \mathcal{I}[t, T(v)]| \geq d\sqrt{n}\log^2 n | \mathcal{F}_t) & \\ & \stackrel{(a)}{\leq} \mathbb{P}(T(v) \geq t' | \mathcal{F}_t) + \mathbb{P}(|\mathcal{N}(v) \cap \mathcal{I}[t, t']| \geq d\sqrt{n}\log^2 n | \mathcal{F}_t) \\ & \stackrel{(b)}{\leq} \frac{1}{n} + \frac{\mathbb{E}[|\mathcal{N}(v) \cap \mathcal{I}[t, t']| | \mathcal{F}_t]}{d\sqrt{n}\log^2 n}. \end{aligned} \quad (14)$$

Above, (a) is due to a union bound, and since $|\mathcal{N}(v) \cap \mathcal{I}[t, T(v)]| \leq |\mathcal{N}(v) \cap \mathcal{I}[t, t']|$ whenever $T(v) < t'$. The inequality (b) is due to (13) and Markov's inequality.

We proceed by bounding $\mathbb{E}[|\mathcal{N}(v) \cap \mathcal{I}[t, t']| | \mathcal{F}_t]$; to do so, we study the probability that a given $u \in \mathcal{N}(v) \setminus \mathcal{I}(t)$ is infected in the interval $[t, t']$. If $u \in \text{highdeg}(G)$, we trivially bound the probability by 1. If $u \notin \text{highdeg}(G)$, $\deg(u) \leq d$ so by Lemma 11, $T(u) - t$ stochastically dominates a $\text{Exp}(\lambda d)$ random variable. As a result, for such u , we have that

$$\mathbb{P}(T(u) \in [t, t'] | \mathcal{F}_t) \leq 1 - e^{-\lambda(t'-t)} \leq \frac{d}{\sqrt{n}}.$$

In the final inequality above, we have used $1 - e^{-x} \leq x$ as well as the definition of t' . It follows that

$$\mathbb{E}[|\mathcal{N}(v) \cap \mathcal{I}[t, t']| | \mathcal{F}_t] = \sum_{u \in \mathcal{N}(v) \setminus \mathcal{I}(t)} \mathbb{P}(T(u) \in [t, t'] | \mathcal{F}_t) \leq m + d\sqrt{n} \leq 2d\sqrt{n}. \quad (15)$$

In the first inequality above, we have used that there are at most m high-degree vertices in $\mathcal{N}(v)$ and at most n low-degree vertices. Together, (12), (14) and (15) show that

$$\mathbb{P}(\mathcal{E}_v | \mathcal{F}_t) \mathbf{1}(T = t) \leq \frac{3}{\log^2 n} \mathbf{1}(T = t).$$

Since $T < \infty$ almost surely (since all vertices become infected in finite time), it holds that

$$\mathbb{P}(\mathcal{E}_v | \mathcal{F}_T) = \int_0^\infty \mathbb{P}(\mathcal{E}_v | \mathcal{F}_t) \mathbf{1}(T = t) dt \leq \frac{3}{\log^2 n}.$$

Taking an expectation over \mathcal{F}_T shows that $\mathbb{P}(\mathcal{E}_v) \leq 3/\log^2(n)$. Finally, we note that

$$\mathcal{A}^c = \bigcup_{v \in V: \deg(v) \geq 2d\sqrt{n} \log^2 n} \mathcal{E}_v,$$

hence a union bound over these at most m events shows that

$$\mathbb{P}(\mathcal{A}^c) \leq \frac{3m}{\log n} = o(1).$$

■

As the following result shows, the event \mathcal{A} is crucially used to obtain sharper characterizations of infection times than what Lemma 11 can provide, especially for high-degree vertices.

Lemma 16 *Let $v \in V$ and let $\delta > 0$ satisfy $\lambda^2 \delta^2 mn \leq 1$. Fix $t \geq 0$, and let $v \in V \setminus \mathcal{I}(t)$. Then*

$$\mathbb{P}(\{T(v) \in [t, t + \delta]\} \cap \mathcal{A} | \mathcal{F}_t) \leq 3d\lambda\delta\sqrt{n} \log^2(n).$$

Proof Define the event

$$\mathcal{A}_t := \{\text{For all } v \in V \text{ such that } T(v) > t, \text{ it holds that } |\mathcal{N}(v) \cap \mathcal{I}(t)| \leq 2d\sqrt{n} \log^2 n\}.$$

Notice that \mathcal{A}_t is \mathcal{F}_t -measurable, and $\mathcal{A} \subset \mathcal{A}_t$. Moreover, for v such that $T(v) > t$,

$$\mathbb{P}(\{T(v) \leq t + \delta\} \cap \mathcal{A} | \mathcal{F}_t) \leq \mathbb{P}(\{T(v) \leq t + \delta\} \cap \mathcal{A}_t | \mathcal{F}_t) = \mathbb{P}(T(v) \leq t + \delta | \mathcal{F}_t) \mathbf{1}(\mathcal{A}_t). \quad (16)$$

We will therefore proceed by bounding $\mathbb{P}(T(v) \leq t + \delta | \mathcal{F}_t)$ under the assumption that \mathcal{A}_t holds.

Suppose that $\deg(v) \leq d$. Then by Lemma 11, $T(v) - t$ stochastically dominates a $\text{Exp}(\lambda d)$ random variable, so that

$$\mathbb{P}(T(v) \leq t + \delta | \mathcal{F}_t) \leq 1 - e^{-\delta\lambda d} \leq \delta\lambda d. \quad (17)$$

Next, suppose that $\deg(v) \geq D$, and let $\mathcal{P}_2(v)$ be the set of paths of length 2 starting from v . Then by Proposition 9,

$$T(v) - t = \min_{P \in \mathcal{P}(v,t)} \text{weight}(P) \quad (18)$$

$$\geq \min \left\{ \min_{u \in \mathcal{N}(v) \cap \mathcal{I}(t)} F(u, v), \min_{P \in \mathcal{P}_2(v)} \text{weight}(P) \right\}, \quad (19)$$

which can be explained as follows. If the minimizing path $P \in \mathcal{P}(v, t)$ in (18) has length 1, then the minimum is achieved by $F(u, v)$ for some $u \in \mathcal{N}(v) \cap \mathcal{I}(t)$. Otherwise, the minimizing path in (18) has length at least 2, which must include a path of length 2 starting at v . In light of (19),

$$\begin{aligned}
 \mathbb{P}(T(v) \leq t + \delta | \mathcal{F}_t) &\leq \mathbb{P}\left(\min\left\{\min_{u \in \mathcal{N}(v) \cap \mathcal{I}(t)} F(u, v), \min_{P \in \mathcal{P}_2(v)} \text{weight}(P)\right\} \leq \delta \mid \mathcal{F}_t\right) \\
 &\leq \sum_{u \in \mathcal{N}(v) \cap \mathcal{I}(t)} \mathbb{P}(F(u, v) \leq \delta | \mathcal{F}_t) + \sum_{P \in \mathcal{P}_2(v)} \mathbb{P}(\text{weight}(P) \leq \delta | \mathcal{F}_t) \\
 &\stackrel{(a)}{\leq} |\mathcal{N}(v) \cap \mathcal{I}(t)| \lambda \delta + |\mathcal{P}_2(v)| \lambda^2 \delta^2 \\
 &\stackrel{(b)}{\leq} 2d\lambda\delta\sqrt{n} \log^2(n) + (dn + mn)\lambda^2 \delta^2 \\
 &\stackrel{(c)}{\leq} 3d\lambda\delta\sqrt{n} \log^2(n), \tag{20}
 \end{aligned}$$

where the inequality (a) follows since $F(u, v) \sim \text{Exp}(\lambda)$ and $\text{weight}(P) \sim \text{Gamma}(2, \lambda)$ for $P \in \mathcal{P}_2(v)$. The inequality (b) uses the bound $|\mathcal{N}(v) \cap \mathcal{I}(t)| \leq 2d\sqrt{n} \log^2 n$ which holds on the event \mathcal{A}_t , and also uses that the number of paths of length 2 starting from v is at most $dn + mn$, since there are at most n neighbors of degree at most d and at most m neighbors of degree at most n . Finally, (c) holds by our assumption that $\lambda^2 \delta^2 mn \leq 1$. \blacksquare

A useful consequence of Lemma 16 is a bound on the probability that $T(u)$ and $T(v)$ are close, for distinct vertices u, v .

Lemma 17 *Let $u, v \in V$ be distinct vertices and let $\delta > 0$ satisfy $\lambda^2 \delta^2 n \leq 1$. Then*

$$\mathbb{P}(\{|T(u) - T(v)| \leq \delta\} \cap \mathcal{A}) \leq 6d\lambda\delta\sqrt{n} \log^2(n).$$

Proof Suppose without loss of generality that $T(u) < T(v)$, so that $|T(u) - T(v)| \leq \delta$ is equivalent to $T(v) \in [T(u), T(u) + \delta]$. Then by Lemma 16,

$$\begin{aligned}
 \mathbb{P}(\{T(v) \in [T(u), T(u) + \delta]\} \cap \mathcal{A} | \mathcal{F}_t) \mathbf{1}(T(u) = t) &= \mathbb{P}(\{T(v) \in [t, t + \delta]\} \cap \mathcal{A} | \mathcal{F}_t) \mathbf{1}(T(u) = t) \\
 &\leq 3d\lambda\delta\sqrt{n} \log^2(n).
 \end{aligned}$$

Since $T(u) < \infty$ almost surely, we have that

$$\begin{aligned}
 &\mathbb{P}(\{T(v) \in [T(u), T(u) + \delta]\} \cap \mathcal{A} | \mathcal{F}_{T(u)}) \\
 &= \int_0^\infty \mathbb{P}(\{T(v) \in [T(u), T(u) + \delta]\} \cap \mathcal{A} | \mathcal{F}_t) \mathbf{1}(T(u) = t) dt \leq 3d\lambda\delta\sqrt{n} \log^2(n).
 \end{aligned}$$

Finally, the claim follows from taking an expectation over $\mathcal{F}_{T(u)}$. \blacksquare

Lemma 18 *For any $0 \leq t_1 < t_2 < \infty$, $|\mathcal{I}(t_2) \setminus \mathcal{I}(t_1)|$ is stochastically bounded by a $\text{Poi}(\lambda m n d(t_2 - t_1))$ random variable.*

Proof Recall that $|E|$ is the number of edges in the graph, and let A_1, A_2, \dots be a sequence of i.i.d. $\text{Exp}(\lambda|E|)$ random variables. As a shorthand, for a positive integer k , let τ_k denote the time at which the k th vertex is infected, after time t_1 . We claim that it suffices to show that, conditioned on \mathcal{F}_{t_1} ,

$$\tau_k - t_1 \succcurlyeq \sum_{i=1}^k A_i. \quad (21)$$

With (21) in hand, it is readily seen that the arrival times of vertices after time t_1 stochastically bounds the arrival times of a homogeneous Poisson point process with rate $\lambda|E|$. As a result, $|\mathcal{I}(t_2) \setminus \mathcal{I}(t_1)|$ is stochastically bounded by a $\text{Poi}(\lambda|E|(t_2 - t_1))$ random variable. Finally, the claim follows since we can upper bound

$$|E| \leq \sum_{v \in V} \deg(v) \leq nd + mn \leq mnd,$$

where the second inequality above follows since there are at most n vertices of degree at most d and there are m vertices of degree at most n .

It remains to prove (21); we do so by induction. Let $E(t_1) \subset E$ be the set of edges where one endpoint is in $\mathcal{I}(t_1)$ and the other is not in $\mathcal{I}(t_1)$. Since the next infection must be a vertex which neighbors a currently infected vertex, we have by Proposition 9 that $\tau_2 - t_1 \geq \min_{e \in E(t_1)} F(e)$, where we recall that $\{F(e)\}_{e \in E}$ is a collection of i.i.d. $\text{Exp}(\lambda)$ random variables. Notice that by properties of Exponential random variables, $\min_{e \in E(t_1)} F(e)$ is equal in distribution to a $\text{Exp}(\lambda|E(t_1)|)$ random variable, which in turn stochastically bounds a $\text{Exp}(\lambda|E|)$ random variable. This proves the base case.

Now let $k \geq 1$, and assume that (21) holds for k . We may assume without loss of generality that $\tau_{k+1} < \infty$ almost surely. Otherwise, since G is connected, this must mean that there are no more infection events that can happen after the k th infection. In this case, $\tau_{k+1} = \infty$, hence (21) trivially holds with k replaced with $k + 1$. Let us now proceed with the assumption that $\tau_{k+1} < \infty$ almost surely. Fix $s \geq t_1$, condition on \mathcal{F}_s , and let us assume that $\tau_k = s$ (an event that is \mathcal{F}_s -measurable). Repeating the same arguments above, we see that, conditioned on \mathcal{F}_s and assuming $\tau_k = s$, we have that $\tau_{k+1} - s$ stochastically dominates a $\text{Exp}(\lambda|E|)$ random variable. Hence, letting $x \geq 0$ and $A \sim \text{Exp}(\lambda|E|)$, we have that

$$\mathbb{P}(A \geq x) \mathbf{1}(\tau_k = s) \leq \mathbb{P}(\tau_{k+1} - s \geq x | \mathcal{F}_s) \mathbf{1}(\tau_k = s) = \mathbb{P}(\tau_{k+1} - \tau_k \geq x | \mathcal{F}_s) \mathbf{1}(\tau_k = s).$$

Taking an integral over $0 \leq s < \infty$, we have that

$$\mathbb{P}(A \geq x) \leq \int_0^\infty \mathbb{P}(\tau_{k+1} - \tau_k \geq x | \mathcal{F}_s) \mathbf{1}(\tau_k = s) ds = \mathbb{P}(\tau_{k+1} - \tau_k \geq x | \mathcal{F}_{\tau_k}).$$

In particular, $\tau_{k+1} - \tau_k \succcurlyeq A$ conditioned on τ_k . As A is independent of τ_k and (21) holds, it immediately follows that (21) also holds with k replaced with $k + 1$. \blacksquare

Appendix B. Analysis of Algorithm 1: Proof of Theorem 6

The following lemma, which characterizes when $\widehat{\deg}_\delta(v)$ is a good approximation for $\deg(v)$, is the main result of this section.

Lemma 19 *Suppose that $\beta, \gamma \in (0, 1)$ satisfy $\gamma > \max\{1 - \beta, (1 + \beta)/2, \beta\}$, and let $\delta := n^{-\beta}$. It holds with probability $1 - o(1)$ that for all $v \in V$, at least one of the following conditions is true:*

$$(C1) \quad \left| \widehat{\deg}_\delta(v) - \lambda \deg(v) \right| \leq n^\gamma.$$

$$(C2) \quad \text{There exists } u \in \text{highdeg}(G) \setminus \{v\} \text{ such that } |T(v) - T(u)| < n^{-\beta}.$$

Before proving Lemma 19, we show how it can be used to prove Theorem 6. We first analyze what can be learned from a single cascade.

Lemma 20 *Suppose that $\alpha, \beta \in (0, 1)$ satisfy $\alpha > \max\{1 - \beta, (1 + \beta)/2, \beta\}$ and $\beta > 1/2$. Define*

$$S := \left\{ v : \widehat{\deg}(v) \geq \frac{n^\alpha}{\log n} \right\}. \quad (22)$$

Then with probability $1 - o(1)$,

$$\text{highdeg}(G) \subseteq S \subseteq \left\{ u \in V : \exists v \in \text{highdeg}(G) \text{ with } |T(u) - T(v)| \leq \frac{1}{n^\beta} \right\}. \quad (23)$$

Proof Throughout this proof, we will assume that the event described in Lemma 19 holds. We may do so since this event holds with probability $1 - o(1)$.

We start by proving that the first inclusion in (23) holds with probability $1 - o(1)$. To this end, define the event

$$\mathcal{E} := \left\{ \text{For all distinct } u, v \in \text{highdeg}(G), \text{ it holds that } |T(u) - T(v)| > \frac{1}{n^\beta} \right\}.$$

Suppose that \mathcal{E} holds, and choose γ such that $\alpha > \gamma > \max\{1 - \beta, (1 + \beta)/2, \beta\}$. Then for all $v \in \text{highdeg}(G)$, we must have that $|\widehat{\deg}(v) - \lambda \deg(v)| \leq n^\gamma$, since the condition (C2) in Lemma 19 cannot hold on \mathcal{E} as $v \in \text{highdeg}(G)$. Hence

$$\widehat{\deg}(v) \geq \lambda \deg(v) - n^\gamma \geq \lambda n^\alpha - n^\gamma \geq \frac{n^\alpha}{\log n}, \quad (24)$$

where we have used the fact that λ is constant with respect to n , and that $\alpha > \gamma$. Moreover, an important consequence of (24) is that $\text{highdeg}(G) \subseteq S$. To complete the proof of the first inclusion in (23), we need to show that \mathcal{E} holds with probability $1 - o(1)$. To this end, let us recall the event \mathcal{A} (see Definition 12). We can bound

$$\mathbb{P}(\mathcal{E}^c \cap \mathcal{A}) \leq \sum_{u, v \in \text{highdeg}(G): u \neq v} \mathbb{P} \left(\left\{ |T(u) - T(v)| \leq \frac{1}{n^\beta} \right\} \cap \mathcal{A} \right) \stackrel{(a)}{\leq} n^{-(\beta-1/2)+o(1)} \stackrel{(b)}{=} o(1). \quad (25)$$

Above, (a) follows since there are at most $m^2 = n^{o(1)}$ terms in the summation, and each term in the summation can be bounded via Lemma 17. The equality (b) holds under our assumption that $\beta > 1/2$. By a union bound, we have that

$$\mathbb{P}(\mathcal{E}^c) \leq \mathbb{P}(\mathcal{E}^c \cap \mathcal{A}) + \mathbb{P}(\mathcal{A}^c) = o(1),$$

where the bound on the right hand side is due to (25) and Lemma 15.

We now turn to the proof of the second inclusion in (23). As a shorthand, let us denote

$$S' := \left\{ u \in V : \exists v \in \text{highdeg}(G) \text{ with } |T(u) - T(v)| \leq \frac{1}{n^\beta} \right\}.$$

Let $u \in S$. If $u \in \text{highdeg}(G)$, then we trivially have that $u \in S'$. On the other hand, if $u \notin \text{highdeg}(G)$, then $\deg(u) \leq n^{o(1)}$ so, under condition (C1) of Lemma 19, it cannot be the case that $u \in S$. Thus, (C2) must hold, which implies that $u \in S'$ in this case as well. Both cases prove that $S \subseteq S'$. \blacksquare

We are now ready to prove the main result of this section.

Proof [Proof of Theorem 6] Let S_i denote the version of the set in (22) corresponding to the i th cascade, for $i \in [K]$. In particular, we have that $\widehat{\text{HD}} = \bigcap_{i=1}^K S_i$, where $\widehat{\text{HD}}$ is the output of Algorithm 1. With this notation in hand, it suffices to show that $\bigcap_{i=1}^K S_i = \text{highdeg}(G)$ with probability $1 - o(1)$. By Lemma 20, we have with probability $1 - o(1)$ that $\bigcap_{i=1}^K S_i \supseteq \text{highdeg}(G)$, provided K is constant with respect to n . In the remainder of the proof, we therefore focus on proving that $\bigcap_{i=1}^K S_i \subseteq \text{highdeg}(G)$ with high probability. It will be useful to define \mathcal{A}_i to be the event described in Definition 12 which concerns the i th cascade, for each $i \in [K]$.

We proceed by bounding the probability that there exists a vertex u satisfying $u \notin \text{highdeg}(G)$ and $u \in S_i$, for a fixed $i \in [K]$. If there exists such a vertex, then by the characterization in Lemma 20, there must exist $v \in \text{highdeg}(G)$ such that $|T_i(u) - T_i(v)| \leq n^{-\beta}$. We have that

$$\begin{aligned} \mathbb{P}(\{u \in S_i\} \cap \mathcal{A}_i) &\leq \mathbb{P}\left(\left\{\exists v \in \text{highdeg}(G) : |T_i(u) - T_i(v)| \leq \frac{1}{n^\beta}\right\} \cap \mathcal{A}_i\right) \\ &\leq \sum_{v \in \text{highdeg}(G)} \mathbb{P}\left(\left\{|T_i(u) - T_i(v)| \leq \frac{1}{n^\beta}\right\} \cap \mathcal{A}_i\right) \\ &\leq n^{-(\beta-1/2)+o(1)}. \end{aligned}$$

Above, the final inequality follows from Lemma 17, and since $|\text{highdeg}(G)| \leq m = n^{o(1)}$. Next, by the independence of cascades, we have that for $u \notin \text{highdeg}(G)$,

$$\mathbb{P}\left(\left\{u \in \bigcap_{i=1}^K S_i\right\} \cap \bigcap_{i=1}^K \mathcal{A}_i\right) = \prod_{i=1}^K \mathbb{P}(\{u \in S_i\} \cap \mathcal{A}_i) = n^{-K(\beta-1/2)+o(1)}.$$

Putting everything together, we have that

$$\begin{aligned} \mathbb{P}\left(\left\{\left|\left(\bigcap_{i=1}^K S_i\right) \setminus \text{highdeg}(G)\right| \geq 1\right\} \cap \bigcap_{i=1}^K \mathcal{A}_i\right) &\leq \mathbb{E}\left[\left|\left(\bigcap_{i=1}^K S_i\right) \setminus \text{highdeg}(G)\right| \mathbf{1}\left(\bigcap_{i=1}^K \mathcal{A}_i\right)\right] \\ &= \sum_{u \notin \text{highdeg}(G)} \mathbb{P}\left(\left\{u \in \bigcap_{i=1}^K S_i\right\} \cap \bigcap_{i=1}^K \mathcal{A}_i\right) \\ &= n^{1-K(\beta-1/2)+o(1)} \\ &= o(1), \end{aligned}$$

where (a) is due to Markov's inequality, and (b) holds whenever $K > 1/(\beta - 1/2)$. Finally, the theorem follows from setting $\beta = \alpha - 1/4$, and by noting that $\mathbb{P}(\bigcap_{i=1}^K \mathcal{A}_i) = 1 - o(1)$ in light of Lemma 15. \blacksquare

In the rest of this section, we will focus on proving Lemma 19. For our analysis, it will be convenient to define the quantities

$$\widehat{dI}_\beta(t) := n^\beta \left(I \left[t, t + \frac{1}{n^\beta} \right] - I \left[t - \frac{1}{n^\beta}, t \right] \right) \quad (26)$$

$$dI(t) := \lambda (\text{cut}(\mathcal{I}(t)) - \text{cut}(\mathcal{I}(t^-))). \quad (27)$$

In words, $\widehat{dI}_\beta(t)$ and $dI(t)$ are the discretized and expected second derivatives of the infection curve, respectively. Recall that we can relate these to the quantities of interest via $\widehat{\text{deg}}_\delta(v) = \widehat{dI}_\beta(T(v))$ for $\delta = n^{-\beta}$ and $dI(T(v)) = \lambda(\text{deg}(v) - 2|\mathcal{N}(v) \cap \mathcal{I}(T(v))|)$.

Our first intermediate result towards proving Lemma 19 controls the difference between $\widehat{dI}_\beta(t)$ and $dI(t)$ for a fixed $t \geq 0$.

Lemma 21 *Let $\beta, \gamma \in (0, 1)$ satisfy $\gamma > \max\{1 - \beta, (1 + \beta)/2\}$. Let us also define the event*

$$\mathcal{B}_t := \left\{ \forall v \in \text{highdeg}(G), T(v) \notin \left(t - \frac{1}{n^\beta}, t \right) \cup \left(t, t + \frac{1}{n^\beta} \right) \right\}. \quad (28)$$

Then for any $t \geq 0$, it holds that

$$\mathbb{P} \left(\left\{ \left| \widehat{dI}_\beta(t) - dI(t) \right| \geq n^\gamma \right\} \cap \mathcal{B}_t \right) \leq \exp \left(-n^{2\gamma - (1 + \beta) + o(1)} \right).$$

Since the proof of the lemma is somewhat involved, we defer the details to Section C.

Our next goal is to show that $\widehat{dI}_\beta(t)$ and $dI(t)$ are close for all $t \geq 0$ whenever \mathcal{B}_t holds, rather than for just a fixed $t \geq 0$. To prove this, we will employ a discretization argument. Define the finite set of time indices

$$\mathbb{T}_n := \left\{ \frac{k}{n^5} : k \in \mathbb{Z}_{\geq 0} \right\} \cap [0, n^2].$$

The following result shows that the discretization \mathbb{T}_n is rich enough to capture the continuous-time dynamics of the cascade.

Lemma 22 *With probability $1 - o(1)$, for each $v \in V$ there exists a unique $t \in \mathbb{T}_n$ such that $t - 1/n^5 \leq T(v) \leq t$.*

Proof We first show that $\max_{v \in V} T(v) \leq n^2$. From the representation in Proposition 9, we can upper bound $\max_{v \in V} T(v)$ by $\sum_{e \in E} F(e)$, so we focus on characterizing the latter quantity. In particular, we have that

$$\mathbb{E} \left[\sum_{e \in E} F(e) \right] = \frac{|E|}{\lambda}.$$

We can upper bound the expectation by noting that $|E| \leq dn + mn = n^{1+o(1)}$, which holds since there are at most n vertices of degree at most d and at most m vertices of degree at most n for $G \in \mathcal{G}(n, m, d, D)$. Markov's inequality now implies

$$\mathbb{P} \left(\max_{v \in V} T(v) \geq n^2 \right) \leq n^{-1+o(1)}.$$

Hence there must exist some interval of the form $[t, t + 1/n^5]$ which contains $T(v)$, for every $v \in V$.

We next show that in each interval of the form $[t, t + 1/n^5]$, there exists at most one infection event. Condition on \mathcal{F}_t . If there are at least 2 infection events in $[t, t + 1/n^5]$, then there must exist distinct $e_1, e_2 \in E$ such that $F(e_1), F(e_2) \leq n^{-5}$ (if not, at most one vertex is infected in the interval). The probability of this occurring is at most

$$\binom{|E|}{2} \left(1 - e^{-\lambda/n^5}\right)^2 \leq n^{2+o(1)} \left(\frac{\lambda}{n^5}\right)^2 = n^{-8+o(1)}.$$

We conclude by taking a union bound over the $O(n^7)$ elements of \mathbb{T}_n . ■

Together, Lemmas 21 and 22 prove Lemma 19.

Proof [Proof of Lemma 19] We start by relating maximal deviations between $\widehat{\deg}_\delta(v)$ and $\lambda \deg(v)$ to maximal deviations between $\widehat{dI}_\beta(t)$ and $dI(t)$. Recall that $\widehat{dI}_\beta(T(v)) = \widehat{\deg}_\delta(v)$ for $\delta = n^{-\beta}$ and that $dI(T(v)) = \lambda(\deg(v) - 2|\mathcal{N}(v) \cap \mathcal{I}(T(v))|)$. By the triangle inequality, we have that

$$\left| \widehat{\deg}_\delta(v) - \lambda \deg(v) \right| \leq \left| \widehat{dI}_\beta(T(v)) - dI(T(v)) \right| + 2\lambda |\mathcal{N}(v) \cap \mathcal{I}(T(v))|.$$

Recall the event \mathcal{A} from Definition 12 as well as the event \mathcal{B}_t from (28). On \mathcal{A} , it holds that

$$\begin{aligned} & \max_{v \in V} \left| \widehat{\deg}_\delta(v) - \lambda \deg(v) \right| \mathbf{1}(\mathcal{B}_{T(v)}) \\ & \leq \max_{v \in V} \left| \widehat{dI}_\beta(T(v)) - dI(T(v)) \right| \mathbf{1}(\mathcal{B}_{T(v)}) + 2\lambda \max_{v \in V} |\mathcal{N}(v) \cap \mathcal{I}(T(v))| \\ & \leq \sup_{t \geq 0} \left| \widehat{dI}_\beta(t) - dI(t) \right| \mathbf{1}(\mathcal{B}_t) + 4\lambda d\sqrt{n} \log^2 n. \end{aligned} \quad (29)$$

In light of (29), we have for any γ satisfying the conditions of the lemma that

$$\begin{aligned} & \mathbb{P} \left(\max_{v \in V} \left| \widehat{\deg}_\delta(v) - \lambda \deg(v) \right| \mathbf{1}(\mathcal{B}_{T(v)}) \geq n^\gamma \right) \\ & \leq \mathbb{P}(\mathcal{A}^c) + \mathbb{P} \left(\sup_{t \geq 0} \left| \widehat{dI}_\beta(t) - dI(t) \right| \mathbf{1}(\mathcal{B}_t) \geq n^\gamma - 4\lambda d\sqrt{n} \log^2 n \right) \\ & \leq o(1) + \mathbb{P} \left(\sup_{t \geq 0} \left| \widehat{dI}_\beta(t) - dI(t) \right| \mathbf{1}(\mathcal{B}_t) \geq n^{\gamma'} \right). \end{aligned} \quad (30)$$

In the second line, the $o(1)$ bound is due to Lemma 15, and we have used the fact that $\gamma > 1/2$ to argue that $n^\gamma - 4\lambda d\sqrt{n} \log^2 n \geq n^{\gamma'}$ for some $\gamma' < \gamma$ satisfying the same conditions as γ in the lemma statement. In the remainder of the proof, we show that the probability on the right hand side in (30) is also $o(1)$. Once this is shown, the proof is concluded, since the event

$$\left\{ \max_{v \in V} \left| \widehat{\deg}_\delta(v) - \lambda \deg(v) \right| \mathbf{1}(\mathcal{B}_{T(v)}) \leq n^\gamma \right\}$$

is equivalent to the one described in the statement of Lemma 19.

We therefore proceed by upper bounding the probability on the right hand side of (30). Assume that the event described in the statement of Lemma 22 holds. Fix $t \geq 0$, and let $t' \in \mathbb{T}_n$ satisfy $t' \leq t < t' + n^{-5}$. We first claim that $dI(t) = dI(t')$ or $dI(t) = dI(t' + n^{-5})$. To see why, we consider three cases. In the first case, suppose that no vertex is infected in the interval $[t', t' + n^{-5})$. Then there is no change in $\mathcal{I}(t)$ across the entire interval, hence $dI(t) = dI(t') = dI(t' + n^{-5})$. In the second case, suppose that there is a single infection event in the interval, which occurs in the sub-interval $[t', t)$. Then there is no change in $\mathcal{I}(t)$ in the interval $[t, t' + n^{-5}]$, hence $dI(t) = dI(t' + n^{-5})$. Finally, in the third case, suppose again that there is a single infection event in the interval, which occurs in the sub-interval $[t, t' + n^{-5})$. By analogous arguments, it follows that $dI(t) = dI(t')$. There are no other cases since, on the event described in the statement of Lemma 22, at most one infection event can occur in $[t', t' + n^{-5})$.

Next, we study $\widehat{dI}(t)$. By the triangle inequality,

$$\begin{aligned} \frac{1}{n^\beta} \left| \widehat{dI}(t) - \widehat{dI}(t') \right| &\leq \left| I \left[t, t + \frac{1}{n^\beta} \right] - I \left[t', t' + \frac{1}{n^\beta} \right] \right| + \left| I \left[t - \frac{1}{n^\beta}, t \right] - I \left[t' - \frac{1}{n^\beta}, t' \right] \right| \\ &= \left| I \left[t' + \frac{1}{n^\beta}, t + \frac{1}{n^\beta} \right] - I[t', t] \right| + \left| I[t', t] - I \left[t' - \frac{1}{n^\beta}, t - \frac{1}{n^\beta} \right] \right| \\ &\leq I \left[t' + \frac{1}{n^\beta}, t + \frac{1}{n^\beta} \right] + I \left[t' - \frac{1}{n^\beta}, t - \frac{1}{n^\beta} \right] + 2I[t', t]. \end{aligned}$$

Notice that the interval $[t' + n^{-\beta}, t + n^{-\beta}]$ intersects at most 2 intervals of the form $[a, b]$ for $a, b \in \mathbb{T}_n$, so we can bound $I[t' + n^{-\beta}, t + n^{-\beta}]$ by 2. The same argument holds for the intervals $[t' - n^{-\beta}, t - n^{-\beta}]$ and $[t', t]$, leading to the bound $|\widehat{dI}(t) - \widehat{dI}(t')| \leq 8n^\beta$. Identical arguments show that $|\widehat{dI}(t) - \widehat{dI}(t' + n^{-5})| \leq 8n^\beta$ as well.

Combining our results on $dI(t)$ and $\widehat{dI}(t)$, we see that

$$\left| \widehat{dI}(t) - dI(t) \right| \leq \max \left\{ \left| \widehat{dI}(t') - dI(t') \right|, \left| \widehat{dI} \left(t' + \frac{1}{n^5} \right) - dI \left(t' + \frac{1}{n^5} \right) \right| \right\} + 8n^\beta. \quad (31)$$

We can therefore bound the probability of interest as follows.

$$\begin{aligned} \mathbb{P} \left(\sup_{t \geq 0} \left| \widehat{dI}(t) - dI(t) \right| \mathbf{1}(\mathcal{B}_t) \geq n^\gamma \right) &\stackrel{(a)}{\leq} \mathbb{P} \left(\max_{t \in \mathbb{T}_n} \left| \widehat{dI}(t) - dI(t) \right| \mathbf{1}(\mathcal{B}_t) \geq n^\gamma - 8n^\beta \right) \\ &\leq \sum_{t \in \mathbb{T}_n} \mathbb{P} \left(\left\{ \left| \widehat{dI}(t) - dI(t) \right| \geq n^\gamma - 8n^\beta \right\} \cap \mathcal{B}_t \right) \\ &\stackrel{(b)}{\leq} n^7 \exp \left(-n^{\gamma+o(1)} \right) \\ &= o(1). \end{aligned}$$

Above, (a) is a direct consequence of (31). Inequality (b) uses that $\gamma > \beta$, that there are n^7 elements of \mathbb{T}_n , and Lemma 21. \blacksquare

Appendix C. Concentration of the second derivative: Proof of Lemma 21

The proof of Lemma 21 relies on the following two results.

Lemma 23 *Let $t \geq 0$, $\beta > 0$ and assume that $\gamma \in ((1 + \beta)/2, 1)$. Then*

$$\mathbb{P} \left(\left| I \left[t, t + \frac{1}{n^\beta} \right] - \int_t^{t+n^{-\beta}} \lambda \text{cut}(\mathcal{I}(s)) ds \right| \geq n^{\gamma-\beta} \right) \leq \exp \left(-n^{2\gamma-(1+\beta)+o(1)} \right).$$

Lemma 24 *Let $\gamma, \beta > 0$ satisfy $\gamma > 1 - \beta$. Recall the event*

$$\mathcal{B}_t := \left\{ \forall v \in \text{highdeg}(G), T(v) \notin \left(t - \frac{1}{n^\beta}, t \right) \cup \left(t, t + \frac{1}{n^\beta} \right) \right\}.$$

Then the following hold:

$$\mathbb{P} \left(\left\{ \left| \int_t^{t+n^{-\beta}} \text{cut}(\mathcal{I}(s)) ds - \frac{\text{cut}(\mathcal{I}(t))}{n^\beta} \right| \geq n^{\gamma-\beta} \right\} \cap \mathcal{B}_t \right) \leq \exp \left(-n^{\gamma+o(1)} \right), \quad (32)$$

$$\mathbb{P} \left(\left\{ \left| \int_{t-n^{-\beta}}^t \text{cut}(\mathcal{I}(s)) ds - \frac{\text{cut}(\mathcal{I}(t^-))}{n^\beta} \right| \geq n^{\gamma-\beta} \right\} \cap \mathcal{B}_t \right) \leq \exp \left(-n^{\gamma+o(1)} \right). \quad (33)$$

Before proving the two lemmas, we show how they can be combined to prove the main result.

Proof [Proof of Lemma 21] It is immediate from Lemmas 23 and 24 and the triangle inequality that

$$\mathbb{P} \left(\left\{ \left| I \left[t, t + \frac{1}{n^\beta} \right] - \frac{\text{cut}(\mathcal{I}(t))}{n^\beta} \right| \geq 2n^{\gamma-\beta} \right\} \cap \mathcal{B}_t \right) \leq 2 \exp \left(-n^{2\gamma-(1+\beta)+o(1)} \right), \quad (34)$$

$$\mathbb{P} \left(\left\{ \left| I \left[t - \frac{1}{n^\beta}, t \right] - \frac{\text{cut}(\mathcal{I}(t^-))}{n^\beta} \right| \geq 2n^{\gamma-\beta} \right\} \cap \mathcal{B}_t \right) \leq 2 \exp \left(-n^{2\gamma-(1+\beta)+o(1)} \right). \quad (35)$$

Above, we have used the fact that $2\gamma - (1 + \beta) < \gamma$ whenever $\gamma < 1$ to simplify the probability bounds. The bounds (34) and (35) along with the definitions of $dI(t)$ and $\widehat{dI}_\beta(t)$ (see (27) and (26)) proves that

$$\mathbb{P} \left(\left\{ \left| \widehat{dI}_\beta(t) - dI(t) \right| \geq 4n^\gamma \right\} \cap \mathcal{B}_t \right) \leq 4 \exp \left(-n^{2\gamma-(1+\beta)+o(1)} \right),$$

which is equivalent to the claimed result. ■

The rest of this section is devoted to the proofs of the two lemmas. The proof of Lemma 23 is in Section C.1, and the proof of Lemma 24 is in Section C.2.

C.1. Approximating the cut: Proof of Lemma 23

The key tool used to prove Lemma 23 is a version of Freedman's inequality due to Shorack and Wellner; see (Shorack and Wellner, 1986, Appendix B).

Lemma 25 (Freedman's inequality for continuous-time martingales) *Suppose that $\{X(t)\}_{t \geq 0}$ is a real-valued continuous-time martingale adapted to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. Suppose further that X_t has jumps that are bounded by $C > 0$ in absolute value, almost surely. Define the quadratic variation of the process*

$$\langle X \rangle_t := \lim_{\xi \rightarrow 0} \sum_{k=0}^{\lfloor t/\xi \rfloor} \mathbb{E} \left[(X((k+1)\xi) - X(k\xi))^2 \middle| \mathcal{F}_{k\xi} \right].$$

Then for any $x \geq 0$ and $\sigma > 0$, it holds that

$$\mathbb{P}(X(t) - X(0) \geq x \text{ and } \langle X \rangle_t \leq \sigma^2 \text{ for some } t \geq 0) \leq \exp \left(-\frac{x^2/2}{\sigma^2 + Cx/3} \right).$$

Proof For $y \geq t$, define the process

$$M(y) := I(y) - I(t) - \int_t^y \lambda \text{cut}(\mathcal{I}(s)) ds.$$

Notice that $M(y)$ is a martingale. This is evident from the formula for the condition expectation of the derivative of I , given in (4). To apply the concentration inequality of Lemma 25, we derive a few basic properties of $M(y)$. The process jumps by 1 precisely when a new vertex is infected, so we may set $C = 1$. To bound the quadratic variation, we have, for $\xi > 0$ sufficiently small,

$$\begin{aligned} & \mathbb{E} \left[(M(y+\xi) - M(y))^2 \middle| \mathcal{F}_y \right] \\ & \leq 2 \mathbb{E} \left[(I(y+\xi) - I(y))^2 \middle| \mathcal{F}_y \right] + 2 \mathbb{E} \left[\left(\int_y^{y+\xi} \lambda \text{cut}(\mathcal{I}(s)) ds \right)^2 \middle| \mathcal{F}_y \right], \end{aligned} \quad (36)$$

where we have used the inequality $(a+b)^2 \leq 2a^2 + 2b^2$. To bound the first term on the right hand side of (36), we recall from Lemma 18 that $I(y+\xi) - I(y)$ is stochastically bounded by a $\text{Poi}(\lambda m n d \xi)$ random variable conditioned on \mathcal{F}_y . It follows that

$$\mathbb{E} \left[(I(y+\xi) - I(y))^2 \middle| \mathcal{F}_y \right] \leq \lambda m n d \xi + (\lambda m n d \xi)^2 = \lambda m n d \xi + o(\xi). \quad (37)$$

To bound the second term on the right hand side of (36), use the (somewhat loose) bound $\text{cut}(\mathcal{I}(s)) \leq |E| \leq n^2$ to arrive at

$$\int_y^{y+\xi} \lambda \text{cut}(\mathcal{I}(s)) ds \leq \xi \lambda n^2. \quad (38)$$

Together, (37) and (38) show that

$$\mathbb{E} \left[(M(y+\xi) - M(y))^2 \middle| \mathcal{F}_y \right] \leq 2\lambda m n d \xi + o(\xi),$$

hence $\langle M \rangle_y \leq 2\lambda m n d (y - t)$.

Next, define the stopped process $\overline{M}(y) := M(y \wedge (t + n^{-\beta}))$. For this process, we have that $M(t + n^{-\beta}) = \overline{M}(t + n^{-\beta})$ and that $\langle \overline{M} \rangle_y \leq 2\lambda m d n^{1-\beta}$ almost surely. Lemma 25 applied to the process \overline{M} shows that

$$\begin{aligned} \mathbb{P} \left(M(t + n^{-\beta}) - M(t) \geq x \middle| \mathcal{F}_t \right) &= \mathbb{P} \left(\overline{M}(t + n^{-\beta}) - \overline{M}(t) \geq x \middle| \mathcal{F}_t \right) \\ &\leq \exp \left(-\frac{x^2/2}{2\lambda m d n^{1-\beta} + x/3} \right). \end{aligned}$$

Applying the same analysis to the martingale $-M$ and taking a union bound shows that, for $x \geq 0$,

$$\mathbb{P} \left(\left| I \left[t, t + \frac{1}{n^\beta} \right] - \int_t^{t+n^{-\beta}} \lambda \text{cut}(\mathcal{I}(s)) ds \right| \geq x \mid \mathcal{F}_t \right) \leq 2 \exp \left(-\frac{x^2/2}{2\lambda m d n^{1-\beta} + x/3} \right).$$

Substituting $x = n^{\gamma-\beta}$ into the expression above yields a probability bound of

$$\exp \left(-n^{\min\{\gamma-\beta, 2\gamma-1-\beta\}+o(1)} \right).$$

The desired claim follows since $2\gamma - 1 - \beta < \gamma - \beta$ whenever $\gamma < 1$. \blacksquare

C.2. Changes in the cut: Proof of Lemma 24

As a shorthand, denote $\delta := n^{-\beta}$. We can bound

$$\begin{aligned} \left| \int_t^{t+\delta} \text{cut}(\mathcal{I}(s)) ds - \delta \text{cut}(\mathcal{I}(t)) \right| &\leq \int_t^{t+\delta} |\text{cut}(\mathcal{I}(s)) - \text{cut}(\mathcal{I}(t))| ds \\ &\leq \delta \sup_{t < s < t+\delta} |\text{cut}(\mathcal{I}(s)) - \text{cut}(\mathcal{I}(t))|. \end{aligned} \quad (39)$$

Recall the event $\mathcal{B}_t := \{\forall v \in \text{highdeg}(G), T(v) \notin (t - \delta, t) \cup (t, t + \delta)\}$. On \mathcal{B}_t , all vertices in $\mathcal{I}(t + \delta) \setminus \mathcal{I}(t)$ have degree at most d . As a result, for all $s \in (t, t + \delta)$, it holds that

$$|\text{cut}(\mathcal{I}(s)) - \text{cut}(\mathcal{I}(t))| \leq \sum_{v \in \mathcal{I}(s) \setminus \mathcal{I}(t)} \deg(v) \leq d |\mathcal{I}(s) \setminus \mathcal{I}(t)| \leq d |\mathcal{I}(t + \delta) \setminus \mathcal{I}(t)|. \quad (40)$$

The first inequality in (40) follows since the size of the cut changes by at most $\deg(v)$ vertices when v is added to the cascade. The second inequality holds since $\mathcal{I}(t) \subseteq \mathcal{I}(s) \subseteq \mathcal{I}(t + \delta)$.

We proceed by studying $|\mathcal{I}(t + \delta) \setminus \mathcal{I}(t)|$. Let $X \sim \text{Poi}(\lambda m n d \delta)$. By Lemma 18, it holds that

$$\begin{aligned} \mathbb{P} \left(|\mathcal{I}(t + \delta) \setminus \mathcal{I}(t)| \geq \frac{n^\gamma}{d} \right) &\stackrel{(a)}{\leq} \mathbb{P} \left(X \geq \frac{n^\gamma}{2d} + \mathbb{E}[X] \right) \\ &\stackrel{(b)}{\leq} \exp \left(-\frac{n^{2\gamma}/(8d^2)}{n^{1-\beta+o(1)} + n^\gamma/(6d)} \right) \\ &\stackrel{(c)}{\leq} \exp \left(-n^{\gamma+o(1)} \right). \end{aligned} \quad (41)$$

Above, (a) holds since $\mathbb{E}[X] \leq \lambda m n^{1-\beta+o(1)} \leq n^\gamma/(2d)$ provided $\gamma > 1 - \beta$. The inequality (b) follows from Bernstein's inequality and the fact that $\lambda m n d = n^{1+o(1)}$. The final inequality (c) holds by our assumption that $\gamma > 1 - \beta$. Together, (39), (40) and (41) imply the result (32).

To prove (32), we may follow identical reasoning. Analogously to (39) and (40), it holds on the event \mathcal{B}_t that

$$\left| \int_{t-\delta}^t \text{cut}(\mathcal{I}(s)) ds - \delta \text{cut}(\mathcal{I}(t^-)) \right| \leq \delta \sup_{t-\delta < s < t} |\text{cut}(\mathcal{I}(s)) - \text{cut}(\mathcal{I}(t))| \leq d |\mathcal{I}(t) \setminus \mathcal{I}(t - \delta)|.$$

We may then analyze $|\mathcal{I}(t) \setminus \mathcal{I}(t - \delta)|$ in the same manner as (41). Putting everything together yields the second result (33).

Appendix D. Lower bounds for $\alpha \in (0, 1/2)$: Proof of Theorem 4

To prove the theorem, we will construct a tailored family of graphs $\mathcal{H} \subset \mathcal{G}$ for which we can tractably analyze the distribution of infection times, while also constituting a “hard” instance for estimating high-degree vertices. To construct \mathcal{H} , we first describe a particular deterministic graph H .

Definition 26 Fix a positive integer N . Let H' be a connected tree with N vertices, maximum degree at most 3, and diameter at most $2 \log N$ (e.g., a balanced binary tree). For each vertex in $u \in H'$, we add a path of length $\log^5 n$ starting at u ; all added paths are disjoint. We call the resulting graph H . The set of endpoints of the paths added to H' , which are also the set of leaves of H , is denoted by L .

The addition of paths which are much longer than the diameter of H' is a crucial aspect of our construction. At a high level, it ensures that the infection times of vertices in L are primarily determined by the infection times of vertices in the corresponding path, rather than the infection times of vertices in H' . Since all the added paths are disjoint, the infection times of vertices in L become *almost* independent, which simplifies much of our analysis. A bit more formally, we have the following useful observation.

Fact 27 The infection times of vertices in L , given by the collection $\{\mathbf{T}(v)\}_{v \in L}$, are conditionally independent with respect to $\{\mathbf{T}(z)\}_{z \in H'}$, the collection of infection times of vertices in H' .

To construct the graphs of interest, we will introduce additional vertices to H which connect to L in a particular manner.

Definition 28 Suppose that N vertices, labelled from 1 to N , are added to H . Each vertex in $[N]$ forms an edge with a single element of L . Define \mathcal{H}_\emptyset to be the set of resulting graphs for which all vertices in L have degree at most $\log^2 n$. For each $v \in L$, define \mathcal{H}_v to be the set of graphs for which each $u \in L \setminus \{v\}$ has degree at most $\log^2 n$, but v has degree at least $D/2$. The union of all such graphs is given by $\mathcal{H} := \bigcup_{x \in L \cup \{\emptyset\}} \mathcal{H}_x$.

We make a few simple observations about the family \mathcal{H} .

- If $G \in \mathcal{H}_\emptyset$ then $\text{highdeg}(G) = \emptyset$, and if $G \in \mathcal{H}_v$ for $v \in L$ then $\text{highdeg}(G) = \{v\}$.
- The total number of vertices in $G \in \mathcal{H}$ is $n := 2N + N \log^5 N$. In particular, $N = n^{1-o(1)}$.

The first step in our analysis is to show that it is challenging to even detect whether the set of observed infection times are due to a graph in \mathcal{H}_\emptyset or in \mathcal{H}_v for some $v \in L$. To do so, we construct convenient distributions μ_\emptyset over elements of \mathcal{H}_\emptyset and μ_v over elements of \mathcal{H}_v , $v \in L$. We let P_\emptyset and P_v , $v \in L$ denote the induced distributions over the infection times. That is, for $x \in L \cup \{\emptyset\}$, P_x is the distribution corresponding to $\mathcal{SI}(G, v_0)^{\otimes K}$ where $G \sim \mu_x$ and $v_0 \perp H'$ is chosen arbitrarily (but is fixed across all cascades). We defer the precise description of the distributions $\{\mu_x\}_{x \in L \cup \{\emptyset\}}$ to Section D.1, and move forward with the proof for now.

Lemma 29 *Let $\epsilon > 0$ and $\alpha \in (0, 1/2)$. Suppose that $D \leq n^\alpha$ for some $\alpha \in (0, 1/2)$ and that*

$$K \leq \left(\frac{1 - 2\alpha - \epsilon}{5} \right) \frac{\log n}{\log \log n}.$$

Then it holds for all $v \in L$ that $\text{TV}(P_\emptyset, P_v) = o(1)$.

At a high level, we prove Lemma 29 by studying the χ^2 -divergence between P_\emptyset and P_v for $v \in L$, which provides an upper bound for the total variation distance and has a useful tensorization property which we leverage in our analysis. Before proving the lemma, we show how it can be used to prove the theorem.

Proof [Proof of Theorem 4] Define the measure μ over the ensemble \mathcal{H} via $\mu := \frac{1}{N+1} \sum_{x \in L \cup \{\emptyset\}} \mu_x$, and similarly define the induced measure over vertex infection times $P := \frac{1}{N+1} \sum_{x \in L \cup \{\emptyset\}} P_x$. Let us also define $\delta := \max_{x, y \in L \cup \{\emptyset\}} \text{TV}(P_x, P_y)$ and note that $\delta = o(1)$ by Lemma 29 and the triangle inequality. For any estimator HD, it holds that

$$\begin{aligned} P(\text{HD} = \text{highdeg}(G)) &= \frac{1}{N+1} \sum_{x \in L \cup \{\emptyset\}} P_x(\text{HD} = \{x\}) \\ &\stackrel{(a)}{=} 1 - \frac{1}{N+1} \sum_{x, y \in L \cup \{\emptyset\}: x \neq y} P_x(\text{HD} = \{y\}) \\ &\stackrel{(b)}{\leq} 1 - \frac{1}{N+1} \sum_{x, y \in L \cup \{\emptyset\}: x \neq y} (P_y(\text{HD} = \{y\}) - \delta) \\ &\stackrel{(c)}{=} 1 + N\delta - NP(\text{HD} = \text{highdeg}(G)), \end{aligned}$$

where (a) is due to the law of total probability, (b) uses the definition of the total variation distance, and (c) follows as each term $P_y(\text{HD} = \{y\})$ is counted N times in the summation. Rearranging terms shows that

$$P(\text{HD} = \text{highdeg}(G)) \leq \frac{1 + N\delta}{1 + N} = o(1). \quad \blacksquare$$

The remainder of this section is devoted to the proof of Lemma 29.

D.1. Distributions of graphs in \mathcal{H}

Before proving the lemma, we must describe the distributions $\{\mu_x\}_{x \in L \cup \{\emptyset\}}$ over graphs in \mathcal{H} . To this end, it is useful to first define corresponding ‘‘auxiliary’’ distributions $\{\mu'_x\}_{x \in L \cup \{\emptyset\}}$. We say that $G \sim \mu'_\emptyset$ if each vertex in $[N]$ chooses a uniform random element of L with which they form an edge. On the other hand, we say that $G \sim \mu'_v$ if, for each $u \in [N]$, u connects to a uniform random element of L with probability $1 - D/N$, otherwise u connects to v with probability D/N . Finally, for $x \in L \cup \{\emptyset\}$, the measure μ_x is equal to μ'_x conditioned on the sampled graph being an element of \mathcal{H}_x .

For $x \in L \cup \{\emptyset\}$, we recall that P_x is the distribution corresponding to $\mathcal{SI}(G, v_0)^{\otimes K}$, where $G \sim \mu_x$. Similarly, we define P'_x to be the distribution corresponding to $\mathcal{SI}(G, v_0)^{\otimes K}$, where $G \sim \mu'_x$. The following result shows that the total variation distance between P_x and P'_x is small.

Lemma 30 *It holds for any $x \in L \cup \{\emptyset\}$ that $\text{TV}(P_x, P'_x) = o(1)$.*

The lemma implies that we can essentially replace P_x with P'_x in our analysis moving forward. Indeed, by the triangle inequality for the total variation distance, we have for any $v \in L$ that

$$\text{TV}(P_\emptyset, P_v) \leq \text{TV}(P'_\emptyset, P'_v) + \text{TV}(P'_\emptyset, P_\emptyset) + \text{TV}(P'_v, P_v) = \text{TV}(P'_\emptyset, P'_v) + o(1), \quad (42)$$

where the final equality is due to Lemma 30. Hence it suffices to bound the total variation between the auxiliary distributions.

Proof [Proof of Lemma 30] Let $x \in L \cup \{\emptyset\}$. By the data processing inequality for the total variation distance, it holds that $\text{TV}(P_x, P'_x) \leq \text{TV}(\mu_x, \mu'_x)$. By the definition of μ_x and μ'_x , the two distributions can be coupled precisely when $G \sim \mu'_x$ satisfies $G \in \mathcal{H}_x$. Hence

$$\text{TV}(P_x, P'_x) \leq \text{TV}(\mu_x, \mu'_x) = \mu'_x(\mathcal{H}_x^c). \quad (43)$$

In light of (43), we proceed by proving that $\mu'_x(\mathcal{H}_x^c) = o(1)$. Notice that for $G \sim \mu'_x$, all vertices except for those in L have a deterministic degree which is at most 3, hence it suffices to study the degrees of vertices in L alone. In the case $x = \emptyset$, it holds for every $v \in L$ that $\deg_G(v) - 1 \sim \text{Bin}(N, 1/N)$. Bernstein's inequality implies

$$\mu'_\emptyset(\deg_G(v) \geq \log^2 N) \leq \mu'_\emptyset\left(\deg_G(v) - 1 \geq 1 + \frac{1}{2} \log^2 N\right) \leq \exp\left(-\frac{2}{3} \log^2 N\right). \quad (44)$$

It follows that

$$\begin{aligned} \mu'_\emptyset(\mathcal{H}_\emptyset^c) &= \mu'_\emptyset(\exists v \in L : \deg_G(v) > \log^2 n) \\ &\stackrel{(a)}{\leq} \mu'_\emptyset(\exists v \in L : \deg_G(v) > \log^2 N) \\ &\stackrel{(b)}{\leq} N \exp\left(-\frac{2}{3} \log^2 N\right) = o(1). \end{aligned} \quad (45)$$

Above, (a) follows since $N \leq n$, and (b) is due to a union bound over the N elements of L .

We now turn to the case where $x \in L$. If $v \in L \setminus \{x\}$, then $\deg_G(v) - 1 \sim \text{Bin}(N, (1 - \frac{D}{N}) \frac{1}{N})$, which is stochastically bounded by $\text{Bin}(N, 1/N)$. Following the same steps as (44), we have that

$$\mu'_x(\deg_G(v) \geq \log^2 n) \leq \mu'_x(\deg_G(v) \geq \log^2 N) \leq \exp\left(-\frac{2}{3} \log^2 N\right).$$

On the other hand, when $v = x$, $\deg_G(v)$ stochastically dominates $\text{Bin}(N, D/N)$. By Bernstein's inequality,

$$\mu'_x\left(\deg_G(v) \leq \frac{D}{2}\right) \leq \exp\left(-\frac{3}{56} D\right).$$

It follows that

$$\begin{aligned} \mu'_x(\mathcal{H}_x^c) &= \mu'_x\left(\exists v \in L \setminus \{x\} : \deg_G(v) > \log^2 n \text{ or } \deg_G(x) < \frac{D}{2}\right) \\ &\leq \sum_{v \in L \setminus \{x\}} \mu'_x(\deg_G(v) > \log^2 n) + \mu'_x\left(\deg_G(x) < \frac{D}{2}\right) \\ &\leq N \exp\left(-\frac{2}{3} \log^2 N\right) + \exp\left(-\frac{3}{56} D\right) = o(1). \end{aligned} \quad (46)$$

Together, (45), (46) and (43) prove the claim. \blacksquare

D.2. Impossibility of detection: Proof of Lemma 29

The advantage of using the auxiliary distributions is that they are product distributions conditioned on the infection times in H , which makes it easier to analyze the infection times in $[N]$. To see this concretely, let us assume without loss of generality that $\lambda = 1$, and define

$$\mathbf{E}(j) = (E_1(j), \dots, E_K(j)) \stackrel{i.i.d.}{\sim} \text{Exp}(1)^{\otimes K}, \quad j \in [N].$$

Let us also define the collection of independent variables $\{U(j)\}_{j \in [N]}$ where $U(j) \sim \text{Unif}(L)$ denotes the random neighbor of the vertex j in the construction of $G \sim P'_\emptyset$. We then have the representation

$$\mathbf{T}(j) = \mathbf{T}(U(j)) + \mathbf{E}(j), \quad j \in [N]. \quad (47)$$

On the other hand, if $G \sim P'_v$ for some $v \in L$, then we have the representation

$$\mathbf{T}(j) = \begin{cases} \mathbf{T}(U(j)) + \mathbf{E}(j) & \text{with probability } 1 - D/N \\ \mathbf{T}(v) + \mathbf{E}(j) & \text{with probability } D/N. \end{cases} \quad (48)$$

We can also succinctly characterize the distributions of infection times in $[N]$ for the cases described in (47) and (48). Define the following (conditional) distributions over $\mathbf{t} := (t_1, \dots, t_K) \in \mathbb{R}_+^K$,

$$f_v(\mathbf{t}) := \prod_{i=1}^K e^{-(t_i - T_i(v))} \mathbf{1}(t_i \geq T_i(v)) = e^{-\sum_{i=1}^K t_i} \prod_{i=1}^K e^{T_i(v)} \mathbf{1}(t_i \geq T_i(v))$$

$$f_\emptyset(\mathbf{t}) := \frac{1}{N} \sum_{w \in L} f_w(\mathbf{t}) = \frac{e^{-\sum_{i=1}^K t_i}}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w)} \mathbf{1}(t_i \geq T_i(w)).$$

Then it is readily seen that

$$\text{Law}(\{\mathbf{T}(j)\}_{j \in [N]} | \{\mathbf{T}(u)\}_{u \in L}) = \begin{cases} f_\emptyset^{\otimes N} & G \sim P'_\emptyset \\ \left[\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \right]^{\otimes N} & G \sim P'_v, v \in L. \end{cases} \quad (49)$$

Our strategy moving forward is to determine when the distributions in (49) are indistinguishable, which amounts to bounding the total variation distance between them. To do so, we will instead bound a larger distance measure – the χ^2 divergence. We recall that the χ^2 divergence between probability measures P and Q is defined as

$$\chi^2(P \| Q) := \int \frac{(dP - dQ)^2}{dQ}.$$

The total variation distance and the χ^2 divergence can be related through the inequality (Polyanskiy and Wu, 2024+, Proposition 7.15)

$$\text{TV}(P, Q) \leq 2\sqrt{\chi^2(P \| Q)}.$$

For our setting, the χ^2 divergence has two advantageous properties. First, it tensorizes; that is, for any positive integer k , it holds that (Polyanskiy and Wu, 2024+, Chapter 7.12)

$$\chi^2(P^{\otimes k} \| Q^{\otimes k}) = (1 + \chi^2(P \| Q))^k - 1.$$

Additionally, the formula for the χ^2 divergence involves the difference between the densities of P and Q , which makes it convenient to analyze mixture distributions, as is the case for our setting.

We apply these ideas to the conditional distributions in (49). Conditioned on $\{\mathbf{T}(u)\}_{u \in H}$, we have that

$$\begin{aligned} \text{TV} \left(f_\emptyset^{\otimes N}, \left[\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \right]^{\otimes N} \right) &\leq 2 \sqrt{\chi^2 \left(\left[\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \right]^{\otimes N} \parallel f_\emptyset^{\otimes N} \right)} \\ &= 2 \sqrt{\left[1 + \chi^2 \left(\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \parallel f_\emptyset \right) \right]^N - 1} \\ &\leq 2 \sqrt{\exp \left\{ N \chi^2 \left(\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \parallel f_\emptyset \right) \right\} - 1}. \end{aligned}$$

In particular, we see that the conditional total variation tends to 0 if

$$\chi^2 \left(\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \parallel f_\emptyset \right) = o\left(\frac{1}{N}\right). \quad (50)$$

We proceed by bounding the χ^2 divergence between the two distributions. Letting $\mathbf{t} = (t_1, \dots, t_K) \in \mathbb{R}_+^K$, it holds that

$$\begin{aligned} \chi^2 \left(\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \parallel f_\emptyset \right) &= \left(\frac{D}{N}\right)^2 \int_{\mathbb{R}_+^K} \frac{(f_v(\mathbf{t}) - f_\emptyset(\mathbf{t}))^2}{f_\emptyset(\mathbf{t})} d\mathbf{t} \\ &\leq \left(\frac{D}{N}\right)^2 \int_{\mathbb{R}_+^K} \left(\frac{f_v(\mathbf{t})}{f_\emptyset(\mathbf{t})}\right) f_v(\mathbf{t}) d\mathbf{t}. \end{aligned} \quad (51)$$

To bound the χ^2 divergence, we therefore need to study the likelihood ratio f_v/f_\emptyset . This is accomplished in the following result.

Lemma 31 *For every $v \in L$, it holds with probability $1 - o(1)$ that*

$$\sup_{\mathbf{t} \in \mathbb{R}_+^K} \frac{f_v(\mathbf{t})}{f_\emptyset(\mathbf{t})} \leq \log^{5K} n.$$

Before proving Lemma 31, we show how it can be used to prove the main result.

Proof [Proof of Lemma 29] Let \mathcal{A}_v denote the event described in Lemma 31; note that \mathcal{A}_v is measurable with respect to $\{\mathbf{T}(u)\}_{u \in L}$. On this event, it follows from (51) that

$$\chi^2 \left(\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \parallel f_\emptyset \right) \leq \left(\frac{D}{N}\right)^2 \int_{\mathbb{R}_+^K} (\log^{5K} n) f_v(\mathbf{t}) d\mathbf{t} \leq \left(\frac{D}{N}\right)^2 \log^{5K} n.$$

Since $D/N = n^{\alpha-1+o(1)}$, for any $\epsilon > 0$ the condition (50) is satisfied if

$$K \leq \frac{1 - 2\alpha - \epsilon}{5} \left(\frac{\log n}{\log \log n} \right). \quad (52)$$

We now translate our results on the conditional distributions $f_\emptyset^{\otimes N}$ and $[(1-D/N)f_\emptyset + (D/N)f_v]^{\otimes N}$ to their *unconditional* counterparts P'_\emptyset and P'_v . If (52) holds, then

$$\begin{aligned} \text{TV}(P'_\emptyset, P'_v) &\stackrel{(a)}{\leq} \mathbb{E} \left[\text{TV} \left(f_\emptyset^{\otimes N}, \left[\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \right]^{\otimes N} \right) \right] \\ &\stackrel{(b)}{\leq} \mathbb{E} \left[\text{TV} \left(f_\emptyset^{\otimes N}, \left[\left(1 - \frac{D}{N}\right) f_\emptyset + \frac{D}{N} f_v \right]^{\otimes N} \right) \mathbf{1}(\mathcal{A}_v) \right] + \mathbb{P}(\mathcal{A}_v^c) = o(1). \end{aligned}$$

Above, the inequality (a) follows from the property that conditioning increases the total variation (Polyanskiy and Wu, 2024+, Theorem 7.5), and inequality (b) holds since the total variation is always upper bounded by 1. Finally, $\text{TV}(P'_\emptyset, P'_v) = o(1)$ implies that $\text{TV}(P_\emptyset, P_v) = o(1)$ as well in light of (42). \blacksquare

In the remainder of this subsection, we study the likelihood ratio

$$\frac{f_v(\mathbf{t})}{f_\emptyset(\mathbf{t})} = \frac{\prod_{i=1}^K e^{T_i(v)} \mathbf{1}(t_i \geq T_i(v))}{\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w)} \mathbf{1}(t_i \geq T_i(w))}.$$

When there exists $i \in [K]$ such that $t_i < T_i(v)$, it is clear from the formula that the likelihood ratio is zero. On the other hand, when $t_i \geq T_i(v)$ for all $i \in [K]$,

$$\begin{aligned} \frac{f_v(\mathbf{t})}{f_\emptyset(\mathbf{t})} &= \frac{\prod_{i=1}^K e^{T_i(v)}}{\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w)} \mathbf{1}(t_i \geq T_i(w))} \\ &\leq \frac{\prod_{i=1}^K e^{T_i(v)}}{\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w)} \mathbf{1}(T_i(v) \geq T_i(w))} \\ &= \left(\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \right)^{-1}. \end{aligned}$$

Above, the inequality on the second line holds since $\mathbf{1}(T_i(v) \geq T_i(w)) \leq \mathbf{1}(t_i \geq T_i(w))$ if $t_i \geq T_i(v)$. Importantly, the upper bound derived for the likelihood ratio does not depend on \mathbf{t} . Hence, to control $\sup_{\mathbf{t} \in \mathbb{R}_+^K} f_v(\mathbf{t})/f_\emptyset(\mathbf{t})$, we establish a lower bound for the quantity

$$\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)). \quad (53)$$

We remark that the analysis of (53) is greatly facilitated by the way in which the graph H is constructed (see Definition 26). To see this concretely, for each $v \in L$ let v' be the first vertex in H' on the unique path in H connecting v to H' . Then we can write

$$T_i(v) = T_i(v') + X_i(v), \quad i \in [K], \quad (54)$$

where $X_i(v)$ is the time it takes for the cascade to spread from v' to v . Importantly, as the paths added to H' to form H are all disjoint, $\{X_i(v)\}_{v \in L}$ is a collection of i.i.d. $\text{Gamma}(\log^5 N)$ random variables. Since the diameter of H' is assumed to be much smaller than $\log^5 N$, we can show that

$T_i(v') = o(\log^5 N)$, and the behavior of $T_i(v)$ is largely dictated by $X_i(v)$, a random variable that is independent across $v \in L$. As we shall see, this reduction to i.i.d. random variables leads to a tractable analysis of (53).

To formalize these ideas, we begin by defining the following ‘‘nice’’ event.

Definition 32 *Let $v \in L$. The event \mathcal{E}_v holds if and only if the following conditions are satisfied:*

1. $T_i(z) \leq \log^2 N$ for all $z \in H'$ and $i \in [K]$.
2. $T_i(v) \geq \frac{2}{3} \log^5 N$ for all $i \in [K]$.
3. $\sum_{i=1}^K (X_i(v) - \log^5 N)^2 \leq 2K \log^5 N$.

Importantly, as the following result shows, \mathcal{E}_v holds with high probability.

Lemma 33 *For all $v \in L$, it holds that $\mathbb{P}(\mathcal{E}_v) = 1 - o(1)$.*

Proof We let $\mathcal{E}_{v,1}$, $\mathcal{E}_{v,2}$ and $\mathcal{E}_{v,3}$ denote the events corresponding to the first, second and third conditions of \mathcal{E}_v in Definition 32. We start by showing that $\mathcal{E}_{v,1}$ holds with probability $1 - o(1)$. Since H' is a tree and the diameter of H' is at most $2 \log N$ (see Definition 26), we have that $T_i(z)$ is stochastically upper bounded by a Gamma($2 \log N, 1$) random variable. Hence

$$\mathbb{P}(T_i(z) \geq \log^2 N) \leq e^{-\frac{1}{2} \log^2 N} 2^{\log N} \leq e^{-\frac{1}{4} \log^2 N},$$

where the first inequality is due to a Chernoff bound, and the second inequality holds for N sufficiently large. Taking a union bound over all $z \in H'$ and $i \in [K]$, we obtain

$$\mathbb{P}(\mathcal{E}_{v,1}^c) = \mathbb{P}\left(\max_{z \in H', i \in [K]} T_i(z) \geq \log^2 N\right) \leq \sum_{z \in H', i \in [K]} \mathbb{P}(T_i(z) \geq \log^2 N) \leq NK e^{-\frac{1}{4} \log^2 N} = o(1). \quad (55)$$

We now turn to the events $\mathcal{E}_{v,2}$ and $\mathcal{E}_{v,3}$. Notice that if $\mathcal{E}_{v,1}$ holds and there exists $i \in [K]$ such that $T_i(v) < \frac{2}{3} \log^5 N$, then

$$\sum_{i=1}^K (X_i(v) - \log^5 N)^2 \geq \left(\frac{\log^5 N}{4}\right)^2 > 2K \log^5 N, \quad (56)$$

where the final inequality holds for N sufficiently large. As a consequence of (56), we have that

$$\mathcal{E}_{v,1} \cap \mathcal{E}_{v,2}^c \subseteq \mathcal{E}_{v,3}^c, \quad (57)$$

so we focus on bounding the probability of the latter event. We do so through a second moment argument. As a shorthand, let $R := \log^5 N$. The first moment of each of the summands is $\mathbb{E}[(X_i(v) - R)^2] = \text{Var}(X_i(v)) = R$, where we have used the fact that $X_i(v) \sim \text{Gamma}(R, 1)$. The second moment is

$$\begin{aligned} \mathbb{E}[(X_i(v) - R)^4] &= \mathbb{E}[X_i(v)^4] - 4R \mathbb{E}[X_i(v)^3] + 6R^2 \mathbb{E}[X_i(v)^2] - 4R^3 \mathbb{E}[X_i(v)] + R^4 \\ &= R(R+1)(R+2)(R+3) - 4R^2(R+1)(R+2) + 6R^3(R+1) - 4R^4 + R^4 \\ &= 3R^2 + 6R. \end{aligned}$$

In particular, $\text{Var}((X_i(v) - R)^2) = 2R^2 + 6R \leq 8R^2$, and furthermore,

$$\text{Var}\left(\sum_{i=1}^K (X_i(v) - R)^2\right) = \sum_{i=1}^K \text{Var}((X_i(v) - R)^2) \leq 8KR^2.$$

Finally, by Chebyshev's inequality,

$$\mathbb{P}(\mathcal{E}_{v,3}^c) = \mathbb{P}\left(\sum_{i=1}^K (X_i(v) - R)^2 > 2KR\right) \leq \frac{1}{K^2 R^2} \text{Var}\left(\sum_{i=1}^K (X_i(v) - R)^2\right) \leq \frac{8}{K} = o(1). \quad (58)$$

Together, (55), (57) and (58) prove that $\mathbb{P}(\mathcal{E}_v^c) = o(1)$ as claimed. \blacksquare

Next, we state a simple consequence of Taylor's theorem that will be useful for the proof of Lemma 31.

Lemma 34 *Suppose that a, b are real numbers such that $b \neq 0$ and $a/b \geq -1/2$. Then*

$$e^{-a} \left(1 + \frac{a}{b}\right)^b \geq e^{-2a^2/b}.$$

Proof Taylor's theorem implies that for any $z \geq -1/2$, it holds that $\log(1+z) \geq z - 2z^2$. Using this inequality, we have that

$$\log\left(e^{-a} \left(1 + \frac{a}{b}\right)^b\right) = -a + b \log\left(1 + \frac{a}{b}\right) \geq -a + b\left(\frac{a}{b} - 2\left(\frac{a}{b}\right)^2\right) = -\frac{2a^2}{b}. \quad \blacksquare$$

We now prove the main result of this section.

Proof [Proof of Lemma 31] Throughout this proof, we shall condition on the infection times in H' , given by the collection $\{\mathbf{T}(z)\}_{z \in H'}$, as well as v , given by $\mathbf{T}(v)$. We will also assume that the ‘‘nice’’ event \mathcal{E}_v holds, which is measurable with respect to $\{\mathbf{T}(z)\}_{z \in H'}$, $\mathbf{T}(v)$. Finally, as a shorthand, we will denote $R := \log^5 N$.

Recall from Fact 27 that $\{\mathbf{T}(u)\}_{u \in L}$ is a collection of conditionally independent variables with respect to $\{\mathbf{T}(z)\}_{z \in H'}$. As a result, if we condition on both $\{\mathbf{T}(z)\}_{z \in H'}$ and $\mathbf{T}(v)$, then the terms

$$\prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)), \quad w \in L, \quad (59)$$

are mutually (conditionally) independent. Moreover, each of the terms in (59) are bounded between 0 and 1. Hoeffding's inequality therefore implies that

$$\begin{aligned} & \frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \\ & \geq \mathbb{E}\left[\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| \{\mathbf{T}(z)\}_{z \in H'}, \mathbf{T}(v)\right] - N^{-1/3}, \quad (60) \end{aligned}$$

with probability at least $1 - 2 \exp(-2N^{1/3})$. We proceed by studying the conditional expectation in (60). To this end, for any $w \in L$, we denote w' to be the last vertex in H' on the unique path starting from v_0 and ending at w ; such a vertex exists since H is connected. We have that

$$\begin{aligned}
 & \mathbb{E} \left[\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| \{\mathbf{T}(z)\}_{z \in H'}, \mathbf{T}(v) \right] \\
 &= \frac{1}{N} \sum_{w \in L} \mathbb{E} \left[\prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| \{\mathbf{T}(z)\}_{z \in H'}, \mathbf{T}(v) \right] \\
 &= \frac{1}{N} \sum_{w \in L} \mathbb{E} \left[\prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| \mathbf{T}(w'), \mathbf{T}(v) \right] \\
 &= \frac{1}{N} \sum_{w \in L} \prod_{i=1}^K \mathbb{E} \left[e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| T_i(w'), T_i(v) \right]. \tag{61}
 \end{aligned}$$

Above, the second equality uses that the conditional expectation depends only on the random variables $T_i(w), T_i(v)$, and that the distribution of $T_i(w)$ conditioned on $\{\mathbf{T}(z)\}_{z \in H'}$ is the same as the distribution conditioned on $T(w')$ in light of (54). The final equality follows from the independence of the $T_i(w)$'s over $i \in [K]$.

We continue by lower bounding the conditional expectations in (61). Recall that for $w \in L$, we have from (54) that $T_i(w) = T_i(w') + X_i(w)$ where $X_i(w) \sim \text{Gamma}(R, 1)$. Then

$$\begin{aligned}
 & \mathbb{E} \left[e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| T_i(w'), T_i(v) \right] \\
 &= e^{T_i(w') - T_i(v)} \mathbb{E} \left[e^{X_i(w)} \mathbf{1}(T_i(v) - T_i(w') \geq X_i(w)) \middle| T_i(w'), T_i(v) \right] \\
 &\stackrel{(a)}{=} \frac{e^{T_i(w') - T_i(v)}}{\Gamma(R)} \int_0^{T_i(v) - T_i(w')} x^{R-1} dx \\
 &= \frac{e^{T_i(w') - T_i(v)} (T_i(v) - T_i(w'))^R}{\Gamma(R+1)} \\
 &\stackrel{(b)}{\geq} \frac{e^{T_i(w') - T_i(v)} (T_i(v) - T_i(w'))^R}{3\sqrt{R}(R/e)^R} \\
 &= \frac{1}{3\sqrt{R}} e^{-(T_i(v) - T_i(w') - R)} \left(1 + \frac{T_i(v) - T_i(w') - R}{R} \right)^R \\
 &\stackrel{(c)}{\geq} \frac{1}{3\sqrt{R}} \exp \left\{ -\frac{2(T_i(v) - T_i(w') - R)^2}{R} \right\} \\
 &\stackrel{(d)}{\geq} \frac{1}{3\sqrt{R}} \exp \left\{ -\frac{4}{R} (X_i(v) - R)^2 - \frac{4}{R} (T_i(v') + T_i(w'))^2 \right\} \\
 &\stackrel{(e)}{\geq} \frac{e^{-16}}{3\sqrt{R}} \exp \left\{ -\frac{4}{R} (X_i(v) - R)^2 \right\}. \tag{62}
 \end{aligned}$$

In the display above, the equality (a) uses that $T_i(v) - T_i(w') \geq 0$, which holds since $T_i(v) \geq 2R/3 \geq \sqrt{R} \geq \max_{z \in G'} T_i(z)$ on the event \mathcal{E}_v . The inequality (b) is due to Stirling's approximation. The inequality (c) follows from Lemma 34, which can be applied here since $T_i(v) - T_i(w') \geq$

$R/2$ on \mathcal{E}_v . The inequality (d) uses $(a + b)^2 \leq 2a^2 + 2b^2$, and the inequality (e) follows since $T_i(v'), T_i(w') \leq \sqrt{R}$ on the event \mathcal{E}_v .

Substituting the lower bound (62) into (61), we obtain

$$\begin{aligned} \mathbb{E} \left[\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \middle| \{\mathbf{T}(z)\}_{z \in H'}, \mathbf{T}(v) \right] \\ \geq \left(\frac{e^{-16}}{3\sqrt{R}} \right)^K \exp \left\{ -\frac{4}{R} \sum_{i=1}^K (X_i(v) - R)^2 \right\} \stackrel{(f)}{\geq} \left(\frac{e^{-30}}{\sqrt{R}} \right)^K, \end{aligned} \quad (63)$$

where (f) holds on the event \mathcal{E}_v .

Together, (60) and (63) imply that conditioned on $\{\mathbf{T}(z)\}_{z \in H'}, \mathbf{T}(v)$, and provided \mathcal{E}_v holds, we have with probability at least $1 - 2 \exp(2N^{-1/3})$ that

$$\frac{1}{N} \sum_{w \in L} \prod_{i=1}^K e^{T_i(w) - T_i(v)} \mathbf{1}(T_i(v) \geq T_i(w)) \geq \left(\frac{e^{-30}}{\sqrt{R}} \right)^K - N^{-1/3} \geq R^{-K}.$$

Finally, the desired result holds since $\mathbb{P}(\mathcal{E}_v) = 1 - o(1)$. ■

Appendix E. Proof of Proposition 9

It suffices to show that the model described in Proposition 9 satisfies the equation (1). As (1) describes a *unique* Markov process, the claim will then follow. We start by showing that it is unlikely for two vertices to become infected in a very small time interval.

Lemma 35 *For any $t \geq 0$, it holds that*

$$\mathbb{P}(|\mathcal{I}[t, t + \epsilon]| \geq 2 | \mathcal{F}_t) = o(\epsilon).$$

Proof Define τ_1 and τ_2 to be the first and second times after t when a new vertex is added to the cascade, respectively. Conditionally on \mathcal{F}_t , the memoryless property of Exponential distributions implies that $\tau_1 - t$ is equal in distribution to $\text{Exp}(\lambda |\text{cut}(\mathcal{I}[0, t])|)$, which stochastically dominates a $\text{Exp}(\lambda n^2)$ random variable (here, we have used that the number of edges in G is at most n^2). By the memoryless property of Exponential distributions, $\tau_2 - \tau_1$ is independent of $\tau_1 - t$. Through similar argument as above, it holds that $\tau_2 - \tau_1$ stochastically dominates a $\text{Exp}(\lambda n^2)$ random variable as well. Consequently, the probability that at least two vertices join the cascade in the time interval $[t, t + \epsilon]$ is

$$\begin{aligned} \mathbb{P}(\tau_2 - t \leq \epsilon | \mathcal{F}_t) &\leq \mathbb{P}(\tau_1 - t \leq \epsilon \text{ and } \tau_2 - \tau_1 \leq \epsilon | \mathcal{F}_t) \\ &= \mathbb{P}(\tau_1 - t \leq \epsilon | \mathcal{F}_t) \mathbb{P}(\tau_2 - \tau_1 \leq \epsilon | \mathcal{F}_t) \\ &\leq \left(1 - e^{-\epsilon \lambda n^2}\right)^2 \\ &\leq \epsilon^2 \lambda^2 n^4 = o(\epsilon). \end{aligned}$$

■

We now turn to the proof of the main result of this section.

Proof [Proof of Proposition 9] We start by defining some relevant notation. Condition on \mathcal{F}_t , and let $v \in V \setminus \mathcal{I}[0, t]$. Recalling that $\{F(e)\}_{e \in E}$ is a collection of i.i.d. $\text{Exp}(\lambda)$ random variables, we define

$$W := \min\{F(w, v) : w \in \mathcal{I}[0, t], (w, v) \in E\}.$$

By basic properties of Exponential distributions, we have that $W \sim \text{Exp}(\lambda|\mathcal{N}(v) \cap \mathcal{I}[0, t]|)$. Additionally define the event \mathcal{E} to hold if and only if at most one vertex joins the cascade in the interval $(t, t + \epsilon]$.

We claim that if \mathcal{E} holds, then $v \in \mathcal{I}[0, t + \epsilon] \iff W \leq \epsilon$. Indeed, if $W \leq \epsilon$, then from the description of $T(v)$ in the proposition statement, we have that $T(v) \leq t + \epsilon$, which implies that $v \in \mathcal{I}[t, t + \epsilon]$. On the other hand, if $W > \epsilon$, then we consider two cases: (1) no vertex becomes infected in $[t, t + \epsilon]$ or (2) one vertex becomes infected in $[t, t + \epsilon]$. Note that no other case is possible if \mathcal{E} holds. In the first case, it holds for all $e \in \text{cut}(\mathcal{I}[0, t])$ that $F(e) > \epsilon$, hence the weight of any path starting in $\mathcal{I}[0, t]$ and ending at v is at least ϵ , implying that $T(v) > t + \epsilon$. In the second case, if $v \in \mathcal{I}[t, t + \epsilon]$ despite $W > \epsilon$, there must exist a path $P \in \mathcal{P}(v, t)$ of length at least 2 such that $\text{weight}(P, t) \leq \epsilon$. However, this implies the existence of another vertex u on the path which is *not* an element of $\mathcal{I}[0, t]$. Since the portion of the path halting at u has a smaller weight than P , it follows that $T(u) \in \mathcal{I}[t, t + \epsilon]$. However, it is not possible for both u and v to be elements of $\mathcal{I}[t, t + \epsilon]$ if \mathcal{E} holds, hence we must have that $T(v) > t + \epsilon$.

As a result of the argument above, we have that

$$\begin{aligned} \mathbb{P}(\{v \in \mathcal{I}[0, t + \epsilon]\} \cap \mathcal{E} | \mathcal{F}_t) &= \mathbb{P}(\{W \leq \epsilon\} \cap \mathcal{E} | \mathcal{F}_t) \\ &= \mathbb{P}(W \leq \epsilon | \mathcal{F}_t) - \mathbb{P}(\{W \leq \epsilon\} \cap \mathcal{E}^c | \mathcal{F}_t) \\ &= \epsilon \lambda |\mathcal{N}(v) \cap \mathcal{I}[0, t]| + o(\epsilon). \end{aligned}$$

Finally, since $\mathbb{P}(\mathcal{E} | \mathcal{F}_t) = 1 - o(\epsilon)$ by Lemma 35, it holds that

$$\mathbb{P}(v \in \mathcal{I}[0, t + \epsilon] | \mathcal{F}_t) = \epsilon \lambda |\mathcal{N}(v) \cap \mathcal{I}[0, t]| + o(\epsilon),$$

hence (1) is satisfied. ■