



A semiparametric Bayesian approach to epidemics, with application to the spread of the coronavirus MERS in South Korea in 2015

Michael Schweinberger  ^a, Rashmi P. Bomiriya  ^b and Sergii Babkin  ^c

^aDepartment of Statistics, University of Missouri, Columbia, MO, USA; ^bRS Metrics Asia, Penn State University, University Park, PA, USA; ^cMicrosoft, Rice University, Houston, TX, USA

ABSTRACT

We consider incomplete observations of stochastic processes governing the spread of infectious diseases through finite populations by way of contact. We propose a flexible semiparametric modelling framework with at least three advantages. First, it enables researchers to study the structure of a population contact network and its impact on the spread of infectious diseases. Second, it can accommodate short- and long-tailed degree distributions and detect potential superspreaders, who represent an important public health concern. Third, it addresses the important issue of incomplete data. Starting from first principles, we show when the incomplete-data generating process is ignorable for the purpose of Bayesian inference for the parameters of the population model. We demonstrate the semiparametric modelling framework by simulations and an application to the partially observed MERS epidemic in South Korea in 2015. We conclude with an extended discussion of open questions and directions for future research.

ARTICLE HISTORY

Received 29 June 2020
Accepted 19 August 2021

KEYWORDS

Contact networks; network sampling; link-tracing; missing data

MSC SUBJECT CLASSIFICATIONS

05C80; 05C81

1. Motivation

The ongoing spread of COVID-19 and other viruses underscores the importance of understanding how infectious diseases spread by way of contact and how the spread of infectious diseases can be curbed. Indeed, COVID-19 is not the first virus to spread around the globe, and it will not be the last: Epidemics have been documented since at least the Middle Ages (e.g. the plague) and are primed to become more frequent rather than less frequent in the interconnected world of the twenty-first century (as the recent spread of COVID-19, MERS, SARS, Ebola, and Zika demonstrates).

CONTACT Michael Schweinberger  msq5c@missouri.edu  Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211, USA

 Supplemental data for this article can be accessed here. <https://doi.org/10.1080/10485252.2021.1972294>

1.1. *Advantages of network-based approaches to epidemics*

We follow a network-based approach to modelling the spread of infectious diseases. A network-based approach is appealing, because the network of contacts in a population determines how infectious diseases can spread (e.g. Keeling and Eames 2005; Danon et al. 2011; Welch, Bansal, and Hunter 2011). One of the advantages of a network-based approach is that heterogeneity in the number of contacts can be captured (e.g. Danon et al. 2011) along with other features of population contact networks (e.g. Welch et al. 2011). Indeed, conventional models of epidemics – including classic and lattice-based Susceptible-Infectious-Recovered (SIR) and Susceptible-Exposed-Infectious-Recovered (SEIR) models (e.g. Andersson and Britton 2000; Danon et al. 2011) – can be considered to be degenerate versions of network-based models, in the sense that such models postulate that with probability 1 the population contact network is of a known form: e.g. with probability 1 each population member is in contact with every other population member. Worse, the postulated form of the population contact network may not resemble real-world contact networks. A second advantage is that a network-based approach helps study the structure of a population contact network and its generating mechanism, helping generalise findings to similar populations.

1.2. *Shortcomings of existing network-based approaches*

Motivated by the shortcomings of classic and lattice-based SIR and SEIR models of epidemics, Britton and O'Neill (2002), Groendyke, Welch, and Hunter (2011, 2012), Groendyke and Welch (2018), Bu, Aiello, Xu, and Volfovsky (2021), and others have explored a network-based approach to epidemics. While a network-based approach is more appealing than classic and lattice-based SIR and SEIR models, existing network-based models of epidemics have shortcomings. As we discuss in Section 3, one of the more important shortcomings is that existing network-based approaches are either not flexible models of degree distributions or induce short-tailed degree distributions, that is, the population does not contain population members who have many more contacts than the bulk of the population members. Short-tailed degree distributions are problematic, because degree distributions of real-world contact networks are thought to be long-tailed (e.g. Laumann, Gagnon, Michael, and Michaels 1994; Albert and Barabàsi 2002; Jones and Handcock 2003a, 2003b, 2004) and the population members in the upper tail of the degree distribution represent an important public health concern: Population members with many contacts can infect many others and are hence potential superspreaders. Indeed, there is circumstantial evidence to suggest that superspreaders have played a role in the SARS epidemic in 2002–2003, the MERS epidemic in 2015, and the ongoing COVID-19 pandemic.

1.3. *Proposed network-based approach*

We introduce a flexible semiparametric modelling framework based on infinite mixture models and Dirichlet process priors (Ferguson 1973; Teh 2010), with at least three advantages. First, it shares with existing network-based approaches the advantage that

it enables researchers to study the structure of a population contact network and its impact on the spread of infectious diseases. Second, in contrast to existing network-based approaches, it is a flexible model of short- and long-tailed degree distributions and can detect potential superspreaders. Third, it addresses the important issue of incomplete data and can deal with a wide range of missing data and sample data. Starting from first principles, we show when the incomplete-data generating process is ignorable for the purpose of Bayesian inference for the parameters of the population model. In addition, we stress the importance of collecting network data with a view to reducing the posterior uncertainty about the population contact network and its generating mechanism along with possible sources of infections. We demonstrate the proposed framework by simulations and an application to the partially observed MERS epidemic in South Korea in 2015 (Ki 2015). The MERS epidemic was driven by the coronavirus MERS, which is related to the coronaviruses SARS and COVID-19. We detect three to five potential superspreaders, who may have had a great impact on the outcome of the outbreak.

1.4. Goal: superpopulation inference for finite populations

The proposed semiparametric modelling framework, based on infinite mixture distributions and Dirichlet process priors (Ferguson 1973; Teh 2010), extends to infinite populations. That said, we assume that the number of population members N is finite and embrace a superpopulation approach to statistical inference along the lines of Hartley and Sielken (1975) and Schweinberger, Krivitsky, Butts, and Stewart (2020), motivated by applications.

The assumption of finite N is motivated by the fact that in epidemiological applications the number of population members N cannot be infinite. For example, when the population of interest consists of all animals or all humans on earth, the size of the population is bounded above by real-world constraints such as geography and the scarcity of natural resources: Planet earth cannot host infinite populations of animals or humans.

Since the population of interest is finite, the natural objective of statistical inference is to learn the stochastic process that generated the population contact network and allows infectious diseases to spread through the population of interest, with a view to understanding and predicting epidemics in the population of interest and similar populations. In other words, it is natural to embrace a superpopulation approach to statistical inference, as discussed by Hartley and Sielken (1975) and Schweinberger et al. (2020). The properties of statistical procedures for superpopulation inference can be understood by developing a non-asymptotic statistical theory that relies on concentration inequalities and other non-asymptotic tools that have been embraced in high-dimensional statistics (see, e.g. Wainwright 2019). We are not aware of non-asymptotic statistical theory for stochastic models of epidemics, although there are asymptotic results in probability theory (e.g. Reinert 1995; Britton, Lindholm, and Turova 2011; Barbour and Reinert 2013; Pang and Pardoux 2020; Ball 2021) and statistical theory (e.g. Britton 1998, 2001) based on $N \rightarrow \infty$ asymptotics. Developing non-asymptotic statistical theory for stochastic models of epidemics constitutes an interesting direction for future research, but is beyond the scope of our paper.

1.5. Structure of the paper

The remainder of the paper is structured as follows. A network-based stochastic model of epidemics is reviewed in Section 2. We discuss shortcomings of parametric population models in Section 3 and introduce semiparametric population models in Section 4. In Section 5, we argue that collecting complete data is all but impossible, and stress the importance of collecting network data. Principled Bayesian inference based on incomplete data is discussed in Section 6. We present simulation results in Section 7 and an application to the partially observed MERS epidemic in South Korea in Section 8. We conclude with an extended discussion of open questions and directions for future research in Section 9.

2. A network-based stochastic model of epidemics

We introduce a network-based stochastic model of epidemics: We first describe a generic data-generating process in Section 2.1 and then review parametric population models in Section 2.2.

2.1. Data-generating process

We consider a population with $N < \infty$ population members, who may be connected by contacts. In the simplest case, contacts among population members are time-invariant and are either absent or present. We discuss in Section 9.6 possible extensions to time-evolving population contact networks.

A basic data-generating process is shown in Figure 1 and can be described as follows:

- Generate a population contact network.
- Conditional on the population contact network, generate an epidemic.

The population contact network is generated by a random graph model as described in Section 2.2 (parametric) or Section 4 (semiparametric). Conditional on the population contact network, an infectious disease spreads through the population by way of contact, governed by a continuous-time stochastic process such as the SIR and SEIR model (Andersson and Britton 2000; Britton and O'Neill 2002; Groendyke et al. 2011, 2012). We focus on the network-based SEIR model, which can be sketched as follows (Britton and O'Neill 2002; Groendyke et al. 2011, 2012). In the simplest case, the initial state of the

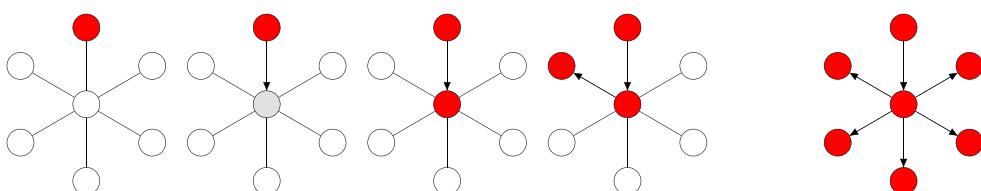


Figure 1. Data-generating process: Conditional on contacts (undirected lines) among population members (circles), infectious population members (red) spread an infectious disease by contact (directed lines) to susceptible population members (white), which are exposed (gray) before turning infectious (red).

stochastic process consists of a population with one infected population member and $N-1$ susceptible population members. Infected population members pass through three states: the exposed state; the infectious state; and the removed state. In the exposed state, population members have been exposed to the infectious disease, but cannot infect others. In the infectious state, population members can infect susceptible population members by way of contact, with transmissions being independent across contacts. In the final state – the removed state – population members have either recovered and are immune to re-infection or have died, and hence cannot infect others. The epidemic continues until all infected population members have been removed from the population. All of the described events – the event that an infectious population member infects a susceptible population member and the transition of an infected population member from the exposed to the infectious state and from the infectious state to the removed state – are independent and occur at random times. The waiting times until these events occur follow Exponential or Gamma distributions. More specific assumptions about the distributions of waiting times and the population contact network are detailed in Section 2.2 and in the monographs of Andersson and Britton (2000) and Britton and Pardoux (2019).

2.2. Parametric population models

Consider an epidemic that started at time 0 and ceased by time $0 < t < \infty$. We assume that the identities of infected population members are known and denoted by $1, \dots, M$, where $M \in \{1, \dots, N\}$. The population contact network is represented by $\mathbf{Y} = \{Y_{i,j}\}_{i < j}^N \in \mathcal{Y} = \{0, 1\}^{\binom{N}{2}}$, where $Y_{i,j} = 1$ if population members i and j are in contact during the epidemic and $Y_{i,j} = 0$ otherwise. Since contacts are undirected and self-contacts are meaningless, we assume that $Y_{i,j} = Y_{j,i}$ and $Y_{i,i} = 0$ hold with probability 1. The transmissions are denoted by $\mathbf{T} = \{T_{i,j}\}_{i \neq j}^N$, where $T_{i,j} = 1$ if i infects j and $T_{i,j} = 0$ otherwise. Observe that $Y_{i,j} = 0$ implies $T_{i,j} = T_{j,i} = 0$ whereas $T_{i,j} = T_{j,i} = 1$ implies $Y_{i,j} = 1$ with probability 1. The starting times of the exposure, infectious, and removal periods of infected population members are denoted by $\mathbf{E} = \{E_i\}_{i=1}^N \in \mathbb{R}_+^N$, $\mathbf{I} = \{I_i\}_{i=1}^N \in \mathbb{R}_+^N$, and $\mathbf{R} = \{R_i\}_{i=1}^N \in \mathbb{R}_+^N$, respectively, where $\mathbb{R}_+ = (0, \infty)$ and $E_i < I_i < R_i$ holds with probability 1; note that E_i, I_i, R_i are undefined if population member i is not infected. We write $\mathbf{X} = (\mathbf{E}, \mathbf{I}, \mathbf{R}, \mathbf{T})$ and, in a mild abuse of language, we refer to \mathbf{X} as an epidemic.

The complete-data likelihood function, given complete observations \mathbf{x} and \mathbf{y} of the epidemic \mathbf{X} and the population contact network \mathbf{Y} , is of the form

$$L(\boldsymbol{\eta}, \boldsymbol{\theta}; \mathbf{x}, \mathbf{y}) \propto L(\boldsymbol{\eta}_E; \mathbf{x}) L(\boldsymbol{\eta}_I; \mathbf{x}) L(\boldsymbol{\beta}; \mathbf{x}, \mathbf{y}) L(\boldsymbol{\theta}; \mathbf{y}),$$

where $\boldsymbol{\eta} = (\boldsymbol{\eta}_E, \boldsymbol{\eta}_I, \boldsymbol{\beta}) \in \Omega_\eta \subseteq \mathbb{R}^{d_1}$ ($d_1 \geq 1$) and $\boldsymbol{\theta} \in \Omega_\theta \subseteq \mathbb{R}^{d_2}$ ($d_2 \geq 1$) are the parameter vectors of the population model generating the epidemic \mathbf{X} and the population contact network \mathbf{Y} , respectively. We describe each component of the complete-data likelihood function in turn, along with its parameters.

The components $L(\boldsymbol{\eta}_E; \mathbf{x})$ and $L(\boldsymbol{\eta}_I; \mathbf{x})$ of the likelihood function are of the form

$$L(\boldsymbol{\eta}_E; \mathbf{x}) \propto \prod_{i=1}^M p(I_i - E_i \mid E_i, \boldsymbol{\eta}_E)$$

$$L(\boldsymbol{\eta}_I; \mathbf{x}) \propto \prod_{i=1}^M p(R_i - I_i \mid I_i, \boldsymbol{\eta}_I),$$

where $p(\cdot \mid E_i, \boldsymbol{\eta}_E)$ and $p(\cdot \mid I_i, \boldsymbol{\eta}_I)$ are densities with suitable support parameterised by $\boldsymbol{\eta}_E$ and $\boldsymbol{\eta}_I$, respectively: e.g. the densities may be Gamma densities, with $\boldsymbol{\eta}_E \in \mathbb{R}_+ \times \mathbb{R}_+$ and $\boldsymbol{\eta}_I \in \mathbb{R}_+ \times \mathbb{R}_+$ being scale and shape parameters of Gamma densities.

Under the assumption that the waiting times until infectious population members infect susceptible population members are independent Exponential(β) random variables with rate of infection $\beta \in \mathbb{R}_+$, the component $L(\beta; \mathbf{x}, \mathbf{y})$ of the likelihood function is given by

$$L(\beta; \mathbf{x}, \mathbf{y}) \propto \beta^{M-1} \exp(-\beta a(\mathbf{x}, \mathbf{y})),$$

where $a(\cdot) > 0$ is defined by

$$a(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^M \sum_{j=1}^M y_{i,j} 1_{I_i < E_j} \max(\min(E_j, R_i) - I_i, 0) + \sum_{i=1}^M d_i(\mathbf{x}, \mathbf{y})(R_i - I_i).$$

Here, $1_{I_i < E_j}$ is 1 in the event of $I_i < E_j$ and is 0 otherwise, and $d_i(\mathbf{x}, \mathbf{y})$ is the number of non-infected population members in contact with population member i . The function $a(\mathbf{x}, \mathbf{y})$ was derived in Britton and O'Neill (2002) and Groendyke et al. (2011).

To represent existing (parametric) and proposed (semiparametric) population models along with possible extensions in a unifying framework, it is convenient to parameterise the distribution of the population contact network Y in exponential-family form (Sundberg 2019). The component $L(\boldsymbol{\theta}; \mathbf{y})$ of the likelihood function can then be written as

$$L(\boldsymbol{\theta}; \mathbf{y}) \propto \exp\left(\boldsymbol{\theta}^\top s(\mathbf{y}) - \psi(\boldsymbol{\theta})\right), \quad \mathbf{y} \in \mathcal{Y},$$

where $\boldsymbol{\theta}$ is a d_2 -vector of natural parameters, $s(\mathbf{y})$ is a d_2 -vector of sufficient statistics, and

$$\psi(\boldsymbol{\theta}) = \log \sum_{\mathbf{y}^* \in \mathcal{Y}} \exp\left(\boldsymbol{\theta}^\top s(\mathbf{y}^*)\right), \quad \boldsymbol{\theta} \in \Omega_\theta = \{\boldsymbol{\theta} \in \mathbb{R}^{d_2} : \psi(\boldsymbol{\theta}) < \infty\} = \mathbb{R}^{d_2}.$$

Britton and O'Neill (2002) and Groendyke et al. (2011) assumed that contacts are independent Bernoulli(μ) ($\mu \in (0, 1)$) random variables, which is equivalent to the one-parameter exponential family with natural parameter $\theta = \text{logit}(\mu) \in \mathbb{R}$ and sufficient statistic $s(\mathbf{y}) = \sum_{i < j} y_{i,j}$. Groendyke et al. (2012) extended the exponential-family framework to include predictors of contacts by assuming that contacts are independent Bernoulli($\mu_{i,j}$) ($\mu_{i,j} \in (0, 1)$) random variables with $\text{logit}(\mu_{i,j}) = \sum_{k=1}^{d_2} \theta_k s_{i,j,k}(y_{i,j}) \in \mathbb{R}$, that is, the log odds of the probability of a contact between two population members is a weighted sum of functions $s_{i,j,k}(y_{i,j})$ of covariates and $y_{i,j}$, weighted by $\theta_k \in \mathbb{R}$ ($k = 1, \dots, d_2$). A specific example is given by $s_{i,j,1}(y_{i,j}) = y_{i,j}$ and $s_{i,j,2}(y_{i,j}) = c_{i,j} y_{i,j}$, where $c_{i,j} \in \{0, 1\}$ is a same-hospital indicator, equal to 1 if population members i and j were in the same hospital during the epidemic and 0 otherwise. The example model is equivalent to a two-parameter exponential family with natural parameters $\theta_1 \in \mathbb{R}$ and $\theta_2 \in \mathbb{R}$ and sufficient statistics $s_1(\mathbf{y}) = \sum_{i < j} s_{i,j,1}(y_{i,j})$ and $s_2(\mathbf{y}) = \sum_{i < j} s_{i,j,2}(y_{i,j})$. If $\theta_2 = 0$, the model of Groendyke et al. (2012) reduces to the model of Britton and O'Neill (2002) and Groendyke et al. (2011).

3. Shortcomings of parametric population models

While the network-based SEIR models reviewed in Section 2.2 are more flexible than classic and lattice-based SIR and SEIR models, such parametric population models have shortcomings. Chief among them is the fact that the induced degree distributions are short-tailed. Here, the degrees of population members are the numbers of contacts of population members.

For example, the model of Britton and O'Neill (2002) and Groendyke et al. (2011) assumes that contacts are independent Bernoulli(μ) ($\mu \in (0, 1)$) random variables. As a consequence, the degrees of population members are Binomial($N - 1, \mu$) and approximately Poisson($N\mu$) distributed provided N is large, μ is small, and $N\mu$ tends to a constant, as one would expect in sparse population contact networks where the expected degrees of population members are bounded above by a finite constant and hence μ is a constant multiple of $1/N$. The model of Groendyke et al. (2012) allows degree distributions to be longer-tailed than the model of Britton and O'Neill (2002) and Groendyke et al. (2011) – depending on available covariates – but the induced degree distribution may nonetheless be shorter-tailed than the degree distributions of real-world contact networks. The degree distributions of real-world contact networks are believed to be long-tailed (e.g. Laumann et al. 1994; Albert and Barabàsi 2002; Jones and Handcock 2003a, 2003b, 2004): e.g. in networks of sexual contacts arising in the study of HIV and other sexually transmitted diseases, some population members tend to have many more sexual contacts than most other population members. The population members in the upper tail of the degree distribution represent an important public health concern, because population members with many contacts can infect many others and are hence potential superspreaders.

Last, but not least, it is worth mentioning that scale-free networks with power law degree distributions (Barabàsi and Albert 1999; Albert and Barabàsi 2002) are known to induce long-tailed degree distributions. However, aside from the fact that the construction of scale-free networks is incomplete and ambiguous (Bollobás, Riordan, Spencer, and Tusnády 2001), those one-parameter models are not flexible models of degree distributions, and proper statistical procedures do not lend much support to informal claims that the degree distributions of many real-world contact networks are scale-free: see, e.g. the work of Jones and Handcock (2003a, 2003b, 2004) on the degree distributions of sexual contact networks arising in the study of HIV, and the discussion of Willinger, Alderson, and Doyle (2009).

Therefore, more flexible population models are needed to accommodate both short- and long-tailed degree distributions.

4. Semiparametric population model

We introduce a semiparametric population model to accommodate both short- and long-tailed degree distributions and detect potential superspreaders. The population model is semiparametric in that the prior of the epidemiological parameter η is parametric, while the prior of the network parameter θ is nonparametric.

Let $s_1(\mathbf{y}), \dots, s_N(\mathbf{y})$ be the degrees of population members $1, \dots, N$, where the degree of population member i is defined as $s_i(\mathbf{y}) = \sum_{j=1: j \neq i}^N y_{i,j}$ ($i = 1, \dots, N$). A simple model of the sequence of degrees $s_1(\mathbf{y}), \dots, s_N(\mathbf{y})$ is given by the exponential family of distributions

with probability mass functions of the form

$$p_{\boldsymbol{\theta}}(\mathbf{y}) = \exp \left(\sum_{i=1}^N \theta_i s_i(\mathbf{y}) - \psi(\boldsymbol{\theta}) \right), \quad \mathbf{y} \in \mathcal{Y}, \quad (1)$$

where the degrees $s_1(\mathbf{y}), \dots, s_N(\mathbf{y})$ are the sufficient statistics and the weights of the degrees $\theta_1, \dots, \theta_N \in \mathbb{R}$ are the natural parameters of the exponential family, while $\psi(\boldsymbol{\theta})$ ensures that $\sum_{\mathbf{y} \in \mathcal{Y}} p_{\boldsymbol{\theta}}(\mathbf{y}) = 1$. The exponential-family form of (1) can be motivated by its maximum entropy property and other attractive mathematical properties (Sundberg 2019). A convenient property is that the resulting likelihood function factorises as follows:

$$L(\boldsymbol{\theta}; \mathbf{y}) \propto \exp \left(\sum_{i=1}^N \theta_i s_i(\mathbf{y}) - \psi(\boldsymbol{\theta}) \right) \propto \prod_{i < j}^N \exp \left(\lambda_{i,j}(\boldsymbol{\theta}) y_{i,j} - \psi_{i,j}(\boldsymbol{\theta}) \right),$$

where

$$\psi(\boldsymbol{\theta}) = \sum_{i < j}^N \psi_{i,j}(\boldsymbol{\theta}),$$

with $\psi_{i,j}(\boldsymbol{\theta})$ and $\lambda_{i,j}(\boldsymbol{\theta})$ defined by

$$\psi_{i,j}(\boldsymbol{\theta}) = \log \left(1 + \exp \left(\lambda_{i,j}(\boldsymbol{\theta}) \right) \right)$$

and

$$\lambda_{i,j}(\boldsymbol{\theta}) = \theta_i + \theta_j,$$

respectively.

To interpret the natural parameters $\theta_1, \dots, \theta_N$ of the exponential family, observe that the model is equivalent to assuming that the contacts $Y_{i,j}$ are independent Bernoulli($\mu_{i,j}$) ($\mu_{i,j} \in (0, 1)$) random variables with $\text{logit}(\mu_{i,j}) = \theta_i + \theta_j \in \mathbb{R}$. Thus, the log odds of the probability of a contact between population members i and j is additive in the propensities of i and j to be in contact with others. It is worth noting that the resulting model can be viewed as an adaptation of the classic p_1 -model (Holland and Leinhardt 1981) to undirected random graphs and is known as the β -model (Chatterjee, Diaconis, and Sly 2011; Rinaldo, Petrović, and Fienberg 2013; Chen, Kato, and Leng 2021).

To cluster population members based on degrees and detect potential superspreaders, we assume that the degree parameters $\theta_1, \dots, \theta_N$ are generated by a Dirichlet process prior (Ferguson 1973; Teh 2010), that is,

$$\begin{aligned} \theta_1 &\sim G \\ \theta_i \mid \theta_1, \dots, \theta_{i-1} &\sim \frac{1}{\alpha + i - 1} \left(\alpha G + \sum_{h=1}^{i-1} \delta_{\theta_h} \right), \quad i = 2, 3, \dots, \end{aligned}$$

where $\alpha > 0$ denotes the concentration parameter and G denotes the base distribution of the Dirichlet process prior, while δ_{θ_h} denotes a point mass at θ_h . A convenient choice of the base distribution G is $N(\mu, \sigma^2)$ ($\mu \in \mathbb{R}, \sigma^2 \in \mathbb{R}_+$).

Draws from a Dirichlet process prior can be generated by generating

- the first draw from G ;
- the i th draw with a probability proportional to α from G , and otherwise drawing one of the existing draws $\theta_1, \dots, \theta_{i-1}$ at random.

4.1. Detecting potential superspreaders

The described construction of Dirichlet process priors reveals that sampling the degree parameters $\theta_1, \dots, \theta_N$ from a Dirichlet process prior implies that some of the degree parameters are resampled. As a consequence, some population members share the same degree parameters. Thus, Dirichlet process priors induce a partition of the population into subpopulations, where subpopulations share the same degree parameters. In applications to real-world data, we can make probabilistic statements about which population members belong to subpopulations with high degree parameters based on the posterior distribution. A short demonstration is provided by the application to the partially observed MERS epidemic in South Korea presented in Section 8.4.

4.2. Short- and long-tailed degree distributions

In addition to detecting potential superspreaders, the semiparametric population model can accommodate short- and long-tailed degree distributions. To demonstrate, we consider a population of size $N = 1,000$ and generate three sets of degree parameters $\theta_1, \dots, \theta_{1000}$ from the Dirichlet process prior with concentration parameter $\alpha = 5$ and base distribution $N(-5, 25)$. Figure 2 shows kernel density plots of the three sets of degree parameters $\theta_1, \dots, \theta_{1000}$ along with the expected degrees of population members. The expected degree of population member i is defined as

$$\mu_i(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \left(\sum_{j=1: j \neq i}^{1000} Y_{i,j} \right) = \sum_{j=1: j \neq i}^{1000} \frac{1}{1 + \exp(-\theta_i - \theta_j)}, \quad i = 1, \dots, N.$$

Figure 2 demonstrates that the distribution of the expected degrees $\mu_1(\boldsymbol{\theta}), \dots, \mu_{1000}(\boldsymbol{\theta})$ can be short- or long-tailed, depending on the degree parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{1000})$. The first set of degree parameters $\theta_1, \dots, \theta_{1000}$ generated from the Dirichlet process prior consists of three subsets of degree parameters, all of them negative. The resulting distribution of the expected degrees resembles a steep mountain with a high peak in a neighbourhood of 0 and a short upper tail. In fact, 90% of all population members have expected degrees of less than 5, and the highest expected degree is less than 39, which is much lower than the highest possible degree of 999 in a population of size 1,000. The second set of generated degree parameters $\theta_1, \dots, \theta_{1000}$ consists of many negative degree parameters and some positive degree parameters between 0 and 5. Since the log odds of the probability of a contact between two population members i and j is $\theta_i + \theta_j$, population members with positive degree parameters can have high to very high expected degrees. Figure 2 shows that the population consists of at least three subpopulations: population members with expected degrees of less than 100; population members with expected degrees between 100 and 200; and population members with expected degree of more than 300. The highest expected

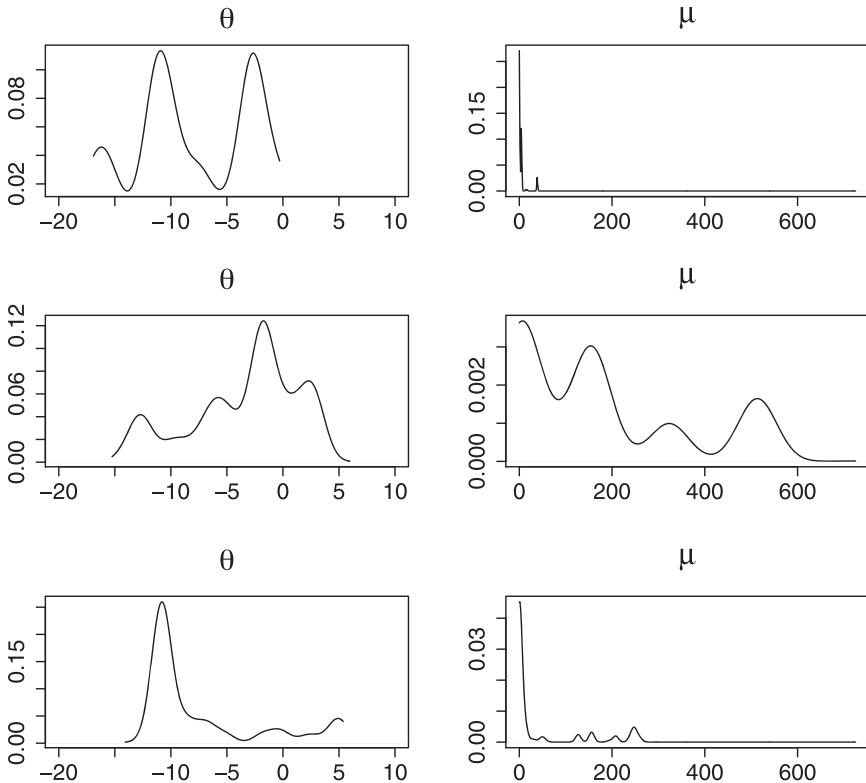


Figure 2. Three sets of draws $\theta_1, \dots, \theta_{1000}$ from the Dirichlet process prior with concentration parameter $\alpha = 5$ and base distribution $N(-5, 25)$. Left: Distribution of degree parameters $\theta_1, \dots, \theta_{1000}$. Right: Distribution of expected degrees $\mu_1(\theta), \dots, \mu_{1000}(\theta)$. Each row corresponds to one set of draws from the Dirichlet process prior.

degree is about 722. The resulting distribution of the expected degrees is both multimodal and long-tailed. The third set of generated degree parameters $\theta_1, \dots, \theta_{1000}$ resembles the second set of generated degree parameters, in that the distribution of the expected degrees is multimodal and long-tailed. That said, the third set of draws is less extreme than the second set of draws: e.g. the highest expected degree is about 259 rather than 722.

The examples presented above demonstrate that Dirichlet process priors with Gaussian base distributions can accommodate both short- and long-tailed degree distributions, despite the fact that Gaussians are symmetric and unimodal distributions with light tails. Other, non-Gaussian base distributions can be chosen. As a consequence, the proposed semiparametric population model is flexible and can accommodate a wide range of degree distributions with countless forms and shapes, including short- and long-tailed degree distributions.

5. Incomplete data

In practice, complete observations of population contact networks and epidemics are rare. While population-level data (e.g. counts of the total number of infected, recovered, or

deceased population members) may be disseminated by public health authorities and can be collected by scraping websites and other channels of communication, collecting individual-level data (e.g. the contacts, exposure, infectious, and removal times of infected population members) requires substantial investments in terms of time and resources, making it all but impossible to collect all relevant data. We first discuss possible reasons for incomplete data in Section 5.1 and then stress the importance of collecting network data in Section 5.2.

5.1. Possible reasons for incomplete data

There are many reasons for the fact that available data are, more often than not, incomplete. Some of the possible reasons are:

- Epidemics are rare events that occur at random times and in random places, and when such rare events do occur, public health officials and scientists may not be well-prepared to collect relevant data without advance notice.
- Ethical and legal considerations can make the collection of data on individual population members challenging, if not impossible: e.g. if there was universal cell phone coverage and all population members carried cell phones at all times, collecting data on contacts among population members would be straightforward by monitoring the locations of cell phones. However, collecting such data would violate laws that protect the privacy of population members.
- Epidemics are not limited to urban areas with excellent infrastructure and ready access to public resources, but may occur in remote corners of the planet: e.g. the most recent outbreaks of Ebola started in remote areas of Africa. Worse, some areas with outbreaks may be war-torn. As a result, researchers may not be able to collect data by visiting areas with outbreaks without exposing themselves and others to unacceptable risks.
- In addition, there are more mundane reasons for incomplete data, such as
 - *design-based mechanisms*: e.g. motivated by financial constraints, researchers may sample population members, which implies that a sampling design determines which data are collected;
 - *out-of-design mechanisms*: e.g. population members refuse to share data when the data are considered sensitive.

5.2. Importance of collecting network data

Data can be incomplete because transmissions, exposure, infectious, or removal times of infected population members are unobserved, but many available data sets share one fundamental weakness: There are no data on contacts among population members. The lack of network data is all the more striking, because collecting network data would help reduce the posterior uncertainty about

- (a) the population contact network Y , which imposes hard constraints on how an infectious disease can spread, because $Y_{i,j} = 0$ (no contact) implies $T_{i,j} = T_{j,i} = 0$ (no transmission) with probability 1;

- (b) the parameter θ of the population model that generated the population contact network Y ;
- (c) possible sources of infections.

Advantages (a) and (b) may not be too surprising, but advantage (c) may be less obvious. To demonstrate that sampling contacts can reduce the uncertainty about possible sources of infections, it is instructive to inspect the full conditional probability of the event that a population member i infected a population member j , assuming that both of them were infected during the course of the epidemic. If $\varphi(i \text{ infected } j)$ denotes the prior probability of the event that population member i infected population member j , then the conditional probability of the event that population member i infected population member j , given everything else, takes the form

$$\mathbb{P}(i \text{ infected } j \mid \mathbf{E}, \mathbf{I}, \mathbf{R}, \mathbf{Y} = \mathbf{y}) = \frac{y_{i,j} \mathbf{1}_{I_i < E_j < R_i} \varphi(i \text{ infected } j)}{\sum_{h=1: h \neq j}^M y_{h,j} \mathbf{1}_{I_h < E_j < R_h} \varphi(h \text{ infected } j)}.$$

If $y_{1,j}, \dots, y_{M,j}$ are observed, then $y_{1,j}, \dots, y_{M,j}$ are fixed and impose hard constraints on who could have infected j : If i was not in contact with j (that is, $y_{i,j} = 0$), i could not have infected j . By contrast, if $y_{1,j}, \dots, y_{M,j}$ are unobserved, then $y_{1,j}, \dots, y_{M,j}$ are not fixed and need to be inferred, increasing the uncertainty about the possible sources of infection of j . In other words, observed contacts help narrow down the possible sources of infections and, in so doing, help reduce the uncertainty about possible sources of infections.

We describe two sampling designs for generating samples of contacts along with epidemiological data: ego-centric sampling and link-tracing. Some background on ego-centric sampling and link-tracing of contacts (but not epidemiological data) can be found in Thompson and Frank (2000), Handcock and Gile (2010), and Krivitsky and Morris (2017). In the literature on network sampling (Frank 1988; Thompson and Frank 2000; Gile and Handcock 2006; Handcock and Gile 2010), popular forms of link-tracing are snowball sampling (Goodman 1961) and respondent-driven sampling (Heckathorn 1997; Salganik and Heckathorn 2004; Gile and Handcock 2010; Gile 2011; Kurant, Markopoulou, and Thiran 2011). Some of them do not generate probability samples in the strict sense of the word, but generate approximate probability samples (e.g. Gile 2011). A recent review of these and other network sampling designs can be found in Schweinberger et al. (2020). We adapt these ideas to sampling contacts along with epidemiological data.

An ego-centric sample of contacts and epidemiological data can be generated as follows:

- (a) Generate a probability sample of population members.
- (b) For each sampled population member, collect data on the contacts of the population member and, should the population member be infected, data on possible sources of infection along with the exposure, infectious, and removal times of the population member.

A probability sample of population members can be generated by any sampling design for sampling from finite populations (e.g. Thompson 2012).

An interesting extension of ego-centric sampling is link-tracing. Link-tracing exploits the observed contacts of sampled population members to include additional population members into the sample. A k -wave link-tracing sample of contacts and epidemiological data can be generated as follows:

- (1) Wave $l = 0$: Generate an ego-centric sample.
- (2) Wave $l = 1, \dots, k$:
 - (a) Add the population members who are connected to the population members of wave $l-1$ to the sample.
 - (b) For each added population member, collect all relevant data.

Ego-centric sampling can be considered to be a special case of k -wave link-tracing with $k = 0$. We use ego-centric sampling in the simulation study in Section 7.

6. Bayesian inference

We discuss Bayesian inference for the parameters η and θ of the population model based on incomplete data. Since interest centers on the population model, it is natural to ask: Under which conditions is the process that determines which data are observed ignorable for the purpose of Bayesian inference for the parameters η and θ of the population model? To answer the question, we start from first principles. We first separate the complete-data generating process from the incomplete-data generating process:

- The *complete-data generating process* is the process that generates the complete data, that is, the process that generates a realisation (\mathbf{x}, \mathbf{y}) of (\mathbf{X}, \mathbf{Y}) .
- The *incomplete-data generating process* is the process that determines which subset of the complete data (\mathbf{x}, \mathbf{y}) is observed.

A failure to separate these processes can lead to misleading conclusions, as pointed out by Rubin (1976), Dawid and Dickey (1977), Koskinen, Robins, and Pattison (2010), Handcock and Gile (2010, 2017), Crane (2018), and Schweinberger et al. (2020).

We therefore proceed as follows. We first separate the complete-data generating process from the incomplete-data generating process in Section 6.1 and then discuss Bayesian inference based on incomplete data in Section 6.2. We then discuss Bayesian computing in Sections 6.3, 6.4, and 6.5.

6.1. Complete- and incomplete-data generating process

To prepare the ground for principled Bayesian inference, we separate the complete-data generating process from the incomplete-data generating process, adapting the generic ideas of Rubin (1976) to stochastic models of epidemics.

To do so, let $\mathbf{A} = \{A_E, A_I, A_R, A_T, A_Y\}$ be indicators of which data are observed, where $A_E = \{A_{E,i}\}_{i=1}^N \in \{0, 1\}^N$, $A_I = \{A_{I,i}\}_{i=1}^N \in \{0, 1\}^N$, $A_R = \{A_{R,i}\}_{i=1}^N \in \{0, 1\}^N$, $A_T = \{A_{T,i,j}\}_{i \neq j}^M \in \{0, 1\}^{M^2 - M}$, and $A_Y = \{A_{Y,i,j}\}_{i < j}^N \in \{0, 1\}^{\binom{N}{2}}$ indicate whether the values of $\{E_i\}_{i=1}^N$, $\{I_i\}_{i=1}^N$, $\{R_i\}_{i=1}^N$, $\{T_{i,j}\}_{i \neq j}^M$, and $\{Y_{i,j}\}_{i < j}^N$ are observed, respectively. The sequence of

indicators \mathbf{A} is considered to be a random variable, with a distribution parameterised by $\boldsymbol{\vartheta} \in \Omega_{\boldsymbol{\vartheta}} \subseteq \mathbb{R}^q$ ($q \geq 1$): e.g. $\boldsymbol{\vartheta} \in [0, 1]^N$ may be a vector of sample inclusion probabilities, where $\vartheta_i \in [0, 1]$ is the probability that population member $i \in \{1, \dots, N\}$ is sampled and data on the contacts of population member i are collected. The parameter $\boldsymbol{\vartheta}$ may be known or unknown. We focus on the more challenging case where $\boldsymbol{\vartheta}$ is unknown. The observed and unobserved subset of the complete data (\mathbf{x}, \mathbf{y}) are denoted by \mathbf{x}_{obs} and \mathbf{x}_{mis} and \mathbf{y}_{obs} and \mathbf{y}_{mis} , respectively, where $\mathbf{x} = (\mathbf{x}_{\text{obs}}, \mathbf{x}_{\text{mis}})$ and $\mathbf{y} = (\mathbf{y}_{\text{obs}}, \mathbf{y}_{\text{mis}})$.

In Bayesian fashion, we build a joint probability model for all knowns and unknowns and condition on all knowns. Let

$$p(\mathbf{a}, \mathbf{x}, \mathbf{y}, \boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}) = p(\mathbf{a}, \mathbf{x}, \mathbf{y} \mid \boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta})p(\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta})$$

be the joint probability density of $\mathbf{a}, \mathbf{x}, \mathbf{y}, \boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}$, where

$$p(\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}) = p(\boldsymbol{\vartheta} \mid \boldsymbol{\eta}, \boldsymbol{\theta})p(\boldsymbol{\eta}, \boldsymbol{\theta})$$

is the prior probability density of $\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}$ and

$$p(\mathbf{a}, \mathbf{x}, \mathbf{y} \mid \boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}) = p(\mathbf{a} \mid \mathbf{x}, \mathbf{y}, \boldsymbol{\vartheta})p(\mathbf{x} \mid \mathbf{y}, \boldsymbol{\eta})p(\mathbf{y} \mid \boldsymbol{\theta})$$

is the conditional probability density of $\mathbf{a}, \mathbf{x}, \mathbf{y}$ given $\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}$; note that $p(\mathbf{y} \mid \boldsymbol{\theta}) \equiv p_{\boldsymbol{\theta}}(\mathbf{y})$. It is worth noting that all of these probability densities are with respect to a suitable dominating measure, but we do not wish to delve into measure-theoretic details, which would distract from the main ideas.

To determine when the incomplete-data generating process is ignorable for the purpose of Bayesian inference for the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ of the population model, we introduce the following definition of a likelihood-ignorable incomplete-data generating process.

Definition (likelihood-ignorable incomplete-data generating process): Assume that the parameters $\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}$ are variation-independent in the sense that the parameter space of $(\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta})$ is given by a product space of the form $\Omega_{\boldsymbol{\vartheta}} \times \Omega_{\boldsymbol{\eta}} \times \Omega_{\boldsymbol{\theta}}$ and that the parameters of the population model $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ and the parameter of the incomplete-data generating process $\boldsymbol{\vartheta}$ are independent under the prior,

$$p(\boldsymbol{\vartheta} \mid \boldsymbol{\eta}, \boldsymbol{\theta}) = p(\boldsymbol{\vartheta}) \quad \text{for all } (\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta}) \in \Omega_{\boldsymbol{\vartheta}} \times \Omega_{\boldsymbol{\eta}} \times \Omega_{\boldsymbol{\theta}}.$$

If the probability of observing data does not depend on the values of the unobserved data,

$$p(\mathbf{a} \mid \mathbf{x}, \mathbf{y}, \boldsymbol{\vartheta}) = p(\mathbf{a} \mid \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}, \boldsymbol{\vartheta}) \quad \text{for all } \mathbf{a}, \mathbf{x}, \mathbf{y} \text{ and all } \boldsymbol{\vartheta} \in \Omega_{\boldsymbol{\vartheta}},$$

then the incomplete-data generating process is called likelihood-ignorable, and otherwise non-ignorable.

We provide two examples of likelihood-ignorable and non-ignorable incomplete-data generating processes.

Example (likelihood-ignorable): All infected population members visit hospitals, which record data on contacts, transmissions, exposure, infectious, and removal times of infected population members. To reduce the posterior uncertainty about the population contact

network and its generating mechanism, investigators generate a probability sample of non-infected population members from the subpopulation of all non-infected population members and collect data on the contacts of sampled population members, rather than limiting the collection of data to infected population members visiting hospitals.

A possible sampling design for generating samples of contacts and epidemiological data is link-tracing, as described in Section 5.2. Link-tracing exploits observed contacts of population members to increase the size of the sample, which implies that the sample inclusion probabilities depend on observed contacts but do not depend on unobserved contacts:

$$p(\mathbf{a} | \mathbf{x}, \mathbf{y}, \boldsymbol{\vartheta}) = p(\mathbf{a} | \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}, \boldsymbol{\vartheta}) \quad \text{for all } \mathbf{a}, \mathbf{x}, \mathbf{y} \text{ and all } \boldsymbol{\vartheta} \in \Omega_{\boldsymbol{\vartheta}}.$$

Therefore, link-tracing sampling designs for generating samples of contacts and epidemiological data are likelihood-ignorable, as are ego-centric sampling designs.

Example (non-likelihood-ignorable): Suppose that there exists a constant $\delta > 0$ such that infected population members i with mild symptoms and fast recovery ($R_i - I_i \leq \delta$) do not visit hospitals, whereas infected population members i with severe symptoms and slow recovery ($R_i - I_i > \delta$) do visit hospitals. Hospitals collect data on infected population members who visit them, but no data are collected on other population members.

Since the incomplete-data generating process excludes all infected population members with mild symptoms and fast recovery, it cannot be ignored. Statistical analyses ignoring it may give rise to misleading conclusions about the rate of infection β and other parameters of the population model, and may generate misleading predictions of future epidemics in the population of interest and similar populations.

6.2. Bayesian inference based on incomplete data

Separating the complete-data generating process from the incomplete-data generating process paves the way for principled Bayesian inference based on incomplete data.

The following result shows that, as long as the incomplete-data generating process is likelihood-ignorable, Bayesian inference for the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ of the population model can be based on the marginal posterior $p(\boldsymbol{\eta}, \boldsymbol{\theta} | \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}})$. In other words, the incomplete-data generating process can be ignored for the purpose of Bayesian inference for the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ of the population model, because the marginal posterior $p(\boldsymbol{\eta}, \boldsymbol{\theta} | \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}})$ can be computed without requiring knowledge about the incomplete-data generating process.

Proposition 6.1: *If the incomplete-data generating process is likelihood-ignorable and the prior $p(\boldsymbol{\eta}, \boldsymbol{\theta})$ is proper, the parameter $\boldsymbol{\vartheta}$ of the incomplete-data generating process and the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ of the population model are independent under the posterior,*

$$p(\boldsymbol{\vartheta}, \boldsymbol{\eta}, \boldsymbol{\theta} | \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}) \propto p(\boldsymbol{\vartheta} | \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}) p(\boldsymbol{\eta}, \boldsymbol{\theta} | \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}),$$

and Bayesian inference for the parameters $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ of the population model can be based on the marginal posterior $p(\boldsymbol{\eta}, \boldsymbol{\theta} \mid \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}})$,

$$p(\boldsymbol{\eta}, \boldsymbol{\theta} \mid \mathbf{a}, \mathbf{x}_{\text{obs}}, \mathbf{y}_{\text{obs}}) = \frac{\sum_{\mathbf{y}_{\text{mis}}} \int p(\mathbf{x} \mid \mathbf{y}, \boldsymbol{\eta}) p(\mathbf{y} \mid \boldsymbol{\theta}) p(\boldsymbol{\eta}, \boldsymbol{\theta}) d\mathbf{x}_{\text{mis}}}{\sum_{\mathbf{y}_{\text{mis}}} \int \int \int p(\mathbf{x} \mid \mathbf{y}, \boldsymbol{\eta}) p(\mathbf{y} \mid \boldsymbol{\theta}) p(\boldsymbol{\eta}, \boldsymbol{\theta}) d\mathbf{x}_{\text{mis}} d\boldsymbol{\eta} d\boldsymbol{\theta}}, \quad (2)$$

which can be computed without requiring knowledge about the incomplete-data generating process.

It is worth noting that the unobserved contacts of non-infected population members can be summed out, both in the numerator and the denominator of the ratio in (2), because the contacts of population members are independent conditional on $\boldsymbol{\theta}$ by construction of the semiparametric population model described in Section 4. In other words, there is no need to use computational methods (e.g. Markov chain Monte Carlo methods) to generate model-based imputations of unobserved contacts of non-infected population members, which saves computing time.

A proof of Proposition 6.1 can be found in the supplement. The case of Britton and O'Neill (2002) and Groendyke et al. (2011, 2012), who considered Bayesian inference from observed infectious and removal times, is a special case of the incomplete-data framework considered here, with $\mathbf{x}_{\text{obs}} = \{\mathbf{I}, \mathbf{R}\}$ and $\mathbf{y}_{\text{obs}} = \{\}$. In general, when \mathbf{x} or \mathbf{y} are partially observed, Bayesian inference can be based on the marginal posterior of $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ given the observed data, provided the incomplete-data generating process is likelihood-ignorable.

6.3. Truncated Dirichlet process priors

To facilitate Markov chain Monte Carlo sampling from the posterior distribution, we approximate Dirichlet process priors by truncated Dirichlet process priors along the lines of Ishwaran and James (2001).

The truncation of Dirichlet process priors takes advantage of the stick-breaking construction of Dirichlet process priors (Sethuraman 1994; Ishwaran and James 2001; Teh 2010). A stick-breaking construction of a Dirichlet process prior with base distribution G and concentration parameter α proceeds as follows. First, we sample parameters from the base distribution G :

$$\gamma_k \stackrel{\text{iid}}{\sim} G, \quad k = 1, 2, \dots$$

We then construct mixing proportions π_1, π_2, \dots by first sampling

$$V_k \mid \alpha \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha), \quad k = 1, 2, \dots$$

and then setting

$$\pi_1 = V_1$$

$$\pi_k = V_k \prod_{j=1}^{k-1} (1 - V_j), \quad k = 2, 3, \dots$$

Last, but not least, we construct an infinite mixture distribution with mixing proportions π_1, π_2, \dots and point masses $\delta_{\gamma_1}, \delta_{\gamma_2}, \dots$ as follows:

$$\mathcal{P}_\infty = \sum_{k=1}^{\infty} \pi_k \delta_{\gamma_k}.$$

The distribution \mathcal{P}_∞ is then a draw from the Dirichlet process prior with concentration parameter α and base distribution G .

A Dirichlet process prior can be truncated by choosing a positive integer K and setting $V_K = 1$, which implies that $\pi_{K+1} = 0, \pi_{K+2} = 0, \dots$. A finite mixture distribution with mixing proportions $\pi_1, \pi_2, \dots, \pi_K$ and point masses $\delta_{\gamma_1}, \delta_{\gamma_2}, \dots, \delta_{\gamma_K}$ can then be constructed as follows:

$$\mathcal{P}_K = \sum_{k=1}^K \pi_k \delta_{\gamma_k}.$$

The distribution \mathcal{P}_K can be regarded as a draw from the Dirichlet process prior with concentration parameter α and base distribution G truncated at K . If K is large, the truncated Dirichlet process prior is expected to be a good approximation of the Dirichlet process prior. Some theoretical guidance regarding the choice of K can be found in Ishwaran and James (2001). In practice, K can be chosen by

- selecting a large positive integer, such as $K = N$;
- exploiting domain knowledge;
- choosing a value of K that leads to acceptable in- or out-of-sample performance.

The truncated stick-breaking construction of π implies that π is generalised Dirichlet distributed, which is conjugate to multinomial sampling (Ishwaran and James 2001) and facilitates Markov chain Monte Carlo sampling from the posterior distribution.

Since the concentration parameter α and the parameters μ and σ^2 of the base distribution $N(\mu, \sigma^2)$ are unknown, it is natural to express the uncertainty about α , μ , and σ^2 by assuming that α , μ , and σ^2 have hyperpriors. We assume that the hyperpriors of the hyperparameters α , μ , and $1/\sigma^2$ are Gamma, Gaussian, and Gamma distributions, respectively, which are conjugate priors and facilitate Markov chain Monte Carlo sampling from the posterior distribution.

6.4. Bayesian Markov chain Monte Carlo algorithm

A Bayesian Markov chain Monte Carlo algorithm for sampling from the posterior distribution is described in the supplement. We list here the main steps of the algorithm, without providing details, and discuss its computing time. Details can be found in the supplement.

To list the main steps of the algorithm, let

- $Z_{i,k} = 1$ if population member i was assigned degree parameter γ_k and $Z_{i,k} = 0$ otherwise ($i = 1, \dots, N, k = 1, \dots, K$);
- $\mathbf{Z}_i = (Z_{i,1}, \dots, Z_{i,K})$ and $\mathbf{Z}_{-i} = (\mathbf{Z}_1, \dots, \mathbf{Z}_{i-1}, \mathbf{Z}_{i+1}, \dots, \mathbf{Z}_N)$ ($i = 1, \dots, N$), and $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_N)$;

- $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_K)$;
- $\boldsymbol{\pi} = (\pi_1, \dots, \pi_K)$.

As a consequence, the degree parameter θ_i of population member i can be expressed as

$$\theta_i = \mathbf{Z}_i^\top \boldsymbol{\gamma}, \quad i = 1, \dots, N.$$

A Markov chain Monte Carlo algorithm for sampling from the posterior distribution can then be constructed by combining the following Markov chain Monte Carlo steps by means of cycling or mixing (Tierney 1994; Liu 2008), assuming all unknown quantities have been initialised:

- (1) Impute the missing data:
 - (a) Sample $\mathbf{X}_{\text{mis}} \mid \mathbf{X}_{\text{obs}} = \mathbf{x}_{\text{obs}}, \mathbf{Y} = \mathbf{y}, \boldsymbol{\eta}$.
 - (b) Sample $\mathbf{Y}_{\text{mis}} \mid \mathbf{X} = \mathbf{x}, \mathbf{Y}_{\text{obs}} = \mathbf{y}_{\text{obs}}, \mathbf{Z} = \mathbf{z}, \boldsymbol{\gamma}, \boldsymbol{\eta}$.
- (2) Sample the parameters of the population model:
 - (a) Sample $\boldsymbol{\gamma} \mid \mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z}$.
 - (b) Sample $\mathbf{Z}_i \mid \mathbf{Y} = \mathbf{y}, \mathbf{Z}_{-i} = \mathbf{z}_{-i}, \boldsymbol{\gamma}, \boldsymbol{\pi}$ ($i = 1, \dots, N$).
 - (c) Set $\theta_i = \mathbf{Z}_i^\top \boldsymbol{\gamma}$ ($i = 1, \dots, N$).
 - (d) Sample $\boldsymbol{\eta} \mid \mathbf{X} = \mathbf{x}, \mathbf{Y} = \mathbf{y}$.
- (3) Sample the hyperparameters:
 - (a) Sample $\alpha \mid \boldsymbol{\pi}$.
 - (b) Sample $\boldsymbol{\pi} \mid \mathbf{Z} = \mathbf{z}, \alpha$.
 - (c) Sample $\mu \mid \sigma^2, \boldsymbol{\gamma}$.
 - (d) Sample $\sigma^2 \mid \mu, \boldsymbol{\gamma}$.

Most of the Markov chain Monte Carlo steps involve Gibbs sampling from full conditional distributions (e.g. Beta, Gamma, and Gaussian distributions), while the others are Metropolis-Hastings steps. More details are provided in the supplement.

The computing time of the algorithm is a function of

- the number of subpopulations K , which satisfies $K \leq N$;
- the number of infected population members M , which satisfies $M \leq N$ and in large populations satisfies $M \ll N$ (unless a non-negligible fraction of the population is infected);
- the sampling design and the number of population members n sampled out of the N population members for the purpose of collecting data on contacts along with epidemiological data, which satisfies $n \leq N$ and in large populations satisfies $n \ll N$ (unless a non-negligible fraction of the population is sampled);
- the sparsity of the population contact network.

As a specific example, consider ego-centric sampling, as described in Section 5.2. An ego-centric sampling design samples n out of the N population members and, for each sampled population member, collects data on the contacts of the sampled population member with the $N-1$ other population members, in addition to data on the transmissions, exposure, infectious, and removal times of infected population members. As a result, the computing time of each iteration of the Bayesian Markov chain Monte Carlo algorithm is

$O(KnN)$, because updates of the K degree parameters $\gamma_1, \dots, \gamma_K$ involve computations of up to $n(N - 1)$ probabilities of the form

$$\mathbb{P}_{\theta_i, \theta_j}(Y_{i,j} = y_{i,j}) = \exp((\theta_i + \theta_j)y_{i,j} - \log(1 + \exp(\theta_i + \theta_j))),$$

where $\theta_i = \mathbf{Z}_i^\top \boldsymbol{\gamma}$ and $\theta_j = \mathbf{Z}_j^\top \boldsymbol{\gamma}$.

Having said that, the computing time of $O(KnN)$ in the case of ego-centric sampling is based on worst-case scenarios and can be reduced to $O(Kn)$ in special cases: e.g. when the population contact network is sparse in the sense that many population members have few contacts, it is possible to reduce the computing time by taking advantage of sparsity. In fact, many real-world networks are sparse: While population members can create up to $N - 1$ contacts, creating physical contacts that enable disease transmission requires geographical proximity and time. As a consequence, it is plausible that the expected degrees of many if not all population members are bounded above by a finite constant (Dunbar 1992; Rohe, Chatterjee, and Yu 2011; Krivitsky, Handcock, and Morris 2011; Lovász 2012; Amini, Chen, Bickel, and Levina 2013; Krivitsky and Kolaczyk 2015; Butts 2019). In such cases, the resulting population contact network is sparse, and one could reduce the worst-case computing time of $O(KnN)$ to $O(Kn)$. Some ideas of how to exploit sparsity for the purpose of reducing computing time can be found in, e.g. Raftery, Niu, Hoff, and Yeung (2012) and Vu, Hunter, and Schweinberger (2013).

6.5. Addressing the label-switching problem

Markov chain Monte Carlo samples from the posterior distribution may show evidence of label-switching, that is, the labelling of subpopulations may have switched from iteration to iteration of the Markov chain Monte Carlo algorithm. The label-switching problem is rooted in the fact that the likelihood function is invariant to permutations of the labels of subpopulations. While the prior is not invariant to permutations of the labels of subpopulations, the prior is dominated by the likelihood function when the data is informative. As a consequence, the labels of subpopulations may switch from iteration to iteration of the Markov chain Monte Carlo algorithm. Label-switching can give rise to misleading conclusions about parameters that depend on the labelling of the subpopulations, including the indicators $\mathbf{Z}_1, \dots, \mathbf{Z}_N$ and the degree parameters $\gamma_1, \dots, \gamma_K$ and $\theta_1 = \mathbf{Z}_1^\top \boldsymbol{\gamma}, \dots, \theta_N = \mathbf{Z}_N^\top \boldsymbol{\gamma}$. The label-switching problem is a well-known problem in the literature concerned with Bayesian Markov chain Monte Carlo estimation of finite mixture models and related models. We follow the Bayesian decision-theoretic approach of Stephens (2000) to undoing the label-switching, by minimising the posterior expectation of a well-chosen loss function. A loss function and relabelling algorithm are stated in Schweinberger and Handcock (2015) and implemented in R package `hernm` (Schweinberger and Luna 2018). We use them in Sections 7 and 8 to undo the label-switching in all Markov chain Monte Carlo samples generated by the Bayesian Markov chain Monte Carlo algorithm.

7. Simulation results

We explore the frequentist properties of Bayesian point estimators and the reduction in statistical error due to collecting network data by using simulations. We consider a

population of size 187 consisting of $K = 3$ subpopulations labelled 1, 2, 3. The three subpopulations consist of low-, moderate-, and high-degree population members. We assign population members i to subpopulations 1, 2, 3 by sampling $\mathbf{Z}_i \stackrel{\text{iid}}{\sim} \text{Multinomial}(1; \pi = (.4, .3, .3))$ ($i = 1, \dots, N$). We then generate a population contact network from the population model described in Section 4 with degree parameters $\theta_i = \mathbf{Z}_i^\top \boldsymbol{\gamma}$ ($i = 1, \dots, N$). Conditional on the population contact network, an epidemic is generated by the stochastic model described in Section 2, assuming that $I_i - E_i$ and $R_i - I_i$ are independent $\text{Gamma}(\eta_{E,1}, \eta_{E,2})$ and $\text{Gamma}(\eta_{I,1}, \eta_{I,2})$ random variables, respectively ($i = 1, \dots, M$). The data-generating values of the parameters are specified in Sections 7.1 and 7.2. Unless stated otherwise, we assume that the exposure, infectious, and removal times $\mathbf{E}, \mathbf{I}, \mathbf{R}$ are observed, whereas the transmissions \mathbf{T} are unobserved, as are the population contact network \mathbf{Y} and the indicators \mathbf{Z} .

7.1. Simulation results quantifying the error of estimation

We generate 1,000 population contact networks and epidemics as described above. The data-generating values of the parameters $\beta, \eta_{E,1}, \eta_{E,2}, \eta_{I,1}, \eta_{I,2}, \gamma_1, \gamma_2$, and γ_3 are shown in Table 1. For each data set, we truncate the Dirichlet process prior at $K = 3$ and $K = 5$ as described in Section 6.3, and estimate the data-generating model by the Bayesian Markov chain Monte Carlo algorithm described in Section 6.4.

Table 1 sheds light on the frequentist coverage properties of 95% posterior credible intervals. The simulation results indicate that the frequentist coverage properties of posterior credible intervals are excellent in the case of the epidemiological parameters $\beta, \eta_{E,1}, \eta_{E,2}, \eta_{I,1}$, and $\eta_{I,2}$, but less so in the case of the network parameters γ_1, γ_2 , and γ_3 . These results underscore the challenge of estimating network parameters without observing network data. Section 7.2 demonstrates that the statistical error can be reduced by collecting network data.

7.2. Simulation results quantifying the effect of network sampling

To demonstrate that collecting network data can reduce the posterior uncertainty about the parameters of the population model, we consider a population consisting of $K = 3$ subpopulations. The $K = 3$ subpopulations correspond to

- a low-degree subpopulation of size 127 with degree parameter $\gamma_1 = -3.5$;
- a moderate-degree subpopulation of size 50 with degree parameter $\gamma_2 = -1.5$;
- a high-degree subpopulation of size 10 with degree parameter $\gamma_3 = .5$.

Table 1. Simulation results. Coverage properties of 95% posterior credible intervals: number of times 95% posterior credible intervals include the data-generating values of the parameters $\beta, \eta_{E,1}, \eta_{E,2}, \eta_{I,1}, \eta_{I,2}, \gamma_1, \gamma_2$, and γ_3 in %, using a Dirichlet process prior truncated at $K = 3$ and $K = 5$.

Parameter	β	$\eta_{E,1}$	$\eta_{E,2}$	$\eta_{I,1}$	$\eta_{I,2}$	γ_1	γ_2	γ_3
True value	2	8	.25	4	.25	-2	-1	0
$K = 3$	95.4%	95.4%	93.5%	95.2%	94.3%	83.4%	97.4%	89.4%
$K = 5$	93.0%	96.3%	96.3%	96.0%	95.7%	98.4%	100%	96.7%

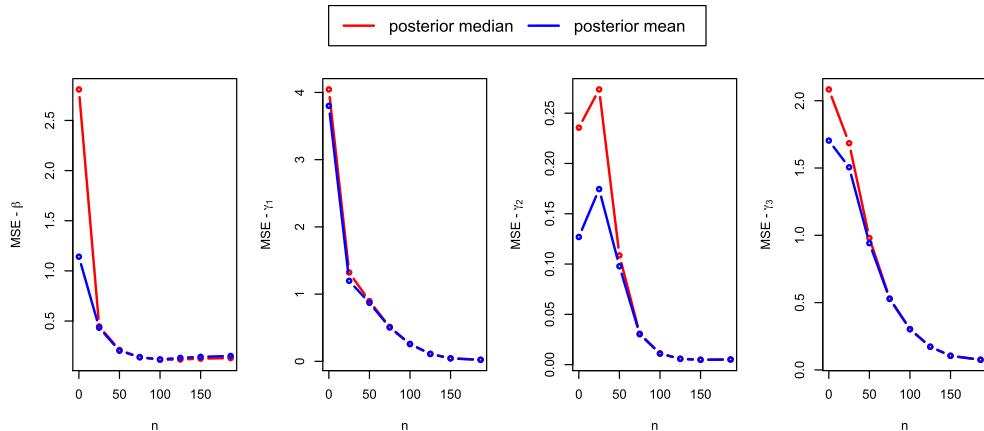


Figure 3. Simulation results: MSE of the posterior median and mean of the parameters β , γ_1 , γ_2 , and γ_3 plotted against the sample size $n = 0, 25, 50, 75, 100, 125, 150, 187$. If $n = 0$, no contacts are observed, otherwise a subset of contacts is observed.

We generate 1,000 ego-centric samples of sizes $n = 25, 50, 75, 100, 125, 150, 187$ from the population of size $N = 187$. We then estimate the population model from each sample of contacts along with observations of the exposure, infectious, and removal times of infected population members. In addition, we estimate the population model without observations of contacts, which corresponds to a sample size of $n = 0$, using observations of the exposure, infectious, and removal times of infected population members. To assess how much the posterior uncertainty about the parameters of the population model is reduced by sampling contacts, we use the mean squared error (MSE) of the posterior median and mean of the parameters.

The MSE of the posterior median and mean of the rate of infection β and the degree parameters γ_1 , γ_2 , and γ_3 is plotted in Figure 3 against the sample size $n = 0, 25, 50, 75, 100, 125, 150, 187$. By construction of the model, estimators of the epidemiological parameters $\eta_{E,1}$, $\eta_{E,2}$, $\eta_{I,1}$, and $\eta_{I,2}$ are not expected to be sensitive to n – which determines how much information is available about the network parameters – and the MSE of the posterior median and mean of $\eta_{E,1}$, $\eta_{E,2}$, $\eta_{I,1}$, and $\eta_{I,2}$ are indeed not sensitive to n (not shown). However, Figure 3 demonstrates that samples of contacts do reduce the MSE of the posterior median and mean of β , γ_1 , γ_2 , and γ_3 : The MSE turns out to be highest when $n = 0$, and rapidly decreases as n increases. These observations underscore the importance of collecting data on contacts or functions of contacts, such as degrees.

8. Partially observed MERS epidemic in South Korea

We showcase the statistical framework introduced in Sections 4, 5, and 6 by applying it to the partially observed MERS epidemic in South Korea in 2015 (Ki 2015).

8.1. Data

The MERS outbreak in South Korea was driven by the coronavirus MERS, which is related to the coronaviruses SARS and COVID-19. We retrieved the data from the website

<http://english.mw.go.kr> of the South Korean Ministry of Health and Welfare on September 22, 2015. The website has been removed since, but the data can be obtained from the authors. The first MERS case in South Korea was confirmed on May 20, 2015 and the last case was confirmed on July 4, 2015. By September 22, 2015, 186 cases had been confirmed, of which 144 had recovered while 35 had died, all of whom are considered to be removed from the population. 7 cases had not removed by September 22, 2015, despite the fact that the last case of MERS was confirmed on July 4, 2015. We assume that the removal times of those 7 cases are unobserved and that the outbreak ceased by September 22, 2015, because the data base has not been updated by the South Korean Ministry of Health and Welfare since July 2015. The data consist of the infectious and removal times of the 186 infected population members, with 7 missing removal times. The exposure times are unobserved and inferred along with the 7 missing removal times. The MERS data set does not include direct observations of transmissions or contacts, but there are two sources of indirect observations:

- The assessments of doctors of who infected whom.
- The observed infectious and removal times, which reveal when infected population members were infectious and, in so doing, help narrow down the possible sources of infections.

Both of these sources help inform who infected whom and who was in contact with whom, because an infection implies a contact.

8.2. Model

We use the semiparametric population model introduced in Sections 2 and 4, assuming that $I_i - E_i$ and $R_i - I_i$ are independent $\text{Gamma}(\eta_{E,1}, \eta_{E,2})$ and $\text{Gamma}(\eta_{I,1}, \eta_{I,2})$ random variables, respectively ($i = 1, \dots, 186$). The priors of the epidemiological parameters are given by $\eta_{E,1} \sim \text{Uniform}(4, 8)$, $\eta_{E,2} \sim \text{Uniform}(.75, 3)$, $\eta_{I,1} \sim \text{Uniform}(1.5, 8)$, $\eta_{I,2} \sim \text{Uniform}(2.5, 7.5)$, and $\beta \sim \text{Uniform}(1, 8)$, which cover a range of plausible values (see e.g. the discussions of Groendyke et al. 2011, 2012). The degree parameters $\theta_1, \dots, \theta_{186}$ are assumed to have been generated by a Dirichlet process prior with concentration parameter α and base distribution $N(\mu, \sigma^2)$. To express the uncertainty about the concentration parameter α and the parameters μ and σ^2 of the base distribution, we assume that the hyperparameters α , μ , and $1/\sigma^2$ have $\text{Gamma}(5, 1)$, $N(0, 1)$, and $\text{Gamma}(1, 10)$ hyperpriors, respectively. To facilitate Markov chain Monte Carlo sampling from the posterior distribution, we truncate the Dirichlet process prior at $K = 3$ as described in Section 6.3. We choose $K = 3$ because we expect two subpopulations, one corresponding to potential superspreaders and one corresponding to all other population members, and $K = 3$ is a convenient upper bound on the number of subpopulations. We consider two specifications of the prior probabilities of who infected whom, because the posterior is sensitive to the prior probabilities of who infected whom. The reason is that the MERS data set does not contain direct observations of transmissions or contacts, and the information on transmissions is therefore limited to two sources of indirect observations: the assessments of doctors of who infected whom and the observed infectious and removal times, as discussed in Section 8.1. While the observed infectious and removal times help narrow down

the possible sources of infections, there may be many possible sources of infections left. As a consequence, it is not surprising that the posterior is sensitive to the choice of prior probabilities of who infected whom. We demonstrate the sensitivity of the posterior to the choice of prior in Section 8.4 by using two specifications of prior probabilities:

- (a) If the doctors assessed that population member i infected population member j , we specify $\varphi(i \text{ infected } j) = 1$ and $\varphi(h \text{ infected } j) = 0$ for all infected population members $h \in \{1, \dots, 186\} \setminus \{i, j\}$. If the doctors did not specify who infected j and there are $M_j \in \{1, \dots, 185\}$ infected population members h satisfying $I_h < E_j < R_h$, we specify $\varphi(i \text{ infected } j) = 1/M_j$ for all infected population members i satisfying $I_i < E_j < R_i$ and $\varphi(h \text{ infected } j) = 0$ for all other infected population members h .
- (b) If there are $M_j \in \{1, \dots, 185\}$ infected population members h satisfying $I_h < E_j < R_h$, we specify $\varphi(i \text{ infected } j) = 1/M_j$ for all infected population members i satisfying $I_i < E_j < R_i$ and $\varphi(h \text{ infected } j) = 0$ for all other infected population members h .

Last, but not least, there is no evidence to suggest that the incomplete-data generating process is non-ignorable, therefore we estimate the population model under the assumption that the incomplete-data generating process is ignorable. It is worth noting that the assumption of ignorability is a strong assumption, but it is convenient and defensible unless there is strong evidence to the contrary. We discuss how to deal with non-ignorable incomplete-data generating processes in Section 9.4.

8.3. Computing

We sample from the posterior distribution by using the Bayesian Markov chain Monte Carlo algorithm described in Section 6.4. The Markov chain Monte Carlo algorithm required 302 minutes. Out of the 2,000,000 sample points generated by the Markov chain Monte Carlo algorithm, we discard the first 200,000 sample points as a burn-in and retain every 100th sample point of the remaining 1,800,000 sample points, giving rise to a Markov chain Monte Carlo sample of size 18,000. The resulting Markov chain Monte Carlo sample shows evidence of label-switching. We undo the label-switching as described in Section 6.5. Trace plots of selected parameters can be found in Figure 4. The trace plots do not show signs of non-convergence and suggest that the label-switching has been undone.

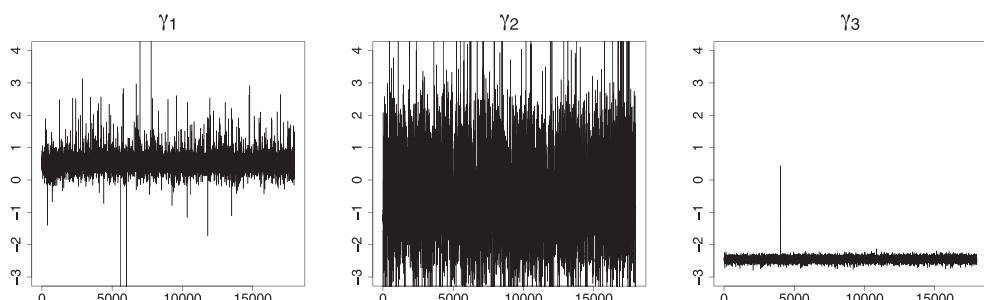


Figure 4. MERS data: Trace plots of the degree parameters γ_1 , γ_2 , and γ_3 .

8.4. Results

The semiparametric population model introduced in Section 4 induces a partition of the population into subpopulations. We therefore start by investigating the local structure of the population, by determining which population members belong to which subpopulations.

Figure 5 shows the posterior classification probabilities of population members along with the transmissions of MERS according to the assessments of doctors. According to doctors, 181 out of the 186 infected population members are suspected of having infected 0, 1, or 2 other population members, while 5 population members are believed to have infected more than 2 other population members: 1 (infected 30 others); 14 (infected 70 others); 15 (infected 6 others); 16 (infected 23 others); and 76 (infected 10 others). In total,

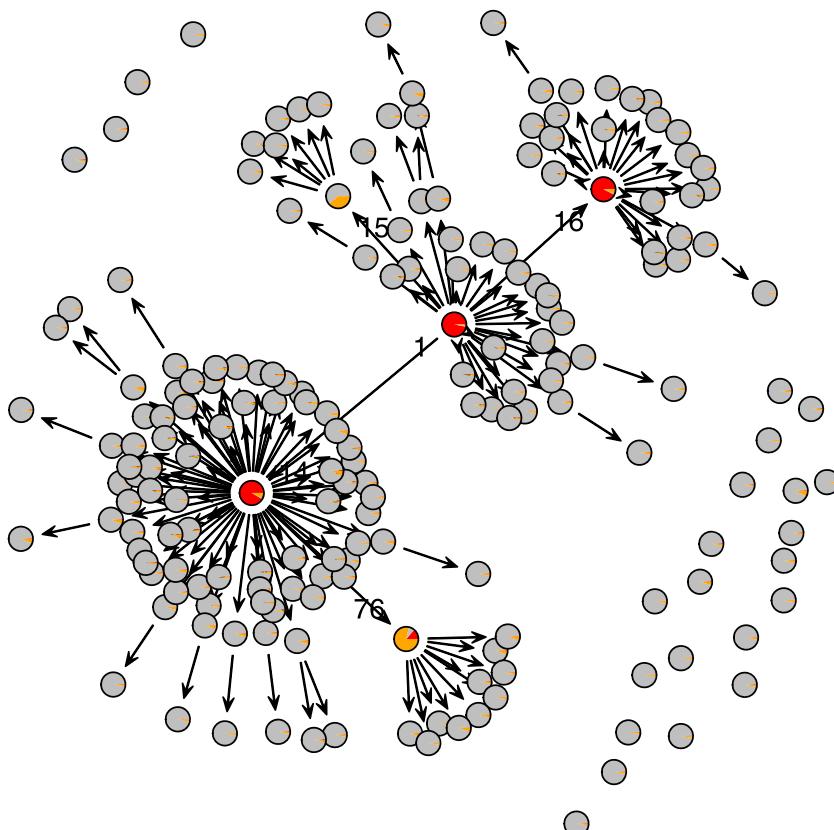


Figure 5. MERS data: Posterior classification probabilities of population members along with transmissions of MERS based on the assessments of doctors. Each population member is represented by a circle. Each circle is divided into coloured slices. The colours of the slices represent subpopulations. The sizes of the slices are proportional to the posterior probabilities of belonging to the corresponding subpopulations. Directed edges between population members indicate transmissions based on the assessments of doctors. The five population members who are believed to have infected 139 out of the 186 infected population members are labelled 1, 14, 15, 16, and 76. The labels of all other population members are omitted.

these 5 population members are believed to have infected 139 out of the 186 infected population members. The posterior classification probabilities in Figure 5 suggest that three subpopulations can be distinguished:

- A subpopulation consisting of population members 1, 14, and 16, which we represent by the colour red in Figure 5. Population member 1 was the first confirmed MERS case in South Korea and is believed to have infected population members 14 and 16. These three population members are suspected of being responsible for the three largest clusters of infections, infecting a total of 123 out of the 186 infected population members, and may therefore be considered to be primary drivers of the MERS outbreak.
- A subpopulation consisting of population members 15 and 76, which we represent by the colour orange in Figure 5. Population member 76 (infected 10 others) belongs with high posterior probability to the orange-coloured subpopulation, whereas there is more uncertainty about the classification of population member 15 (infected 6 others). Both of them are thought to be responsible for clusters of infections – albeit smaller clusters of infections than population members 1, 14, and 16 – and may therefore be considered to be secondary drivers of the MERS outbreak.
- A subpopulation consisting of all other population members, which we represent by the colour grey in Figure 5.

To gain insight into the propensities of population members to form contacts within and between these subpopulations, we inspect the marginal posterior densities of the degree parameters γ_1 , γ_2 , and γ_3 . The degree parameters γ_1 , γ_2 , and γ_3 correspond to the red-, orange-, and grey-coloured subpopulations in Figure 5, respectively. Figure 6 shows Markov chain Monte Carlo approximations of the marginal posterior densities of γ_1 , γ_2 , and γ_3 . The 95% posterior credible intervals of γ_1 , γ_2 , and γ_3 are given by [.140, .959], [-2.575, 1.759], and [-2.587, -2.317], respectively. Since the log odds of the probability of a contact between two population members i and j is $\theta_i + \theta_j = \mathbf{Z}_i^\top \boldsymbol{\gamma} + \mathbf{Z}_j^\top \boldsymbol{\gamma}$, members of the red-coloured subpopulation have a high propensity to be in contact with other population members, in particular other members of the red- and orange-coloured subpopulations, in addition to members of the grey-coloured subpopulation. It is worth noting that the posterior uncertainty – as measured by the lengths of the 95% posterior credible

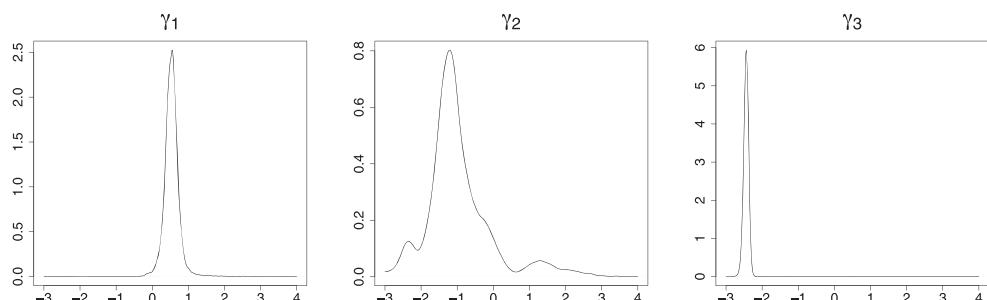


Figure 6. MERS data: Markov chain Monte Carlo approximations of the marginal posterior densities of the degree parameters γ_1 , γ_2 , and γ_3 . The degree parameters γ_1 , γ_2 , and γ_3 correspond to the red-, orange-, and grey-coloured subpopulations in Figure 5, respectively.

intervals – is highest for γ_2 and lowest for γ_3 , which makes sense because no more than 1 or 2 population members belong to the orange-coloured subpopulation with degree parameter γ_2 (with a high posterior probability), whereas 181 population members belong to the grey-coloured subpopulation with degree parameter γ_3 (with a high posterior probability). Taken together, these results suggest that there were three to five potential superspreaders who may have had a great impact on the outcome of the MERS outbreak in South Korea. These observations underscore the importance of detecting potential superspreaders.

We proceed with posterior predictions of the degree distribution of the population contact network. While the population contact network and its degree distribution are unobserved, posterior predictions of the degree distribution can be generated, provided draws from the posterior distribution are available. Probabilistic statements about the degrees of population members based on the posterior distribution are informed by the observed infectious and removal times along with the assessments of doctors of who infected whom. Both of these sources of information help inform who was in contact with whom, because

- the observed infectious and removal times reveal when infected population members were infectious, which helps narrow down the possible sources of infections;
- an infection implies a contact.

The posterior predictions of the degree distribution shown in Figure 7 suggest that the degree distribution is long-tailed: The bulk of population members has no more than 10 contacts, but some population members have as many as 80 contacts. As pointed out before, population members with many contacts can infect many other population members and therefore represent an important public health concern.

Last, but not least, it may be of interest to inspect posterior predictions of the epidemic itself. We focus on the maximum of the ‘epidemic curve,’ that is, the maximum number of infectious population members during the height of the MERS outbreak. The observed number is 129, that is, the number of infectious population members reached 129 at the

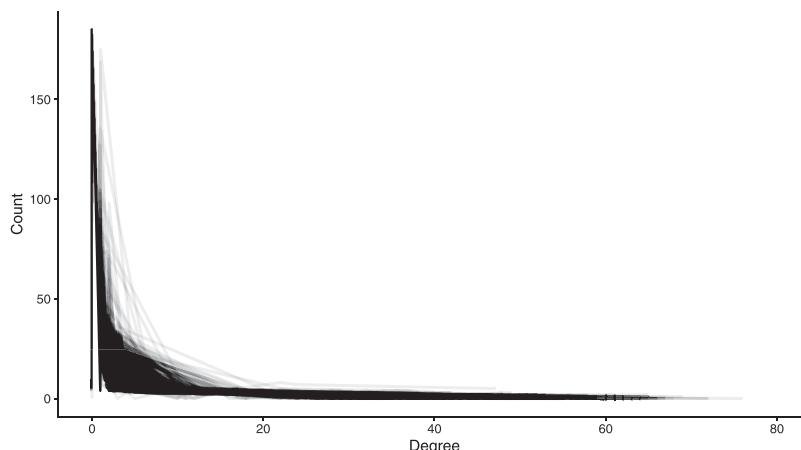
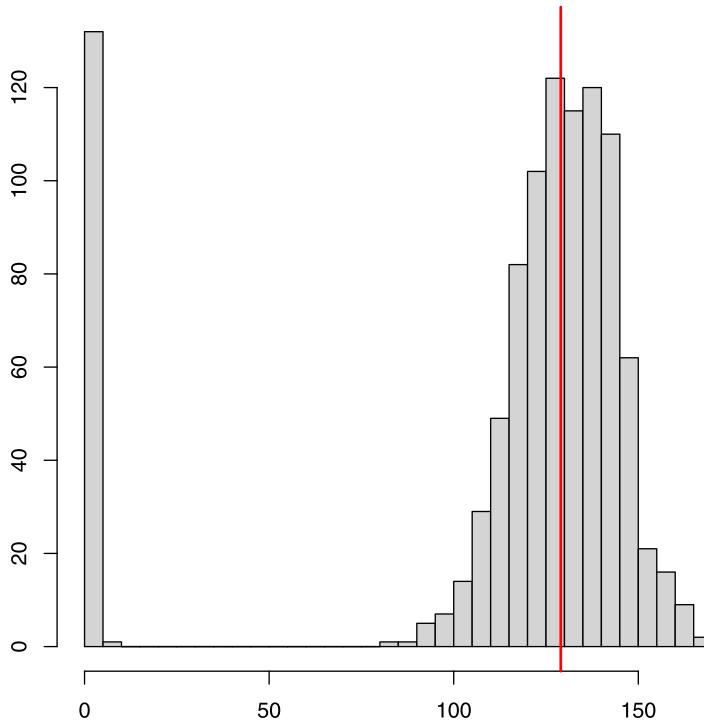


Figure 7. MERS data: Posterior predictions of the degree distribution of the population contact network.



the assessments of doctors and choosing prior specification (b) leads to underpredictions of the maximum of the ‘epidemic curve,’ whereas utilising the assessments of doctors and choosing prior specification (a) leads to predictions that match, on average, the observed maximum of the ‘epidemic curve’ rather well. These observations underscore that

- posterior predictions of the epidemic are sensitive to the choice of prior of who infected whom, at least in the absence of observations of transmissions or contacts;
- data on transmissions and contacts should be collected to reduce the posterior uncertainty about quantities of interest and the sensitivity of posterior predictions of the epidemic to the choice of prior.

9. Open questions and directions for future research

We have introduced a semiparametric population model that can accommodate both short- and long-tailed degree distributions and detect potential superspreaders, in addition to dealing with a wide range of missing and sample data.

Having said that, many open questions remain. Some of them are rooted in the lack of data, while others stem from computational and statistical challenges arising from the lack of data and the complexity of the models. We review a selection of open questions and directions for future research below.

9.1. *What is the population of interest?*

To analyze the MERS outbreak in South Korea, we applied the proposed semiparametric modelling framework to the 186 infected population members. In so doing, we made the implicit assumption that the population of interest consists of those 186 infected population members. There is no denying that such an assumption is unappealing. In fact, the assumption was motivated by convenience rather than substantive considerations, including the challenge of determining the population of interest: e.g. does the population of interest consist of all residents of South Korea, all residents of East Asia, or the whole world?

9.2. *Incomplete data*

As pointed out before, collecting complete data on population contact networks and epidemics is all but impossible. As a consequence, public health officials and researchers face a recurring question in the event of epidemics (whether outbreaks of coronaviruses such as COVID-19, MERS, or SARS, Ebola viruses, or other viruses): Which data to collect? We believe that, to learn the structure of a population contact network and its impact on the spread of an infectious disease, investigators should attempt to collect contact and epidemiological data on all infected population members and collect samples of contacts from non-infected population members by likelihood-ignorable sampling designs. We discuss some of the challenges arising in practice along with possible solutions.

First, while it is challenging to collect data on transmissions, it is advisable to collect data that help reduce the posterior uncertainty about epidemiological parameters. One possible source of data are viral genetic sequence data, among other possible data sources. We refer interested readers to Bouckaert et al. (2019) for a recent review of possible data sources.

Second, it is prudent to sample from population contact networks to reduce the posterior uncertainty about network parameters. The two likelihood-ignorable sampling designs discussed above can be used to do so, though both require a sampling frame. If no sampling frame is available, an alternative would be respondent-driven sampling (e.g. Gile and Handcock 2010; Gile 2011), which is a form of link-tracing without a sampling frame. Location data collected by mobile phones and other electronic sources would be alternative sources of data, but raise data privacy issues (Fienberg and Slavković 2010). In the past decade, substantial progress has been made on data privacy in the statistical literature: see, e.g. the work of Karwa and Slavković (2016) on data privacy in scenarios where network data are generated by the β -model (albeit without Dirichlet process priors and without epidemics). Studying data privacy for epidemics would be an important direction for future research.

9.3. Computational challenges arising from incomplete data

The lack of data has computational implications. Statistical algorithms for likelihood-based statistical inference (e.g. EM algorithms, Dempster, Laird, and Rubin 1977) have to integrate over the unobserved data, hence the computing time tends to increase with the amount of missing data. How to develop scalable statistical algorithms, with statistical guarantees, is an open question. One idea would be to develop a two-step estimation algorithm, first estimating the network parameters and then estimating the epidemiological parameters, leveraging computational advances in the statistical analysis of network data (e.g. Raftery et al. 2012; Salter-Townshend and Murphy 2013) and epidemiological data (Bouckaert et al. 2019). A two-step estimation algorithm may require data on contacts, however, because without data on contacts the posterior correlations of the network parameters and the rate of infection can be high (Groendyke et al. 2011), in which case two-step estimation algorithms may not work well. A related problem is how to update posteriors in reasonable time as more data on infections and contacts come in.

9.4. Non-ignorable incomplete-data generating processes

We have considered here ignorable incomplete-data generating processes. If the incomplete-data generating process is non-ignorable, then either the incomplete-data generating process must be modelled or it must be demonstrated that Bayesian inference for the parameters of the population model is insensitive to the incomplete-data generating process. Both approaches require insight into the incomplete-data generating process and may require additional model assumptions, some of which may be untestable.

9.5. Population models capturing additional network features

We have focussed here on the degrees of population members as network features, but there are many other important network features: e.g. if population contact networks exhibit closure (e.g. transitive closure, Wasserman and Faust 1994; Kolaczyk 2009), then infectious diseases may spread rapidly within subpopulations but may spread only slowly through the whole population, which has potential policy implications. There are two

broad approaches to capturing closure in population contact networks: latent space models (Hoff, Raftery, and Handcock 2002; Handcock, Raftery, and Tantrum 2007; Smith, Asta, and Calder 2019) and related latent variable models (e.g. Salter-Townshend, White, Gollini, and Murphy 2012; Rastelli, Friel, and Raftery 2016; Fosdick and Hoff 2015; Hoff 2021); and exponential-family models of random graphs (Frank and Strauss 1986; Snijders, Pattison, Robins, and Handcock 2006; Lusher, Koskinen, and Robins 2013). Both classes of models can accommodate the degree terms we have used (see, e.g. Krivitsky, Handcock, Raftery, and Hoff 2009; Thiemichen, Friel, Caimo, and Kauermann 2016) along with additional terms that reward closure in population contact networks. However, both of them come at costs in terms of computing time (Bhamidi, Bresler, and Sly 2011; Chatterjee and Diaconis 2013), although the computing time depends on the class of models under consideration (see, e.g. Karwa, Petrović, and Bajić 2016). In addition, Welch (2011) pointed out that closure in population contact networks may not be detectable unless data on the population contact network are collected. This, too, underscores the importance of collecting network data.

9.6. *Time-evolving population contact networks*

We have assumed that the population contact network is time-invariant, motivated by the lack of data on the population contact network and the desire to keep the model as simple and parsimonious as possible. In practice, the population contact network may evolve over time, because population members may create or discontinue contacts and because authorities may enforce social distancing measures. As a consequence, it would be natural to allow the population contact network to change over time. Extensions to time-evolving population contact networks could be based on temporal stochastic block and latent space models (e.g. Fu, Song, and Xing 2009; Sewell and Chen 2015, 2016; Sewell, Chen, Bernhard, and Sulkin 2016); temporal exponential-family random graph models (Robins and Pattison 2001; Hanneke, Fu, and Xing 2010; Ouzienko, Guo, and Obradovic 2011; Krivitsky and Handcock 2014); continuous-time Markov processes (Snijders 2001); relational event models (Butts 2008); and other models (e.g. Katz and Proctor 1959; Durante and Dunson 2014; Sewell 2017). An interesting approach to time-evolving population contact networks was recently introduced by Bu et al. (2021), albeit without the ability to accommodate long-tailed degree distributions and detect potential superspreaders.

Acknowledgements

We are indebted to two anonymous referees, whose constructive comments and suggestions led to numerous improvements of the manuscript. We are grateful to Chris Groendyke, David Welch, and David R. Hunter for valuable discussions, comments, and suggestions.

Data availability

We retrieved the data from the website <http://english.mw.go.kr> of the South Korean Ministry of Health and Welfare on September 22, 2015. The website has been removed since, but the data can be obtained from the authors.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the National Institutes of Health under Grant NIH R01 GM083603 (MS, RB); the Office of Naval Research under Grant ONR N00014-08-1-1015 (MS); the National Science Foundation under Grants NSF DMS-1513644 and NSF DMS-1812119 (MS, SB); and the Department of Defense under Grant ARO W911NF-21-1-0335 (MS).

ORCID

Michael Schweinberger  <http://orcid.org/0000-0003-3649-5386>

References

Albert, R., and Barabàsi, A.L. (2002), ‘Statistical Mechanics of Complex Networks’, *Review of Modern Physics*, 74, 47–97.

Amini, A.A., Chen, A., Bickel, P.J., and Levina, E. (2013), ‘Pseudo-likelihood Methods for Community Detection in Large Sparse Networks’, *The Annals of Statistics*, 41, 2097–2122.

Andersson, H., and Britton, T. (2000), *Stochastic Epidemic Models and Their Statistical Analysis*, New York: Springer-Verlag.

Ball, F. (2021), ‘Central Limit Theorems for SIR Epidemics and Percolation on Configuration Model Random Graphs’, *The Annals of Applied Probability*, 1, 71.

Barabási, A.L., and Albert, R (1999), ‘Emergence of Scaling in Random Networks’, *Science*, 286, 509–512.

Barbour, A.D., and Reinert, G. (2013), ‘Approximating the Epidemic Curve’, *Electronic Journal of Probability*, 18, 1–30.

Bhamidi, S., Bresler, G., and Sly, A. (2011), ‘Mixing Time of Exponential Random Graphs’, *The Annals of Applied Probability*, 21, 2146–2170.

Bollobás, B., Riordan, O., Spencer, J., and Tusnády, G. (2001), ‘The Degree Sequence of a Scale-free Random Graph Process’, *Random Structures & Algorithms*, 18, 279–290.

Bouckaert, R., Vaughan, T.G., Barido-Sottani, J., Duchne, S., Fourment, M., Gavryushkina, A., Heled, J., Jones, G., Kuhnert, D., De Maio, N., Matschiner, M., Mendes, F.K., Moller, N.F., Ogilvie, H.A., Popinga, A., Rambaut, A., Rasmussen, D., Siveroni, I., Suchard, M.A., Wu, C.-H., Xie, D., Zhang, C., Stadler, T., and Drummond, A.J. (2019), ‘BEAST 2.5: An Advanced Software Platform for Bayesian Evolutionary Analysis’, *PLOS Computational Biology*, 15, 1–28.

Britton, T. (1998), ‘Estimation in Multitype Epidemics’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60, 663–679.

Britton, T. (2001), ‘Epidemics in Heterogeneous Communities: Estimation of R0 and Secure Vaccination Coverage’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63, 705–715.

Britton, T., Lindholm, M., and Turova, T. (2011), ‘A Dynamic Network in a Dynamic Population: Asymptotic Properties’, *Journal of Applied Probability*, 48, 1163–1178.

Britton, T., and O’Neill, P.D (2002), ‘Statistical Inference for Stochastic Epidemics in Populations with Network Structure’, *Scandinavian Journal of Statistics*, 29, 375–390.

Britton, T., and Pardoux, E. (eds.) (2019), *Stochastic Epidemic Models with Inference*, Cham: Springer.

Bu, F., Aiello, A.E., Xu, J., and Volfovsky, A. (2021), ‘Likelihood-based Inference for Partially Observed Epidemics on Dynamic Networks’, *Journal of the American Statistical Association*, accepted.

Butts, C.T. (2008), ‘A Relational Event Framework for Social Action’, *Sociological Methodology*, 38, 155–200.

Butts, C.T. (2019), 'A Dynamic Process Interpretation of the Sparse ERGM Reference Model', *Journal of Mathematical Sociology*, 43, 40–57.

Chatterjee, S., and Diaconis, P. (2013), 'Estimating and Understanding Exponential Random Graph Models', *The Annals of Statistics*, 41, 2428–2461.

Chatterjee, S., Diaconis, P., and Sly, A. (2011), 'Random Graphs with a Given Degree Sequence', *The Annals of Applied Probability*, 21, 1400–1435.

Chen, M., Kato, K., and Leng, C. (2021), 'Analysis of Networks Via the Sparse β -model', *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, arXiv:1908.03152.

Crane, H. (2018), *Probabilistic Foundations of Statistical Network Analysis*, Boca Raton, FL: Chapman and Hall/CRC.

Danon, L., Ford, A.P., House, T., Jewell, C.P., Keeling, M.J., Roberts, G.O., Ross, J.V., and Vernon, M.C. (2011), 'Networks and the Epidemiology of Infectious Disease', *Interdisciplinary Perspectives on Infectious Diseases*, 2011, 1–28.

Dawid, A.P., and Dickey, J.M. (1977), 'Likelihood and Bayesian Inference From Selectively Reported Data', *Journal of the American Statistical Association*, 72, 845–850.

Dempster, A.P., Laird, N.M., and Rubin, R.B. (1977), 'Maximum Likelihood From Incomplete Data Via the EM-algorithm', *Journal of the Royal Statistical Society, Series B*, 39, 1–38.

Dunbar, R.I.M. (1992), 'Neocortex Size As a Constraint on Group Size in Primates', *Journal of Human Evolution*, 22, 469–493.

Durante, D., and Dunson, D.B. (2014), 'Nonparametric Bayes Dynamic Modelling of Relational Data', *Biometrika*, 101, 125–138.

Ferguson, T. (1973), 'A Bayesian Analysis of Some Nonparametric Problems', *The Annals of Statistics*, 1, 209–230.

Fienberg, S.E., and Slavković, A. (2010), 'Data Privacy and Confidentiality', in *International Encyclopedia of Statistical Science*, ed. M. Lovric, Berlin: Heidelberg.

Fosdick, B.K., and Hoff, P.D. (2015), 'Testing and Modeling Dependencies Between a Network and Nodal Attributes', *Journal of the American Statistical Association*, 110, 1047–1056.

Frank, O. (1988), 'Random Sampling and Social Networks: A Survey of Various Approaches', *Mathématiques, Informatique et Sciences Humaines*, 26, 19–33.

Frank, O., and Strauss, D. (1986), 'Markov Graphs', *Journal of the American Statistical Association*, 81, 832–842.

Fu, W., Song, L., and Xing, E.P. (2009), 'Dynamic Mixed Membership Blockmodel for Evolving Networks', in *Proceedings of the 26th Annual International Conference on Machine Learning*.

Gile, K. (2011), 'Improved Inference for Respondent-driven Sampling Data with Application to HIV Prevalence Estimation', *Journal of the American Statistical Association*, 106, 135–146.

Gile, K., and Handcock, M.S. (2006), 'Model-based Assessment of the Impact of Missing Data on Inference for Networks', Technical Report, Center for Statistics and the Social Sciences, University of Washington, Seattle, available at <https://www.csss.washington.edu/Papers/wp66.pdf>.

Gile, K., and Handcock, M.H. (2010), 'Respondent-driven Sampling: An Assessment of Current Methodology', *Sociological Methodology*, 40, 285–327.

Goodman, L.A. (1961), 'Snowball Sampling', *Annals of Mathematical Statistics*, 32, 148–170.

Groendyke, C., and Welch, D. (2018), 'epinet: An R Package to Analyze Epidemics Spread Across Contact Networks', *Journal of Statistical Software*, 83, 1–22.

Groendyke, C., Welch, D., and Hunter, D.R. (2011), 'Bayesian Inference for Contact Networks Given Epidemic Data', *Scandinavian Journal of Statistics*, 38, 600–616.

Groendyke, C., Welch, D., and Hunter, D.R. (2012), 'A Network-based Analysis of the 1861 Hagelloch Measles Data', *Biometrics*, 68, 755–765.

Handcock, M.S., and Gile, K. (2010), 'Modeling Social Networks From Sampled Data', *The Annals of Applied Statistics*, 4, 5–25.

Handcock, M.S., and Gile, K. (2017), 'Analysis of Networks with Missing Data with Application to the National Longitudinal Study of Adolescent Health', *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 66, 501–519.

Handcock, M.S., Raftery, A.E., and Tantrum, J.M. (2007), 'Model-based Clustering for Social Networks', *Journal of the Royal Statistical Society, Series A (with discussion)*, 170, 301–354.

Hanneke, S., Fu, W., and Xing, E.P. (2010), 'Discrete Temporal Models of Social Networks', *Electronic Journal of Statistics*, 4, 585–605.

Hartley, H., and Sielken Jr., R. (1975), 'A "super-population Viewpoint" for Finite Population Sampling', *Biometrics*, 31, 411–422.

Heckathorn, D.D. (1997), 'Respondent-driven Sampling: A New Approach to the Study of Hidden Populations', *Social Problems*, 44, 174–199.

Hoff, P.D. (2021), 'Additive and Multiplicative Effects Network Models', *Statistical Science*, 36, 34–50.

Hoff, P.D., Raftery, A.E., and Handcock, M.S. (2002), 'Latent Space Approaches to Social Network Analysis', *Journal of the American Statistical Association*, 97, 1090–1098.

Holland, P.W., and Leinhardt, S. (1981), 'An Exponential Family of Probability Distributions for Directed Graphs', *Journal of the American Statistical Association*, 76, 33–65.

Ishwaran, H., and James, L.F. (2001), 'Gibbs Sampling Methods for Stick-breaking Priors', *Journal of the American Statistical Association*, 96, 161–173.

Jones, J.H., and Handcock, M.S. (2003a), 'An Assessment of Preferential Attachment As a Mechanism for Human Sexual Network Formation', *Proceedings of the Royal Society Series B*, 270, 1123–1128.

Jones, J.H., and Handcock, M.S. (2003b), 'Social Networks: Sexual Contacts and Epidemic Thresholds', *Nature*, 423, 605–606.

Jones, J.H.H., and Handcock, M.S. (2004), 'Likelihood-Based Inference for Stochastic Models of Sexual Network Formation', *Population Biology*, 65, 413–422.

Karwa, V., Petrović, S., and Bajić, D. (2016), 'DERGMs: Degeneracy-Restricted Exponential Random Graph Models', *arXiv preprint arXiv:1612.03054*.

Karwa, V., and Slavković, A.B. (2016), 'Inference Using Noisy Degrees: Differentially Private β -model and Synthetic Graphs', *The Annals of Statistics*, 44, 87–112.

Katz, L., and Proctor, C.H. (1959), 'The Configuration of Interpersonal Relations in a Group As a Time-dependent Stochastic Process', *Psychometrika*, 24, 317–327.

Keeling, M.J., and Eames, K.T.D. (2005), 'Networks and Epidemic Models', *Journal of the Royal Society Interface*, 2, 295–307.

Ki, M. (2015), '2015 MERS Outbreak in Korea: Hospital-to-hospital Transmission', *Epidemiology and Health*, 37, 1–4.

Kolaczyk, E.D. (2009), *Statistical Analysis of Network Data: Methods and Models*, New York: Springer-Verlag.

Koskinen, J.H., Robins, G.L., and Pattison, P.E. (2010), 'Analysing Exponential Random Graph (p-star) Models with Missing Data Using Bayesian Data Augmentation', *Statistical Methodology*, 7, 366–384.

Krivitsky, P.N., and Handcock, M.S. (2014), 'A Separable Model for Dynamic Networks', *Journal of the Royal Statistical Society B*, 76, 29–46.

Krivitsky, P.N., Handcock, M.S., and Morris, M. (2011), 'Adjusting for Network Size and Composition Effects in Exponential-family Random Graph Models', *Statistical Methodology*, 8, 319–339.

Krivitsky, P.N., Handcock, M.S., Raftery, A.E., and Hoff, P.D. (2009), 'Representing Degree Distributions, Clustering, and Homophily in Social Networks With Latent Cluster Random Effects Models', *Social Networks*, 31, 204–213.

Krivitsky, P.N., and Kolaczyk, E.D. (2015), 'On the Question of Effective Sample Size in Network Modeling: An Asymptotic Inquiry', *Statistical Science*, 30, 184–198.

Krivitsky, P.N., and Morris, M. (2017), 'Inference for Social Network Models From Egocentrically-sampled Data, with Application to Understanding Persistent Racial Disparities in HIV Prevalence in the US', *Annals of Applied Statistics*, 11, 427–455.

Kurant, M., Markopoulou, A., and Thiran, P. (2011), 'Towards Unbiased BFS Sampling', *IEEE Journal on Selected Areas in Communications*, 29, 1799–1809.

Laumann, E.O., Gagnon, J.H., Michael, T., and Michaels, S. (1994), *The Social Organization of Sexuality: Sexual Practices in the United States*, Chicago: University of Chicago Press.

Liu, J.S. (2008), *Monte Carlo Strategies in Scientific Computing*, New York: Springer-Verlag.

Lovász, L. (2012), *Large Networks and Graph Limits*, Providence: American Mathematical Society.

Lusher, D., Koskinen, J., and Robins, G. (2013), *Exponential Random Graph Models for Social Networks*, Cambridge, UK: Cambridge University Press.

Ouzienko, V., Guo, Y., and Obradovic, Z. (2011), 'A Decoupled Exponential Random Graph Model for Prediction of Structure and Attributes in Temporal Social Networks', *Statistical Analysis and Data Mining*, 4, 470–486.

Pang, G., and Pardoux, E. (2020), 'Functional Limit Theorems for Non-Markovian Epidemic Models', *arXiv preprint arXiv:2003.03249*, 1–45.

Raftery, A.E., Niu, X., Hoff, P.D., and Yeung, K.Y. (2012), 'Fast Inference for the Latent Space Network Model Using a Case-control Approximate Likelihood', *Journal of Computational and Graphical Statistics*, 21, 901–919.

Rastelli, R., Friel, N., and Raftery, A.E. (2016), 'Properties of Latent Variable Network Models', *Network Science*, 4, 407–432.

Reinert, G. (1995), 'The Asymptotic Evolution of the General Stochastic Epidemic', *The Annals of Applied Probability*, 5, 1061–1086.

Rinaldo, A., Petrović, S., and Fienberg, S.E. (2013), 'Maximum Likelihood Estimation in the β -model', *The Annals of Statistics*, 41, 1085–1110.

Robins, G., and Pattison, P. (2001), 'Random Graph Models for Temporal Processes in Social Networks', *Journal of Mathematical Sociology*, 25, 5–41.

Rohe, K., Chatterjee, S., and Yu, B. (2011), 'Spectral Clustering and the High-dimensional Stochastic Block Model', *The Annals of Statistics*, 39, 1878–1915.

Rubin, D.B. (1976), 'Inference and Missing Data', *Biometrika*, 63, 581–592.

Salganik, M.J., and Heckathorn, D.D. (2004), 'Sampling and Estimation in Hidden Populations Using Respondent-driven Sampling', *Sociological Methodology*, 34, 193–239.

Salter-Townshend, M., and Murphy, T.B. (2013), 'Variational Bayesian Inference for the Latent Position Cluster Model for Network Data', *Computational Statistics and Data Analysis*, 57, 661–671.

Salter-Townshend, M., White, A., Gollini, I., and Murphy, T.B. (2012), 'Review of Statistical Network Analysis: Models, Algorithms, and Software', *Statistical Analysis and Data Mining*, 5, 243–264.

Schweinberger, M., and Handcock, M.S. (2015), 'Local Dependence in Random Graph Models: Characterization, Properties and Statistical Inference', *Journal of the Royal Statistical Society, Series B*, 77, 647–676.

Schweinberger, M., Krivitsky, P.N., Butts, C.T., and Stewart, J.R. (2020), 'Exponential-family Models of Random Graphs: Inference in Finite, Super, and Infinite Population Scenarios', *Statistical Science*, 35, 627–662.

Schweinberger, M., and Luna, P. (2018), 'HERGM: Hierarchical Exponential-family Random Graph Models', *Journal of Statistical Software*, 85, 1–39.

Sethuraman, J. (1994), 'A Constructive Definition of Dirichlet Priors', *Statistica Sinica*, 4, 639–650.

Sewell, D.K. (2017), 'Network Autocorrelation Models with Egocentric Data', *Social Networks*, 49, 113–123.

Sewell, D.K., and Chen, Y. (2015), 'Latent Space Models for Dynamic Networks', *Journal of the American Statistical Association*, 110, 1646–1657.

Sewell, D.K., and Chen, Y. (2016), 'Latent Space Approaches to Community Detection in Dynamic Networks', *Bayesian Analysis*, 12, 351–377.

Sewell, D.K., Chen, Y., Bernhard, W., and Sulkin, T. (2016), 'Model-based Longitudinal Clustering with Varying Cluster Assignments', *Statistica Sinica*, 26, 205–233.

Smith, A.L., Asta, D.M., and Calder, C.A. (2019), 'The Geometry of Continuous Latent Space Models for Network Data', *Statistical Science*, 34, 428–453.

Snijders, T.A.B. (2001), 'The Statistical Evaluation of Social Network Dynamics', in *Sociological Methodology*, eds. M. Sobel and M. Becker, Boston and London: Basil Blackwell, pp. 361–395.

Snijders, T.A.B., Pattison, P.E., Robins, G.L., and Handcock, M.S. (2006), 'New Specifications for Exponential Random Graph Models', *Sociological Methodology*, 36, 99–153.

Stephens, M. (2000), 'Dealing with Label-switching in Mixture Models', *Journal of the Royal Statistical Society, Series B*, 62, 795–809.

Sundberg, R. (2019), *Statistical Modelling by Exponential Families*, Cambridge, UK: Cambridge University Press.

Teh, Y.W. (2010), 'Dirichlet Processes', in *Encyclopedia of Machine Learning*, eds. C. Sammut and G. I. Webb, New York: Springer-Verlag.

Thiemichen, S., Friel, N., Caimo, A., and Kauermann, G. (2016), 'Bayesian Exponential Random Graph Models with Nodal Random Effects', *Social Networks*, 46, 11–28.

Thompson, S. (2012), *Sampling* (3rd ed.), Hoboken, NJ: John Wiley & Sons.

Thompson, S., and Frank, O. (2000), 'Model-based Estimation with Link-tracing Sampling Designs', *Survey Methodology*, 26, 87–98.

Tierney, L. (1994), 'Markov Chains for Exploring Posterior Distributions', *The Annals of Statistics*, 22, 1701–1728.

Vu, D.Q., Hunter, D.R., and Schweinberger, M. (2013), 'Model-based Clustering of Large Networks', *The Annals of Applied Statistics*, 7, 1010–1039.

Wainwright, M.J. (2019), *High-Dimensional Statistics. A Non-Asymptotic Viewpoint*, Cambridge, UK: Cambridge University Press.

Wasserman, S., and Faust, K. (1994), *Social Network Analysis: Methods and Applications*, Cambridge: Cambridge University Press.

Welch, D. (2011), 'Is Network Clustering Detectable in Transmission Trees?', *Viruses*, 3, 659–676.

Welch, D., Bansal, S., and Hunter, D.R. (2011), 'Statistical Inference to Advance Network Models in Epidemiology', *Epidemics*, 3, 38–45.

Willinger, W., Alderson, D., and Doyle, J.C. (2009), 'Mathematics and the Internet: A Source of Enormous Confusion and Great Potential', *Notices of the American Mathematical Society*, 56, 586–599.