



Heriot-Watt University
Research Gateway

Persistence time of SIS infections in heterogeneous populations and networks

Citation for published version:

Clancy, D 2018, 'Persistence time of SIS infections in heterogeneous populations and networks', *Journal of Mathematical Biology*, vol. 77, no. 3, pp. 545-570. <https://doi.org/10.1007/s00285-018-1222-1>

Digital Object Identifier (DOI):

[10.1007/s00285-018-1222-1](https://doi.org/10.1007/s00285-018-1222-1)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Journal of Mathematical Biology

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Persistence time of SIS infections in heterogeneous populations and networks

Damian Clancy¹

Received: 11 July 2017 / Revised: 5 February 2018
© The Author(s) 2018. This article is an open access publication

Abstract For a susceptible–infectious–susceptible infection model in a heterogeneous population, we present simple formulae giving the leading-order asymptotic (large population) behaviour of the mean persistence time, from an endemic state to extinction of infection. Our model may be interpreted as describing an infection spreading through either (1) a population with heterogeneity in individuals’ susceptibility and/or infectiousness; or (2) a heterogeneous directed network. Using our asymptotic formulae, we show that such heterogeneity can only reduce (to leading order) the mean persistence time compared to a corresponding homogeneous population, and that the greater the degree of heterogeneity, the more quickly infection will die out.

Keywords Stochastic epidemic models · Large deviations · Endemic fade-out · Directed configuration model · Superspreaders

Mathematics Subject Classification 92D30 · 60J28

1 Introduction

In modelling endemic infections, a quantity of particular interest is the persistence time until infection dies out from the population. For discrete state-space Markov chain models, the expected persistence time for an infection that has become endemic in the population (i.e. starting from quasi-stationarity) may be found as an eigenvalue of the transition rate matrix (van Doorn and Pollett 1993). However, for large popula-

✉ Damian Clancy
d.clancy@hw.ac.uk

¹ Department of Actuarial Mathematics and Statistics, Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK

tions and for more complicated models, numerical computation of this exact solution can be very time-consuming, and may also suffer from numerical instability. Moreover, it is not straightforward to use this eigenvalue characterization to investigate, for instance, the effect of population heterogeneities upon the expected persistence time. Approximation methods are therefore required. For a number of infection models it has been shown (Andersson and Djehiche 1998; Ball et al. 2016; van Herwaarden and Grasman 1995) that, denoting by N the typical size of the population, the expected time from endemicity to extinction, τ , is asymptotically given by an expression of the form

$$\tau \sim \frac{C}{\sqrt{N}} \exp(AN) \quad (1)$$

where the values of A , C depend upon parameters of the model, but not upon N . It is assumed here that the process is super-critical, so that long-term endemicity is possible.

For the classic susceptible–infectious–susceptible (SIS) model of Weiss and Dishon (1971), Andersson and Djehiche (1998) found simple explicit expressions for both A and C in terms of the basic reproduction number R_0 (the expected number of secondary cases caused by a typical primary case in an otherwise susceptible population), under the assumption of super-criticality (that is, $R_0 > 1$); specifically, $A = (1/R_0) - 1 + \ln R_0$ and $C = R_0 \sqrt{2\pi} / (R_0 - 1)^2$, assuming time is scaled such that individual infectious periods are of mean 1. This was extended by Ball et al. (2016) to allow for a general infectious period distribution in place of the exponential distribution assumed by Andersson and Djehiche (1998); they showed that leading-order behaviour is unchanged, so that $A = (1/R_0) - 1 + \ln R_0$ as before, while the value of C depends upon the infectious period distribution, and may be straightforwardly evaluated provided this distribution is known. Pre-dating the above work, van Herwaarden and Grasman (1995) showed that relationship (1) holds true for a particular susceptible–infectious–removed (SIR) infection model. In this case, however, evaluation of the constant A requires numerical solution of a system of ordinary differential equations, while no method for evaluating C is given.

The system of ordinary differential equations used in van Herwaarden and Grasman (1995) to evaluate A may be regarded as the equations of motion corresponding to a particular Hamiltonian system. More recently, a number of authors (Assaf and Meerson 2010, 2017; Dykman et al. 1994; Elgart and Kamenev 2004; Kamenev and Meerson 2008; Lindley et al. 2014) have applied this Hamiltonian approach to a range of infection models to derive results of the form

$$\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} = A. \quad (2)$$

Equation (2) is not as precise as relationship (1), but does at least give the leading-order behaviour of τ in the large population ($N \rightarrow \infty$) limit. Evaluation of A generally requires numerical solution of the equations of motion, and consequently much of the research effort has focused upon developing efficient numerical procedures.

We shall apply the Hamiltonian approach to approximate the expected persistence time from endemicity, τ , for an SIS model incorporating heterogeneity in individual

infectivities and susceptibilities. Such heterogeneity is a common feature of real-world infections. For instance, for a number of infections (eg SARS) it has been hypothesised that there exists a subgroup of ‘super-spreaders’ within the population, being individuals of higher infectivity than the rest. Heterogeneous susceptibilities may arise, for instance, through individuals having differing histories of prior exposure to infection or vaccination. Alternatively, our model may be interpreted as a model for infection spreading on an uncorrelated (that is, with no correlations between degrees of neighbouring individuals) directed network (Dorogovtsev et al. 2008).

In contrast to almost all previous work, we are able to find an explicit formula for the constant A in Eq. (2), at least provided the heterogeneity is in either infectivity, or susceptibility, but not both. As well as being much quicker and easier to evaluate than the solution to a (typically high-dimensional) system of ordinary differential equations, a further advantage of such an explicit formula is that it may be used to establish qualitative results about the effects of model assumptions. Specifically, we investigate the effect of increasing heterogeneity upon the persistence time of infection in the population.

The remainder of the paper is structured as follows. In Sect. 2, we define precisely our heterogeneous population SIS model, and describe how it may be interpreted as approximating a directed network model. In Sect. 3 we recall some general theory that will be required in the sequel. Our main result, Theorem 1, is derived in Sect. 4, establishing explicit asymptotic formulae for $\ln \tau$ in the large-population limit, provided that heterogeneity is in either infectivity or susceptibility, but not both. Using these explicit formulae, we go on in Sect. 5 to demonstrate that the greater the level of heterogeneity in either infectivity or susceptibility, the more rapidly extinction of infection will occur (on average, to leading order in a large population). In the case that both heterogeneities are present simultaneously, we present an approximate formula for $\lim_{N \rightarrow \infty} (\ln \tau)/N$, valid provided R_0 is only slightly greater than 1, together with numerical work (Fig. 2) indicating that mean persistence time is again maximised in the homogeneous-population case. In Sect. 6 we demonstrate that if heterogeneity is in susceptibilities, our asymptotic formula for $\ln \tau$, and hence also our conclusion that greater heterogeneity reduces mean persistence time, remain valid for more general infectious period distributions than the classic exponential distribution. We present numerical evidence (Fig. 3) suggesting that this is also true when instead heterogeneity is in infectivities. Finally, in Sect. 7, we discuss the directed network interpretation of our results (Theorem 4) and suggest some directions for further work.

2 The SIS infection model in a heterogeneous population or directed network

We first formulate our model in terms of a population divided into a fixed number of groups, and then describe how the same model may be interpreted as modelling an infection spreading on a directed network.

Consider a closed population of N individuals divided into k groups, with group i ($i = 1, 2, \dots, k$) consisting of N_i individuals. Denote by $f_i = N_i/N$ the proportion of the population belonging to group i , so that $\sum_i f_i = 1$. When a group i individ-

Table 1 Transition rates for the k -group SIS model

Event	State transition	Transition rate
Infection in group j	$I_j \rightarrow I_j + 1$	$\frac{\beta}{N} \left(\sum_{m=1}^k \lambda_m I_m \right) \mu_j (N_j - I_j)$
Recovery in group j	$I_j \rightarrow I_j - 1$	γI_j

ual becomes infected, it remains so for a time distributed as an exponential random variable with mean $1/\gamma$ (assumed for simplicity to be the same for each group). During this infectious period, the group i infective makes contacts with each individual in each group $j = 1, 2, \dots, k$ at the points of a Poisson process of rate $\beta \lambda_i \mu_j / N$, where β is some overall measure of infectiousness, λ_i represents the infectivity of group i individuals, and μ_j represents the susceptibility of group j individuals. (The assumption that the group i to group j infection rate factorises in this way is sometimes referred to as ‘separable mixing’.) Without loss of generality, we scale the λ_i , μ_j values so that $\sum_i \lambda_i f_i = \sum_j \mu_j f_j = 1$. These Poisson processes and infectious periods are all mutually independent. If a contacted individual is susceptible, then it becomes infected (and infectious); if the contacted individual is already infected then the contact has no effect. Denoting by $I_j(t)$ the number of infected individuals in group j at time t , then the process $\{\mathbf{I}(t) = (I_1(t), I_2(t), \dots, I_k(t)) : t \geq 0\}$ is a continuous-time Markov chain with transition rates given in Table 1, and the number of susceptible individuals in group j at time $t \geq 0$ is $S_j(t) = N_j - I_j(t)$. We will assume throughout that $\beta, \gamma > 0$, and that $f_i, \lambda_i, \mu_i > 0$ for all i . Note that our model is a special case of the model of Clancy and Pearce (2013), although we use slightly different notation here. The basic reproduction number R_0 is given by the dominant eigenvalue of the matrix M with entries $m_{ij} = \beta \lambda_i \mu_j f_j$, so that

$$R_0 = \frac{\beta}{\gamma} \sum_{i=1}^k \lambda_i \mu_i f_i.$$

We now describe how the above model may be interpreted as describing infection spreading through a network. Each of the N individuals in the population is assigned an in-degree d_{in} and out-degree d_{out} according to some joint probability mass function $p(d_{\text{in}}, d_{\text{out}})$ on \mathbb{Z}_+^2 . These degrees are assigned independently to distinct individuals, but the in and out degree need not be independent for a single individual. To each individual we attach ‘stubs’, or half-edges, with d_{in} stubs pointing inwards and d_{out} stubs pointing outwards. Inward-pointing stubs are then paired with outward-pointing stubs throughout the population, to create links between individuals. In order that this process can produce a valid network, with no left over half-edges, we clearly require that $E[d_{\text{in}}] = E[d_{\text{out}}]$. We do not concern ourselves with the precise mechanism by which stubs are paired off [see Britton et al. (2007) and Chen and Olvera-Cravioto (2013) for relevant discussion]; rather, we shall simply assume that the resulting network is uncorrelated, so that the so-called ‘annealed’ network approximation is valid for an ensemble of such networks. This is a mean-field approximation for heterogeneous networks, and may be described as follows (see Dorogovtsev et al. 2008 for more comprehensive discussion).

Denote by κ the rate at which infection transmits along each link from an infectious individual to a susceptible individual. Suppose for simplicity that there are a finite number k of (d_{in}, d_{out}) pairs having non-zero probability, and define a bijective function $c(d_{in}, d_{out}) : \mathbb{Z}_+^2 \rightarrow \{1, 2, \dots, k\}$ that assigns a unique number to each of the possible (d_{in}, d_{out}) pairs. We say an individual is of ‘group j ’ if they have degrees $(d_{in}, d_{out}) = c^{-1}(j)$, and define $d_{in}(j), d_{out}(j)$ to be the in and out degrees, respectively, of a group j individual. For $j = 1, 2, \dots, k$, denote by N_j the total number of group j individuals in the population, and by $I_j(t)$ the number of group j individuals who are infectious at time t . Then under the annealed network approximation, the total rate at which group j individuals become infected is given by

$$\frac{\kappa}{NE[d_{in}]} \left(\sum_{m=1}^k d_{out}(m) I_m \right) d_{in}(j) (N_j - I_j).$$

When individuals become infected, they remain so for an exponentially distributed time of mean $1/\gamma$ before returning to the susceptible state.

This network model may be approximated by the k -group model with transition rates given in Table 1 by taking

$$\begin{aligned} \beta &= \kappa E[d_{out}], \\ f_j &= p(c^{-1}(j)), \\ \mu_j &= d_{in}(j)/E[d_{in}], \\ \lambda_j &= d_{out}(j)/E[d_{out}], \end{aligned}$$

for $j = 1, 2, \dots, k$.

The undirected version of the above annealed network approximation (with $\lambda = \mu$) has been studied by Hindes and Schwartz (2016), via numerical solution of Hamilton’s equations of motion, Eqs. (15,16) below. We note that real-world social networks are often only partially directed, containing both directed and undirected edges (that is to say, bi-directional edges are over-represented compared to what would be expected by chance), see Spricer and Britton (2015). Each individual then has three-dimensional degree $(d_{in}, d_{out}, d_{un})$, where d_{un} is the number of undirected edges connected to the individual in question. The above annealed network approximation will remain valid if we replace (d_{in}, d_{out}) with $(d_{in} + d_{un}, d_{out} + d_{un})$, provided, as before, that the resulting network is uncorrelated.

In the next section we present some relevant general theory, before going on in Sect. 4 to apply these general methods to the model described above.

3 General theory regarding persistence time from endemicity

Consider an infection modelled by a continuous-time Markov process $\{X(t) : t \geq 0\}$ on finite state-space $S \subset \mathbb{Z}^k$ with transition rate matrix Q . Suppose that S is made up of an absorbing set of states A (corresponding to absence of disease) and a single transient

communicating class C . We denote by Q_C the transition rate matrix restricted to C . Then the infection will almost surely die out (i.e. the process will leave C) within finite time, and (Darroch and Seneta 1967) there exists a unique quasi-stationary distribution $\mathbf{q} = \{q_x : x \in C\}$ such that, for any initial state within C ,

$$q_x = \lim_{t \rightarrow \infty} \Pr(X(t) = x | X(0) \in C) \quad \text{for } x \in C.$$

That is, provided the infection does not die out, it will settle to the endemic distribution \mathbf{q} . The distribution \mathbf{q} may be found as the unique solution of

$$\mathbf{q} Q_C = -(1/\tau)\mathbf{q} \quad \text{with } \sum_{x \in C} q_x = 1, \quad (3)$$

where $-(1/\tau)$ is the eigenvalue of Q_C with largest real part. The time to extinction from quasi-stationarity is exponentially distributed with mean τ .

Although τ may be computed exactly from Eq. (3), this can become impractical when the state-space is large, and it is not straightforward from (3) to establish qualitative results. Approximation methods are therefore valuable, and in particular, methods from Hamiltonian statistical mechanics may be used to study the leading order asymptotic (large population) behaviour of τ , as follows.

Suppose that $X(t)$ is a density-dependent process in the sense of chapter 11 of Ethier and Kurtz (2005); that is, the transition rates are of the form

$$P(X(t + \delta t) = \mathbf{x} + \mathbf{l} | X(t) = \mathbf{x}) = N W_{\mathbf{l}}\left(\frac{\mathbf{x}}{N}\right) + o(\delta t) \quad \text{for } \mathbf{x} \in S, \mathbf{l} \in L, \quad (4)$$

for some functions $W_{\mathbf{l}} : \mathbb{R}^k \rightarrow \mathbb{R}^+$, where L is the set of possible jumps from each state $\mathbf{x} \in S$ and N is some parameter indicating overall size of the system (in our applications, N will be the size of the population). Under mild technical conditions (Ethier and Kurtz 2005, Theorem 11.2.1), the scaled process $X(t)/N$ converges almost surely over finite time intervals, as $N \rightarrow \infty$, to the solution $\mathbf{y}(t)$ of the ordinary differential equation system

$$\frac{d\mathbf{y}}{dt} = \sum_{\mathbf{l} \in L} \mathbf{l} W_{\mathbf{l}}(\mathbf{y}). \quad (5)$$

For our application, we suppose that the system (5) possesses two equilibrium points: a stable endemic equilibrium point \mathbf{y}^* with all components strictly positive, and an unstable disease-free equilibrium point at $\mathbf{y} = \mathbf{0}$. We next summarise some key results from the Hamiltonian approach, in a form suited to our application. Detailed justifications and extensions of the method may be found in the review paper (Assaf and Meerson 2017) and references therein.

The Hamiltonian of the system is defined to be

$$H(\mathbf{y}, \boldsymbol{\theta}) = \sum_{\mathbf{l} \in L} W_{\mathbf{l}}(\mathbf{y}) \left(e^{\boldsymbol{\theta}^T \mathbf{l}} - 1 \right). \quad (6)$$

This Hamiltonian determines the following two complementary Hamilton–Jacobi partial differential equations:

$$H\left(\mathbf{y}, \frac{\partial V}{\partial \mathbf{y}}\right) = 0 \text{ and } H\left(\frac{\partial U}{\partial \boldsymbol{\theta}}, \boldsymbol{\theta}\right) = 0. \tag{7}$$

Each of these Hamilton–Jacobi equations is a way of expressing the eigenvector equation (3) while retaining only leading order terms in the limit $N \rightarrow \infty$ (see ‘‘Appendix’’ for a brief outline of the derivations).

If we can solve either of the Hamilton–Jacobi equations (7), the leading-order asymptotic behaviour of the mean time to extinction τ is given by

$$\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} = V(\mathbf{0}) - V(\mathbf{y}^*) = U(\mathbf{0}) - U(\boldsymbol{\theta}^*), \tag{8}$$

where \mathbf{y}^* is the endemic equilibrium point of the deterministic system (5), and $\boldsymbol{\theta}^*$ is the (assumed unique) non-zero equilibrium point of the complementary system

$$\frac{d\boldsymbol{\theta}}{dt} = - \left. \frac{\partial H}{\partial \mathbf{y}} \right|_{\mathbf{y}=\mathbf{0}}. \tag{9}$$

Note that system (5) may be recovered as $\left. \frac{d\mathbf{y}}{dt} = \frac{\partial H}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\mathbf{0}}$.

The solutions $U(\boldsymbol{\theta})$, $V(\mathbf{y})$ to equations (7) are related via the Legendre transform; that is,

$$U(\boldsymbol{\theta}) = \sup_{\mathbf{y}} \left\{ \mathbf{y}^T \boldsymbol{\theta} - V(\mathbf{y}) \right\}, \quad V(\mathbf{y}) = \sup_{\boldsymbol{\theta}} \left\{ \boldsymbol{\theta}^T \mathbf{y} - U(\boldsymbol{\theta}) \right\},$$

see Masoero (2014).

When (as is usually the case) it is not possible to find an analytical solution to either of the Hamilton–Jacobi equations (7), they may be solved numerically using the method of characteristics. That is, we write down the following $2k$ -dimensional system of ordinary differential equations:

$$\left. \begin{aligned} \frac{d\mathbf{y}}{dt} &= \frac{\partial H}{\partial \boldsymbol{\theta}} = \sum_{l \in L} l W_l(\mathbf{y}) e^{\boldsymbol{\theta}^T l}, \\ \frac{d\boldsymbol{\theta}}{dt} &= - \frac{\partial H}{\partial \mathbf{y}} = - \sum_{l \in L} \frac{\partial W_l}{\partial \mathbf{y}} \left(e^{\boldsymbol{\theta}^T l} - 1 \right), \end{aligned} \right\} \tag{10}$$

referred to as the ‘equations of motion’ of the system, and apply an appropriate numerical solver to (10). We then have $\lim_{N \rightarrow \infty} (\ln \tau)/N = A$, where A is the ‘action’ integral,

$$A = \int_{-\infty}^{\infty} \boldsymbol{\theta}^T \frac{d\mathbf{y}}{dt} dt = - \int_{-\infty}^{\infty} \mathbf{y}^T \frac{d\boldsymbol{\theta}}{dt} dt, \tag{11}$$

the integral in each case being evaluated along a trajectory from $(\mathbf{y}^*, \mathbf{0})$ to $(\mathbf{0}, \boldsymbol{\theta}^*)$. Note that $A = V(\mathbf{0}) - V(\mathbf{y}^*) = U(\mathbf{0}) - U(\boldsymbol{\theta}^*)$.

Having set out the general Hamiltonian approach, we will now apply this technique to the infection model described in Sect. 2 above.

4 Asymptotic persistence time formulae

Recall the infection model $\{\mathbf{I}(t) : t \geq 0\}$ described in Sect. 2, with transition rates given in Table 1. In the large population limit, the scaled infection process $\mathbf{I}(t)/N$ converges almost surely, over finite time intervals, to the deterministic process $\mathbf{y}(t)$ satisfying the system of ordinary differential equations (5); that is,

$$\frac{dy_i}{dt} = \beta \left(\sum_{j=1}^k \lambda_j y_j \right) \mu_i (f_i - y_i) - \gamma y_i \quad \text{for } i = 1, 2, \dots, k. \tag{12}$$

For $R_0 > 1$ there is a unique non-zero equilibrium point \mathbf{y}^* of the system (12), and it is globally asymptotically stable (Lajmanovich and Yorke 1976). This endemic equilibrium point \mathbf{y}^* is given by (Nold 1980)

$$y_i^* = \frac{\mu_i f_i D(\boldsymbol{\lambda}, \boldsymbol{\mu})}{1 + \mu_i D(\boldsymbol{\lambda}, \boldsymbol{\mu})} \quad \text{for } i = 1, 2, \dots, k, \tag{13}$$

where $D(\boldsymbol{\lambda}, \boldsymbol{\mu})$ is the unique positive solution of

$$\frac{\beta}{\gamma} \sum_{j=1}^k \frac{\mu_j f_j \lambda_j}{1 + \mu_j D(\boldsymbol{\lambda}, \boldsymbol{\mu})} = 1. \tag{14}$$

The Hamiltonian (6) corresponding to the process $\mathbf{I}(t)$ is

$$H(\mathbf{y}, \boldsymbol{\theta}) = \beta \left(\sum_{j=1}^k \lambda_j y_j \right) \left(\sum_{i=1}^k \mu_i (f_i - y_i) (e^{\theta_i} - 1) \right) + \gamma \sum_{i=1}^k y_i (e^{-\theta_i} - 1).$$

The corresponding equations of motion (10) are, for $i = 1, 2, \dots, k$,

$$\frac{dy_i}{dt} = \beta \left(\sum_{j=1}^k \lambda_j y_j \right) \mu_i (f_i - y_i) e^{\theta_i} - \gamma y_i e^{-\theta_i}, \tag{15}$$

$$\begin{aligned} \frac{d\theta_i}{dt} = & -\beta \lambda_i \sum_{j=1}^k \mu_j (f_j - y_j) (e^{\theta_j} - 1) + \beta \left(\sum_{j=1}^k \lambda_j y_j \right) \mu_i (e^{\theta_i} - 1) \\ & - \gamma (e^{-\theta_i} - 1). \end{aligned} \tag{16}$$

The non-zero equilibrium point θ^* given by (9) satisfies

$$\beta \lambda_i \sum_{j=1}^k \mu_j f_j \left(e^{\theta_j^*} - 1 \right) + \gamma \left(e^{-\theta_i^*} - 1 \right) = 0 \quad \text{for } i = 1, 2, \dots, k. \quad (17)$$

Setting $B = (\beta/\gamma) \sum_j f_j \mu_j \left(1 - e^{\theta_j^*} \right)$, then (17) implies that

$$e^{-\theta_i^*} = 1 + \lambda_i B \quad \text{for } i = 1, 2, \dots, k.$$

Substituting back into Eq. (17), we find that either $B = 0$ (corresponding to $\theta = \mathbf{0}$) or $B = D(\mu, \lambda)$. The elements of θ^* are thus

$$\theta_i^* = -\ln(1 + \lambda_i D(\mu, \lambda)) \quad \text{for } i = 1, 2, \dots, k.$$

So far, we have allowed for heterogeneities in both infectivity and susceptibility simultaneously. If we restrict to only one type of heterogeneity, then it becomes possible to find an explicit formula for the action A . Our main result is the following.

Theorem 1 *Consider the heterogeneous SIS infection model defined in Sect. 2, with transition rates given in Table 1, and suppose $R_0 > 1$. Recall that τ denotes the mean time from quasi-stationarity to disease extinction, and that $D(\lambda, \mu)$ is defined to be the unique positive solution of Eq. (14).*

(i) *If heterogeneity is in infectivity alone ($\mu = \mathbf{1}$), then*

$$\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} = \sum_{i=1}^k f_i \ln(1 + \lambda_i D(\mathbf{1}, \lambda)) - \frac{\gamma}{\beta} D(\mathbf{1}, \lambda). \quad (18)$$

(ii) *If heterogeneity is in susceptibility alone ($\lambda = \mathbf{1}$), then*

$$\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} = \sum_{i=1}^k f_i \ln(1 + \mu_i D(\mathbf{1}, \mu)) - \frac{\gamma}{\beta} D(\mathbf{1}, \mu).$$

(Note: under the network interpretation, the assumption $\mu = \mathbf{1}$ corresponds to every individual having the same in-degree, whereas $\lambda = \mathbf{1}$ corresponds to every individual having the same out-degree.)

Proof (i) Suppose that $\mu_i = 1$ for all i , and consider a trajectory $\{\theta(z) : 0 \leq z \leq D(\mathbf{1}, \lambda)\}$ along which $\theta_i = -\ln(1 + \lambda_i z)$ for $i = 1, 2, \dots, k$. Along such a trajectory, the Hamiltonian simplifies to

$$H(\mathbf{y}, \theta) = \beta \left(\sum_{j=1}^k \lambda_j y_j \right) \left(\sum_{i=1}^k (f_i - y_i) \left(\frac{1}{1 + \lambda_i z} - 1 \right) \right) + \gamma \sum_{i=1}^k y_i \lambda_i z$$

$$= \gamma z \left(\sum_{j=1}^k \lambda_j y_j \right) \left(1 - \frac{\beta}{\gamma} \sum_{i=1}^k \frac{f_i \lambda_i}{1 + \lambda_i z} + \frac{\beta}{\gamma} \sum_{i=1}^k \frac{y_i \lambda_i}{1 + \lambda_i z} \right).$$

Since $z > 0$ and $\sum_j \lambda_j y_j > 0$ (except at endpoints of the trajectory) the Hamilton–Jacobi equation $H\left(\frac{\partial U}{\partial \theta}, \theta\right) = 0$ reduces to

$$1 - \frac{\beta}{\gamma} \sum_{i=1}^k \frac{f_i \lambda_i}{1 + \lambda_i z} + \frac{\beta}{\gamma} \sum_{i=1}^k \frac{\lambda_i}{1 + \lambda_i z} \frac{\partial U}{\partial \theta_i} = 0. \quad (19)$$

Now along the trajectory under consideration, we have

$$\frac{dU}{dz} = - \sum_{i=1}^k \frac{\lambda_i}{1 + \lambda_i z} \frac{\partial U}{\partial \theta_i}$$

and so Eq. (19) becomes

$$\begin{aligned} \frac{dU}{dz} &= \frac{\gamma}{\beta} - \sum_{i=1}^k \frac{f_i \lambda_i}{1 + \lambda_i z} \\ \Rightarrow U(\theta(z)) - U(\theta(0)) &= \left(\frac{\gamma}{\beta}\right) z - \sum_{i=1}^k \int_0^z \frac{f_i \lambda_i}{1 + \lambda_i x} dx \\ \Rightarrow U(\theta) - U(\mathbf{0}) &= \left(\frac{\gamma}{\beta}\right) z - \sum_{i=1}^k f_i \ln(1 + \lambda_i z) \\ \text{along } \theta_i &= -\ln(1 + \lambda_i z). \end{aligned}$$

The action is therefore given by

$$A = U(\mathbf{0}) - U(\theta^*) = \sum_{i=1}^k f_i \ln(1 + \lambda_i D(\mathbf{1}, \lambda)) - \frac{\gamma}{\beta} D(\mathbf{1}, \lambda),$$

as required.

- (ii) For the SIS model on a finite network, in which each individual u makes contact with each other individual v at rate β_{uv} , it is known that, provided infectious periods are exponentially distributed, the decay parameter of the process is unchanged under transposition of the matrix of infection rates $\{\beta_{uv}\}$. This follows from the property of ‘network duality’, see Wilkinson and Sharkey (2013), Holley and Liggett (1975) and Harris (1976). In our context, this implies that the mean time to extinction from quasi-stationarity, τ , is identical if we interchange the roles of λ, μ . Hence part (ii) of the theorem follows immediately from part (i).

We can confirm this as follows. With $\lambda = \mathbf{1}$, the Hamiltonian may be written as

$$H(\mathbf{y}, \boldsymbol{\theta}) = \gamma \sum_{i=1}^k (e^{\theta_i} - 1) \left(\frac{\beta}{\gamma} \left(\sum_{j=1}^k y_j \right) \mu_i (f_i - y_i) - y_i e^{-\theta_i} \right).$$

With the convention that $y \ln y = 0$ when $y = 0$, take

$$\begin{aligned} V(\mathbf{y}) &= \sum_{i=1}^k y_i \left(1 + \ln y_i - \ln \left(\frac{\beta}{\gamma} \mu_i \right) \right) - \left(\sum_{i=1}^k y_i \right) \ln \left(\sum_{i=1}^k y_i \right) \\ &\quad + \sum_{i=1}^k (f_i - y_i) \ln (f_i - y_i). \end{aligned} \tag{20}$$

Then

$$\frac{\partial V}{\partial y_i} = \ln \left(\frac{y_i}{\frac{\beta}{\gamma} \mu_i (f_i - y_i) \left(\sum_j y_j \right)} \right) \quad \text{for } i = 1, 2, \dots, k,$$

and so

$$H \left(\mathbf{y}, \frac{\partial V}{\partial \mathbf{y}} \right) = 0.$$

That is, $V(\mathbf{y})$ satisfies the relevant Hamilton–Jacobi equation. The action is then given by

$$A = V(\mathbf{0}) - V(\mathbf{y}^*) = \sum_{i=1}^k f_i \ln (1 + \mu_i D(\mathbf{1}, \boldsymbol{\mu})) - \frac{\gamma}{\beta} D(\mathbf{1}, \boldsymbol{\mu}).$$

As expected, we recover the formula for the case of heterogeneous infectivity, but with the roles of $\lambda, \boldsymbol{\mu}$ interchanged.

Having found the solution $V(\mathbf{y})$ for the case $\lambda = \mathbf{1}$, we can find the corresponding function $U(\boldsymbol{\theta})$ as the Legendre transform of $V(\mathbf{y})$. For $i = 1, 2, \dots, k$, we have

$$\frac{d}{dy_i} \left(\mathbf{y}^T \boldsymbol{\theta} - V(\mathbf{y}) \right) = \theta_i - \ln \left(\frac{y_i}{\frac{\beta}{\gamma} \mu_i (f_i - y_i) \left(\sum_j y_j \right)} \right),$$

and so any stationary point satisfies

$$y_i = \frac{\frac{\beta}{\gamma} \mu_i f_i \left(\sum_j y_j \right) e^{\theta_i}}{1 + \frac{\beta}{\gamma} \mu_i \left(\sum_j y_j \right) e^{\theta_i}}. \tag{21}$$

For $\theta \in \mathbb{R}^k$, define the function $Q(\mu, \theta)$ to be the solution of

$$\frac{\beta}{\gamma} \sum_{j=1}^k \frac{\mu_j f_j}{e^{-\theta_j} + \mu_j Q(\mu, \theta)} = 1.$$

Setting $R = (\beta/\gamma) \sum_j y_j$ and substituting from (21) into the definition of R , we find that $R = Q(\mu, \theta)$, and hence the function $y^T \theta - V(y)$ has a stationary point at

$$y_i = \frac{\mu_i f_i e^{\theta_i} Q(\mu, \theta)}{1 + \mu_i e^{\theta_i} Q(\mu, \theta)}. \quad (22)$$

Evaluating the function $y^T \theta - V(y)$ at the point (22), we find

$$U(\theta) = \sum_{i=1}^k f_i \ln(1 + \mu_i e^{\theta_i} Q(\mu, \theta)) - \frac{\gamma}{\beta} Q(\mu, \theta) \quad (23)$$

and can easily verify that the function (23) does indeed satisfy $H\left(\frac{\partial U}{\partial \theta}, \theta\right) = 0$. Now $Q(\mu, \mathbf{0}) = D(\mathbf{1}, \mu)$ and $Q(\mu, \theta^*) = 0$, so we once again find that

$$A = U(\mathbf{0}) - U(\theta^*) = \sum_{i=1}^k f_i \ln(1 + \mu_i D(\mathbf{1}, \mu)) - \frac{\gamma}{\beta} D(\mathbf{1}, \mu).$$

□

Although we did not actually need to find the functions $U(\theta)$, $V(y)$ in order to prove Theorem 1(ii), we include them because knowledge of these functions can be of assistance in generalising and extending our results. We will demonstrate this in Theorem 3 below.

Figure 1 illustrates Theorem 1 in the case of $k = 2$ groups with heterogeneity in infectivity (the graph for the corresponding case with heterogeneity in susceptibility is identical, by network duality). The exact value of $(\ln \tau)/N$ is computed from Eq. (3) for total population sizes $N = 100, 150, \dots, 650$. The action A is computed from Eq. (18). For comparison, we also show the action $A_0 = (1/R_0) - 1 + \ln R_0$ computed for the homogeneous population SIS model with the same value for R_0 . We see that formula (18) gives a good approximation to $(\ln \tau)/N$ for population sizes from around $N = 300$ upwards. We can also see that if we were to ignore heterogeneity and use the homogeneous population result, we would drastically over-estimate the persistence time of infection. We demonstrate this point in Theorem 2 below, as well as comparing heterogeneous populations of greater or lesser degrees of heterogeneity.




Fig. 1 Values of $(\ln \tau)/N$ and asymptotic formulae plotted against population size N . Fixed parameter values $k = 2$, $\mathbf{f} = (0.5, 0.5)$, $\lambda = \frac{2}{51}(50, 1)$, $\boldsymbol{\mu} = (1, 1)$, $R_0 = 1.5$, $\gamma = 1$. For these parameter values, $D(\lambda, \boldsymbol{\mu}) = 0.5$, $\mathbf{y}^* = (1/6, 1/6)$, $D(\boldsymbol{\mu}, \lambda) = 0.2625$, $\boldsymbol{\theta}^* = (-0.4152 - 0.0102)$, with action $A \approx 0.0377$ and corresponding action for the homogeneous case $A_0 \approx 0.0721$. The dots, labelled ‘eigenvalue computation’, are the true values of $(\ln \tau)/N$ computed from Eq. (3); the action A is computed from Eq. (18); the homogeneous action is computed as $A_0 = (1/R_0) - 1 + \ln R_0$

5 The effect of increasing heterogeneity

Using the formulae of the previous section, we are now in a position to investigate the effect of increasing heterogeneity upon the persistence time of infection. First, in order to compare different levels of heterogeneity, we recall the definition of majorization (Marshall et al. 2011). For any $\mathbf{x} \in (\mathbb{R}^+)^k$, denote by $x_{[1]} \geq x_{[2]} \geq \dots \geq x_{[k]}$ the ordered components of \mathbf{x} . Then for $\mathbf{x}^{(1)}, \mathbf{x}^{(2)} \in (\mathbb{R}^+)^k$, we say $\mathbf{x}^{(1)}$ is majorized by $\mathbf{x}^{(2)}$, denoted $\mathbf{x}^{(1)} < \mathbf{x}^{(2)}$, if $\sum_{i=1}^k x_i^{(1)} = \sum_{i=1}^k x_i^{(2)}$ and $\sum_{i=1}^j x_{[i]}^{(1)} \leq \sum_{i=1}^j x_{[i]}^{(2)}$ for $j = 1, 2, \dots, k - 1$. An equivalent definition is that $\sum_{i=1}^k \phi(x_i^{(1)}) \leq \sum_{i=1}^k \phi(x_i^{(2)})$ for all convex functions $\phi(\cdot)$. Intuitively, $\mathbf{x}^{(2)}$ is ‘more heterogeneous’ than $\mathbf{x}^{(1)}$. More generally, given a probability vector (with components summing to 1) $\mathbf{p} \in (\mathbb{R}^+)^k$, then $\mathbf{x}^{(1)}$ is \mathbf{p} -majorized by $\mathbf{x}^{(2)}$, written $\mathbf{x}^{(1)} <_{\mathbf{p}} \mathbf{x}^{(2)}$, if there exists a permutation σ such that $x_{\sigma(1)}^{(1)} \geq x_{\sigma(2)}^{(1)} \geq \dots \geq x_{\sigma(k)}^{(1)}$ and $x_{\sigma(1)}^{(2)} \geq x_{\sigma(2)}^{(2)} \geq \dots \geq x_{\sigma(k)}^{(2)}$ with $\sum_{i=1}^k p_i x_i^{(1)} = \sum_{i=1}^k p_i x_i^{(2)}$ and $\sum_{i=1}^j p_{\sigma(i)} x_{\sigma(i)}^{(1)} \leq \sum_{i=1}^j p_{\sigma(i)} x_{\sigma(i)}^{(2)}$ for $j = 1, 2, \dots, k - 1$.

Theorem 2 Consider two populations, with $\beta^{(1)} = \beta^{(2)} = \beta$, $\gamma^{(1)} = \gamma^{(2)} = \gamma$, and each having the same group structure $\mathbf{f}^{(1)} = \mathbf{f}^{(2)} = \mathbf{f}$, where we use superscripts (1), (2) to denote the population under consideration. Recall that τ denotes the mean time from quasi-stationarity to disease extinction.

(i) *With heterogeneity in infectivity alone,*

$$\lambda^{(1)} \prec_f \lambda^{(2)} \Rightarrow \lim_{N \rightarrow \infty} \frac{\ln \tau^{(1)}}{N} \geq \lim_{N \rightarrow \infty} \frac{\ln \tau^{(2)}}{N}.$$

(ii) *With heterogeneity in susceptibility alone,*

$$\mu^{(1)} \prec_f \mu^{(2)} \Rightarrow \lim_{N \rightarrow \infty} \frac{\ln \tau^{(1)}}{N} \geq \lim_{N \rightarrow \infty} \frac{\ln \tau^{(2)}}{N}.$$

In particular, provided heterogeneity is in either infectivity or susceptibility but not both, then $\lim_{N \rightarrow \infty} (\ln \tau) / N$ is maximised in the homogeneous case.

Proof Consider the case of heterogeneity in infectivity, and suppose that $\lambda^{(1)} \prec_f \lambda^{(2)}$. The function $h(x) = x / (1 + xD(\mathbf{1}, \lambda^{(1)}))$ is concave for $x > 0$, and so applying proposition 14.A.3 of Marshall et al. (2011),

$$\sum_{i=1}^k \frac{f_i \lambda_i^{(2)}}{1 + \lambda_i^{(2)} D(\mathbf{1}, \lambda^{(1)})} \leq \sum_{i=1}^k \frac{f_i \lambda_i^{(1)}}{1 + \lambda_i^{(1)} D(\mathbf{1}, \lambda^{(1)})} = \frac{\gamma}{\beta},$$

the final equality coming from the definition (14) of $D(\mu, \lambda)$ for population 1. The expression $\sum_i f_i \lambda_i^{(2)} / (1 + \lambda_i^{(2)} z)$ is a decreasing function of z , and so from Eq. (14) for population 2 it follows that $D(\mathbf{1}, \lambda^{(1)}) \geq D(\mathbf{1}, \lambda^{(2)})$.

Now define the function $\psi(z) = \sum_i f_i \ln(1 + \lambda_i^{(1)} z) - (\gamma z / \beta)$. Then

$$\frac{d\psi}{dz} = \sum_i \frac{f_i \lambda_i^{(1)}}{1 + \lambda_i^{(1)} z} - \frac{\gamma}{\beta},$$

so that $\psi'(z) > 0$ for $0 < z < D(\mathbf{1}, \lambda^{(1)})$, and hence $\psi(D(\mathbf{1}, \lambda^{(2)})) \leq \psi(D(\mathbf{1}, \lambda^{(1)})) = A^{(1)}$. That is,

$$\sum_{i=1}^k f_i \ln(1 + \lambda_i^{(1)} D(\mathbf{1}, \lambda^{(2)})) - \frac{\gamma D(\mathbf{1}, \lambda^{(2)})}{\beta} \leq A^{(1)}. \tag{24}$$

The function $g(x) = \ln(1 + D(\mathbf{1}, \lambda^{(2)}) x)$ is concave for $x > 0$, and so again applying proposition 14.A.3 of Marshall et al. (2011),

$$\sum_{i=1}^k f_i \ln(1 + \lambda_i^{(2)} D(\mathbf{1}, \lambda^{(2)})) \leq \sum_{i=1}^k f_i \ln(1 + \lambda_i^{(1)} D(\mathbf{1}, \lambda^{(2)})).$$

Combining with (24) yields

$$A^{(2)} \leq A^{(1)},$$




Fig. 2 Contour plot showing the action A (solid contours) as a function of λ_1, μ_1 . Fixed parameter values $k = 2, f = (0.5, 0.5), R_0 = 1.2$. The action is maximised at $(\lambda_1, \mu_1) = (1, 1)$, with value $A_0 = (1/R_0) - 1 + \ln R_0 \approx 0.0157$. Dashed contours show the approximation \hat{A} computed from formula (25). Dotted contours show a finite-population approximation—see main text for details

and the result follows.

Part (ii) of the theorem follows immediately by interchanging the roles of λ, μ . \square

Figure 2 illustrates Theorem 2, as well as showing the effect of allowing heterogeneity in both infectivity and susceptibility simultaneously, for the case of $k = 2$ equal-sized groups ($f_1 = f_2 = 0.5$). The constraints on the elements of λ, μ in this case reduce to $\lambda_1 + \lambda_2 = \mu_1 + \mu_2 = 2$, and so we plot the action as a function of $(\lambda_1, \mu_1) \in (0, 2)^2$. We choose to keep R_0 fixed, with the value of β being varied in order to achieve this. With both heterogeneities present, we have no explicit formula for the action A , and instead compute it by first solving the equations of motion (15,16) numerically using the Matlab `bvp4c` command, and then integrating the numerical solution along the trajectory, Eq. (11). The solid contours in Fig. 2 show the action values A computed in this way. Note that the transformation $(\lambda_1, \mu_1) \rightarrow (2 - \lambda_1, 2 - \mu_1)$ here amounts to simply re-labelling the groups, so that Fig. 2 is invariant under a rotation of half a turn around the point $(1, 1)$; also, we know from network duality that the action is unchanged under the transformation $(\lambda_1, \mu_1) \rightarrow (\mu_1, \lambda_1)$, so that Fig. 2 is invariant under reflection in the line $\lambda_1 = \mu_1$.

Although we do not have an explicit formula for A when both types of heterogeneity are present, we can obtain an approximate formula valid for R_0 close to (and above) 1,

as follows. For R_0 only slightly larger than 1, the trajectory from $(\mathbf{y}^*, \mathbf{0})$ to $(\mathbf{0}, \boldsymbol{\theta}^*)$ satisfying equations (15,16) may be approximated by a straight line. The integral (11) can easily be evaluated along the straight line connecting $(\mathbf{y}^*, \mathbf{0})$ to $(\mathbf{0}, \boldsymbol{\theta}^*)$, and we thus find that $A \approx \tilde{A}$, where

$$\tilde{A} = \frac{D(\boldsymbol{\lambda}, \boldsymbol{\mu})}{2} \sum_{i=1}^k \frac{\mu_i f_i \ln(1 + \lambda_i D(\boldsymbol{\mu}, \boldsymbol{\lambda}))}{1 + \mu_i D(\boldsymbol{\lambda}, \boldsymbol{\mu})}. \quad (25)$$

Values of \tilde{A} are shown as dashed contours in Fig. 2, and we see that for these parameter values \tilde{A} does indeed provide a reasonable approximation to the true value A . Note that although the action A is known (from network duality) to be invariant under the interchange of $\boldsymbol{\lambda}$, $\boldsymbol{\mu}$, the approximating formula (25) does not possess this symmetry. Nevertheless, we see that in Fig. 2 the contours corresponding to \tilde{A} do appear symmetrical under reflection in the line $\lambda_1 = \mu_1$, so that the expected symmetry does at least hold approximately here.

For comparison, the dotted contours in Fig. 2 were computed by solving the eigenvalue equation (3) numerically for $N = 400$ and $N = 500$, and assuming (without proof) that asymptotic formula (1) is valid for our model. Denoting by τ_N the mean time from quasi-stationarity to disease extinction in a population of size N , formula (1) implies that the action A may be approximated by

$$\left(\ln(\tau_{500} \sqrt{500}) - \ln(\tau_{400} \sqrt{400}) \right) / 100, \quad (26)$$

and the dotted contours show computed values of formula (26). The fact that the dotted contours closely follow the solid contours provides some confirmation both that the action A gives a good approximation to $(\ln \tau)/N$ for population sizes above $N = 400$, and that formula (1) does indeed apply to our model.

We see from Fig. 2 that the action decreases as we move away from the point $(\lambda_1, \mu_1) = (1, 1)$, not only along the lines $\lambda_1 = 1$ and $\mu_1 = 1$, as ensured by Theorem 2, but in any direction. That is, heterogeneity in infectivity, or susceptibility, or any combination of the two, reduces the value of $\lim_{N \rightarrow \infty} (\ln \tau)/N$ compared to the homogeneous case. We discuss this further in Sect. 7 below.

6 Generalising the infectious period distribution

So far, we have made the conventional assumption that individuals' infectious periods are exponentially distributed. This is purely a mathematical convenience, not motivated by biological realism. Realism can be greatly improved by allowing infectious periods to follow an Erlang distribution, using the 'method of stages'. That is, when an individual becomes infected, it passes through s infectious stages, remaining in each stage for an exponentially distributed time of mean $(s\gamma)^{-1}$, before returning to susceptibility. As before, we denote by N_j the (constant) number of individuals in group j , and by $N = N_1 + N_2 + \dots + N_k$ the total population size. Denoting by $I_{jv}(t)$ the number of group j individuals in infectious stage v at time t , then

Table 2 Transition rates for the k -group, s -stage SIS model

Event	State transition	Transition rate
Infection in group j	$I_{j1} \rightarrow I_{j1} + 1$	$\frac{\beta}{N} \left(\sum_{m=1}^k \lambda_m \sum_{v=1}^s I_{mv} \right) \mu_j (N_j - \sum_{v=1}^s I_{jv})$
Transition to next infectious stage	$(I_{j,v-1}, I_{jv}) \rightarrow (I_{j,v-1} - 1, I_{jv} + 1)$	$s\gamma I_{j,v-1}$ for $v = 2, 3, \dots, s$
Recovery in group j	$I_{js} \rightarrow I_{js} - 1$	$s\gamma I_{js}$

$\{I_{jv}(t) : j = 1, 2, \dots, k, v = 1, 2, \dots, s, t \geq 0\}$ is a continuous-time Markov chain with transition rates given in Table 2. The number of susceptible individuals in group j is $S_j(t) = N_j - \sum_{v=1}^s I_{jv}(t)$.

Writing $\mathbf{y} = \{y_{iv} : i = 1, 2, \dots, k, v = 1, 2, \dots, s\}$ and $\boldsymbol{\theta} = \{\theta_{iv} : i = 1, 2, \dots, k, v = 1, 2, \dots, s\}$, the corresponding Hamiltonian is

$$\begin{aligned}
 H(\mathbf{y}, \boldsymbol{\theta}) = & \beta \left(\sum_{j=1}^k \lambda_j \sum_{v=1}^s y_{jv} \right) \left(\sum_{i=1}^k \mu_i \left(f_i - \sum_{v=1}^s y_{iv} \right) \right) (e^{\theta_{i1}} - 1) \\
 & + s\gamma \sum_{i=1}^k \sum_{v=1}^{s-1} y_{iv} (e^{-\theta_{iv} + \theta_{i,v+1}} - 1) + s\gamma \sum_{i=1}^k y_{is} (e^{-\theta_{is}} - 1).
 \end{aligned}$$

It is immediate from equation (17) of Clancy (2015) that the deterministic endemic equilibrium point is given by

$$y_{iv}^* = \frac{y_i^*}{s} \quad \text{for } i = 1, 2, \dots, k, v = 1, 2, \dots, s,$$

where y_i^* is the solution (13) for the model with exponentially distributed infectious periods ($s = 1$).

It is straightforward to show that the elements of $\boldsymbol{\theta}^*$ are given by

$$\theta_{iv}^* = -(s + 1 - v) \ln(1 + \lambda_i D_s(\boldsymbol{\mu}, \boldsymbol{\lambda}))$$

where $D_s(\boldsymbol{\mu}, \boldsymbol{\lambda})$ is the solution of

$$\frac{\beta}{s\gamma} \sum_{j=1}^k \mu_j f_j \lambda_j \sum_{v=1}^s \left(\frac{1}{1 + \lambda_j D_s(\boldsymbol{\mu}, \boldsymbol{\lambda})} \right)^v = 1. \tag{27}$$

For the SIS model with Erlang-distributed infectious periods in a homogeneous population ($k = 1$), the solution $U(\boldsymbol{\theta})$ to the relevant Hamilton–Jacobi equation was found in Clancy and Tjia (2018) to be

$$U(\boldsymbol{\theta}) = \ln \left(\sum_{v=1}^s e^{\theta_v} \right) + \frac{\gamma}{\beta} \left(\frac{s}{\sum_{v=1}^s e^{\theta_v}} \right).$$

Taking the Legendre transform, we find that $V(\mathbf{y})$ for this homogeneous-population model is given by

$$\begin{aligned} V(\mathbf{y}) &= \sup_{\boldsymbol{\theta}} \left\{ \boldsymbol{\theta}^T \mathbf{y} - U(\boldsymbol{\theta}) \right\} \\ &= \sum_{v=1}^s y_v \left(1 + \ln y_v - \ln \left(\frac{\beta}{s\gamma} \right) \right) - \left(\sum_{v=1}^s y_v \right) \ln \left(\sum_{v=1}^s y_v \right) \\ &\quad + \left(1 - \sum_{v=1}^s y_v \right) \ln \left(1 - \sum_{v=1}^s y_v \right). \end{aligned} \quad (28)$$

Comparing solution (28) for the SIS model with Erlang-distributed infectious periods in a homogeneous population and solution (20) for the SIS model with exponentially distributed infectious periods and heterogeneous susceptibilities, one may now guess the form of the solution $V(\mathbf{y})$ for the SIS model with Erlang-distributed infectious periods and heterogeneous susceptibilities, and verify that the relevant Hamilton–Jacobi equation is indeed satisfied. The solution is thus found to be

$$\begin{aligned} V(\mathbf{y}) &= \sum_{i=1}^k \sum_{v=1}^s y_{iv} \left(1 + \ln y_{iv} - \ln \left(\frac{\beta}{s\gamma} \mu_i \right) \right) - \left(\sum_{i=1}^k \sum_{v=1}^s y_{iv} \right) \ln \left(\sum_{i=1}^k \sum_{v=1}^s y_{iv} \right) \\ &\quad + \sum_{i=1}^k \left(f_i - \sum_{v=1}^s y_{iv} \right) \ln \left(f_i - \sum_{v=1}^s y_{iv} \right). \end{aligned} \quad (29)$$

Taking the Legendre transform, we find

$$U(\boldsymbol{\theta}) = \sum_{i=1}^k f_i \ln \left(1 + \mu_i \left(\sum_{v=1}^s e^{\theta_{iv}} \right) Q_s(\boldsymbol{\mu}, \boldsymbol{\theta}) \right) - \frac{\gamma s}{\beta} Q_s(\boldsymbol{\mu}, \boldsymbol{\theta})$$

where $Q_s(\boldsymbol{\mu}, \boldsymbol{\theta})$ is the solution of

$$\frac{\beta}{\gamma s} \sum_{j=1}^k \frac{\mu_j f_j \left(\sum_{v=1}^s e^{\theta_{jv}} \right)}{1 + \mu_j \left(\sum_{v=1}^s e^{\theta_{jv}} \right) Q_s(\boldsymbol{\mu}, \boldsymbol{\theta})} = 1.$$

The action A in this case is thus

$$\begin{aligned} A &= V(\mathbf{0}) - V(\mathbf{y}^*) \\ &= U(\mathbf{0}) - U(\boldsymbol{\theta}^*) = \sum_i f_i \ln(1 + \mu_i D(\mathbf{1}, \boldsymbol{\mu})) - \frac{\gamma}{\beta} D(\mathbf{1}, \boldsymbol{\mu}), \end{aligned}$$

as before, and the following result is immediate.




Fig. 3 The effect of the infectious period distribution upon the mean persistence time of infection τ . Fixed parameter values $k = 2$, $f = (0.5, 0.5)$, $R_0 = 1.2$, $\gamma = 1$. Solid line ('Action') has gradient $A \approx 0.0110$ given by Eq. (18) with $\lambda = (5/3, 1/3)$, $\mu = (1, 1)$, intercept chosen arbitrarily; dashed line ('Exponential infectious periods') computed from Eq. (3) with $\lambda = (5/3, 1/3)$, $\mu = (1, 1)$; crosses ('Heterogeneous susceptibilities') computed via simulation with $\lambda = (1, 1)$, $\mu = (5/3, 1/3)$ and constant infectious periods; circles ('Heterogeneous infectivities') computed via simulation with $\lambda = (5/3, 1/3)$, $\mu = (1, 1)$ and constant infectious periods

Theorem 3 *Theorem 1(ii) remains valid if infectious periods are allowed to follow an Erlang, rather than exponential, distribution. Consequently, Theorem 2(ii) likewise remains valid with Erlang-distributed infectious periods.*

Figure 3 illustrates the effect of the infectious period distribution in the case of $k = 2$ equal-sized groups ($f_1 = f_2 = 0.5$). In constructing this figure, we have assumed (without proof) that asymptotic formula (1) is valid for our model. Consequently, we plot the function $\frac{1}{2} \ln N + \ln \tau$, which according to formula (1) should, as N increases, approach a straight line of gradient A and intercept $\ln C$. The dashed line, corresponding to exponentially distributed infectious periods, was computed using the eigenvalue characterisation (3). By network duality, this dashed line may be interpreted as corresponding to either heterogeneous infectivity ($\lambda = (5/3, 1/3)$, $\mu = (1, 1)$) or heterogeneous susceptibility ($\lambda = (1, 1)$, $\mu = (5/3, 1/3)$). We used Monte Carlo simulation to estimate the mean persistence time τ with constant (non-random) infectious periods for the cases of heterogeneous infectivity and heterogeneous susceptibility separately. This infectious period distribution corresponds to an Erlang distribution with s stages in the limit as $s \rightarrow \infty$. An issue that arises is that the time until extinction of infection, starting from quasi-stationarity, is exponentially distributed with mean increasing exponentially in population size, so that to simulate the process to extinction can be very time-consuming. To get around this, we fixed times t_0 (the burn-in period) and t_{max} such that (i) by time t_0 the state of the process is approximately quasi-stationary (having started at time zero close to the re-scaled deterministic equilibrium

point Ny^* , and conditioning upon survival to time t_0); and (ii) by time t_{max} a substantial proportion of all simulations have reached extinction. We then estimated the mean time to extinction τ using the maximum likelihood estimator. That is, denoting by T_1, T_2, \dots, T_r the extinction times of those simulations that went extinct within the time window (t_0, t_{max}) , and by m the number of simulations that had not gone extinct by time t_{max} , our estimate is

$$\hat{\tau} = \frac{m(t_{max} - t_0) + \sum_{i=1}^r (T_i - t_0)}{r}.$$

We have included in Fig. 3 a solid line with gradient equal to the action A computed from formula (18). Note that the intercept of this line was chosen arbitrarily, since we have no way to evaluate the constant C in formula (1) for our model.

We see from Fig. 3 that the dashed line corresponding to exponentially distributed infectious periods does indeed appear to be a straight line of gradient A , providing some confirmation both that the action A gives a reasonable approximation to $(\ln \tau)/N$ for population sizes above $N = 200$ and that formula (1) is valid for our model. With constant infectious periods, we see that heterogeneous infectivity and heterogeneous susceptibility result in almost identical estimates of τ , and that these estimates lie close to a straight line of gradient A . It appears that, as with exponentially distributed infectious periods, the value of τ is unchanged if we interchange λ, μ . We therefore conjecture that, similarly to Theorem 3 above, Theorem 1(i) and Theorem 2(i) remain valid with Erlang-distributed infectious periods. The model with constant infectious periods has reduced mean persistence time τ compared to the model with exponentially distributed infectious periods, but the difference is in the pre-factor constant C and not the leading-order constant A , in line with the results of Ball et al. (2016) for the homogeneous population case.

7 Discussion and possible extensions

The main result of this paper, Theorem 1, provides a simple explicit formula for $\lim_{N \rightarrow \infty} (\ln \tau)/N$, where τ is the expected time from endemicity to extinction for an SIS infection model with heterogeneity in either infectivity or susceptibility of individuals, in a population of size N . The only infection model for which such a formula has previously been available is the SIS model in a homogeneous population, either with exponentially distributed infectious periods (Andersson and Djehiche 1998) or with arbitrary infectious period distribution (Ball et al. 2016). Theorem 1 thus represents a significant advance, but many open questions remain.

Firstly, for the SIS model in a homogeneous population, both of Andersson and Djehiche (1998) and Ball et al. (2016) established asymptotic approximations for τ of the form (1), with explicit formulae for the pre-factor constant C . Our result is less precise than this; we have not shown that an asymptotic formula of the form (1) is valid for our model (although we conjecture, and have presented some numerical evidence, that this is the case), nor have we attempted to evaluate the pre-factor constant C . The asymptotic form (1) has been shown by Assaf and Meerson (2010) to be valid (and

formulae given for the constant C) for general 1-dimensional processes of bounded jump size. The technique of Assaf and Meerson (2010) is an extension of the approach employed here, retaining terms beyond the leading order in N . The analysis is considerably more intricate than the leading-order treatment we have restricted ourselves to, and the extension of the approach of Assaf and Meerson (2010) to multi-dimensional processes such as the infection model considered here is the subject of ongoing work.

Secondly, our model as described in Sect. 2 incorporates heterogeneity in both infectivity and susceptibility simultaneously, but we have only been able to provide an explicit asymptotic formula for the cases in which only one of these two heterogeneities is present. When both heterogeneities are present, provided R_0 is only slightly above one then \tilde{A} given by formula (25) can be used to approximate $A = \lim_{N \rightarrow \infty} (\ln \tau)/N$, and in Fig. 2 we saw that \tilde{A} can indeed provide a reasonable approximation to A . However, even a small error in A can correspond to a very large error in our estimate of τ , so that an approximate formula such as (25) is of considerably less utility than an exact formula such as (18). In particular, the fact that we have not been able to find an exact formula such as (18) valid when both types of heterogeneity are present severely restricts the class of networks to which our results may be applied under the annealed network approximation—we require either that every individual has the same in-degree, or that every individual has the same out-degree. Nevertheless, for the class of directed networks to which they apply our results represent an interesting step forward, and since network models are of great ongoing interest in infection modelling, we now present our results in a form suited to the network interpretation.

As a preliminary to the statement of our results, we require the concept of convex ordering of random variables, defined as follows (Shaked and Shanthikumar 2007, section 3.A.1). Given two random variables $X^{(1)}$, $X^{(2)}$, then $X^{(2)}$ is greater than $X^{(1)}$ in the sense of convex ordering, denoted $X^{(1)} \leq_{cv} X^{(2)}$, if $E[\phi(X^{(1)})] \leq E[\phi(X^{(2)})]$ for all convex functions $\phi(\cdot)$. If $X^{(1)}$, $X^{(2)}$ take values in $\{1, 2, \dots, d_{\max}\}$, then an equivalent definition is that

$$\sum_{i=1}^j P(X^{(1)} \leq i) \leq \sum_{i=1}^j P(X^{(2)} \leq i) \quad \text{for } j = 1, 2, \dots, d_{\max}.$$

Note that $X^{(1)} \leq_{cv} X^{(2)}$ implies $E[X^{(1)}] = E[X^{(2)}]$; intuitively, $X^{(2)}$ is ‘more variable’ than $X^{(1)}$.

In comparing two populations, we need to define our ‘groups’ slightly differently than in Sect. 2. Specifically, partition the population into groups in such a way that two individuals belong to the same group if they share the same values of both $(d_{\text{in}}^{(1)}, d_{\text{out}}^{(1)})$ and $(d_{\text{in}}^{(2)}, d_{\text{out}}^{(2)})$; the condition $f^{(1)} = f^{(2)} = f$ required by Theorem 2 is thus satisfied. (As before, superscripts (1), (2) denote the population under consideration.)

Theorem 4 *Consider an SIS infection in a population of N individuals connected by an uncorrelated directed network. Each individual has in-degree and out-degree distributed as $(d_{\text{in}}, d_{\text{out}})$, the degrees of distinct individuals being mutually independent,*

with $E [d_{in}] = E [d_{out}] = \mu$ and $d_{in}, d_{out} \leq d_{max}$ for some $d_{max} \in \mathbb{N}$. Infection transmits along each link from an infectious to a susceptible individual at rate κ , and when an individual becomes infected it remains so for a time of mean $1/\gamma$ before returning to the susceptible state. Recall that τ denotes the expected time from quasi-stationarity to extinction of infection.

(a) Suppose that $P (d_{in} = \mu) = 1$, so every individual has the same in-degree, and that infectious periods are exponentially distributed. Then

(i) $\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} \approx A_{out}$, with

$$A_{out} = \sum_{i=1}^{d_{max}} P (d_{out} = i) \ln (1 + i D_{out}) - \frac{\gamma}{\kappa} D_{out}$$

where D_{out} is the solution of

$$\frac{\kappa}{\gamma} \sum_{j=1}^{d_{max}} \frac{j P (d_{out} = j)}{1 + j D_{out}} = 1;$$

(ii) for two populations with $\kappa^{(1)} = \kappa^{(2)}, \gamma^{(1)} = \gamma^{(2)}$,

$$d_{out}^{(1)} \leq_{cv} d_{out}^{(2)} \Rightarrow A_{out}^{(1)} \geq A_{out}^{(2)};$$

in particular, A_{out} is maximised when every individual has the same out-degree, $P (d_{out} = \mu) = 1$.

(b) Suppose that $P (d_{out} = \mu) = 1$, so every individual has the same out-degree, and that infectious periods follow an Erlang distribution. Then

(i) $\lim_{N \rightarrow \infty} \frac{\ln \tau}{N} \approx A_{in}$, with

$$A_{in} = \sum_{i=1}^{d_{max}} P (d_{in} = i) \ln (1 + i D_{in}) - \frac{\gamma}{\kappa} D_{in}$$

where D_{in} is the solution of

$$\frac{\kappa}{\gamma} \sum_{j=1}^{d_{max}} \frac{j P (d_{in} = j)}{1 + j D_{in}} = 1;$$

(ii) for two populations with $\kappa^{(1)} = \kappa^{(2)}, s^{(1)} = s^{(2)}, \gamma^{(1)} = \gamma^{(2)}$,

$$d_{in}^{(1)} \leq_{cv} d_{in}^{(2)} \Rightarrow A_{in}^{(1)} \geq A_{in}^{(2)};$$

in particular, A_{in} is maximised when every individual has the same in-degree, $P (d_{in} = \mu) = 1$.

Note that our asymptotic results (Theorem 1) are exact for the model with transition rates given by Table 1, but approximate under the annealed network interpretation.

A third open issue is to allow for more general infectious period distributions in the case of heterogeneous infectivity (or heterogeneous out-degree, under the network interpretation). We conjecture that Theorem 1(i) remains valid, and hence also Theorem 2(i) and Theorem 4(a), if infectious periods are allowed to follow an Erlang, rather than exponential, distribution. Indeed, Fig. 3 suggests that the mean persistence time τ is unchanged when λ , μ are interchanged. A difficulty here is that, in contrast to the case of heterogeneous susceptibilities, we have not been able to find complete solutions $U(\theta)$, $V(y)$ to the relevant Hamilton–Jacobi equations even for the case of exponentially distributed infectious periods, but only, in proving Theorem 1(i), to evaluate $U(\theta)$ along one particular trajectory.

One advantage of simple asymptotic formulae such as provided by Theorem 1, as opposed to the exact formula (3), is that they provide a route to qualitative results such as Theorem 2, that increasing heterogeneity reduces (at least to leading order) the expected persistence time of infection, and in particular that persistence time is maximised in a homogeneous population. Theorem 2 establishes this ordering when heterogeneity is in either infectivity or susceptibility; figure 2 suggests that the result remains true even when both types of heterogeneity are present. It is interesting to compare with the results contained in section 5 of Clancy and Pearce (2013) regarding the effect of such heterogeneities upon the (large population) mean endemic prevalence level $y^* = \sum_{i=1}^k y_i^*$. Theorem 7(i) and theorem 10 of Clancy and Pearce (2013) show, respectively, that heterogeneous infectivity alone has no effect upon the endemic prevalence level y^* , whereas heterogeneous susceptibility alone can only decrease y^* , with $\mu^{(1)} <_f \mu^{(2)} \Rightarrow y^{*(1)} \geq y^{*(2)}$, corresponding to our Theorem 2(ii) for persistence times. When both types of heterogeneity are combined, theorem 8 of Clancy and Pearce (2013) shows that if infectivity and susceptibility are non-negatively correlated ($\sum_{i=1}^k \lambda_i \mu_i f_i \geq 1$) then the endemic prevalence level cannot be greater than for a homogeneous population with the same R_0 value. However, theorem 9 of Clancy and Pearce (2013) together with numerical work shown in figure 3 of Clancy and Pearce (2013) demonstrates that when infectivity and susceptibility are negatively correlated it is possible for the endemic prevalence level to be greater than in the homogeneous case (with R_0 values matched). This presents an interesting contrast to our numerical results in Fig. 2, where heterogeneities were found to always decrease (to leading order) the expected persistence time, regardless of whether λ , μ are positively or negatively correlated.

There is a slightly counter-intuitive aspect to the above results, in that an increase in endemic prevalence level may correspond to a decrease in expected persistence time. This is easily resolved by observing that an increase in prevalence level may be accompanied by a corresponding increase in variability, leading to faster extinction of infection. The effect of heterogeneities upon the variability of the quasi-stationary distribution is studied in section 6 of Clancy and Pearce (2013), via an Ornstein–Uhlenbeck diffusion approximation that leads to a multivariate normal approximation to the quasi-stationary distribution \mathbf{q} . The variability of this approximating normal distribution is then used as a proxy measure of persistence time. This approach seems reasonable in terms of qualitative comparisons between infection models, and is com-

mon in the literature. However, the approach is known to give a very bad numerical approximation to mean persistence time, with incorrect leading-order asymptotic behaviour, due to the failure in the lower tail of the normal approximation to the quasi-stationary distribution (Doering et al. 2005; Clancy and Tjia 2018). The methods of the current paper, in contrast, deal directly with the expected persistence time and yield correct leading-order asymptotic formulae.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

The Hamilton–Jacobi equations (7) amount to two alternative ways of expressing the eigenvector equation (3), retaining only terms of leading order in N in the limit as $N \rightarrow \infty$. We briefly outline the derivation of each equation, referring the reader to Assaf and Meerson (2017) and references therein for full details.

For a process with transition rates of the form (4), Eq. (3) may be written as

$$\sum_{l \in L} \left(q_{x-l} W_l \left(\frac{\mathbf{x}-l}{N} \right) - q_x W_l \left(\frac{\mathbf{x}}{N} \right) \right) = -(\tau N)^{-1} q_x \quad \text{for } \mathbf{x} \in C. \quad (30)$$

Suppose (*ansatz*) that $q_x = \exp(-NV(\mathbf{x}/N) + o(N))$ for some function $V(\cdot)$. Writing $\mathbf{y} = \mathbf{x}/N$, then $V\left(\mathbf{y} - \frac{l}{N}\right) = V(\mathbf{y}) - (l^T/N) \frac{\partial V}{\partial \mathbf{y}} + o(1/N)$ and hence

$$q_{x-l} = \left(\exp \left(l^T \frac{\partial V}{\partial \mathbf{y}} \right) + o(1) \right) q_x.$$

Similarly, $W_l\left(\mathbf{y} - \frac{l}{N}\right) = W_l(\mathbf{y}) + o(1/N)$. Retaining only terms of leading order in N , and assuming that τ is sufficiently large for the right hand side of Eq. (30) to be neglected, then Eq. (30) reduces to

$$\sum_{l \in L} W_l(\mathbf{y}) \left(\exp \left(l^T \frac{\partial V}{\partial \mathbf{y}} \right) - 1 \right) = 0.$$

That is, $H\left(\mathbf{y}, \frac{\partial V}{\partial \mathbf{y}}\right) = 0$.

Next, consider the moment generating function $M(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in C} e^{\boldsymbol{\theta}^T \mathbf{x}} q_x$. Multiplying Eq. (30) by $e^{\boldsymbol{\theta}^T \mathbf{x}}$ and summing over \mathbf{x} we find

$$\sum_{\mathbf{x} \in C} \sum_{l \in L} e^{\boldsymbol{\theta}^T \mathbf{x}} q_x W_l \left(\frac{\mathbf{x}}{N} \right) \left(e^{\boldsymbol{\theta}^T l} - 1 \right) = -(\tau N)^{-1} M(\boldsymbol{\theta}).$$

$$\text{That is, } \sum_{l \in L} \left(e^{\theta^T l} - 1 \right) W_l \left(\frac{1}{N} \frac{\partial}{\partial \theta} \right) M(\theta) = -(\tau N)^{-1} M(\theta). \quad (31)$$

Suppose (*ansatz*) that $M(\theta) = \exp(NU(\theta) + o(N))$ for some function $U(\cdot)$, so that $U(\cdot)$ gives the leading-order term in the cumulant generating function of the quasi-stationary distribution \mathbf{q} . Retaining only terms of leading order in N and assuming that τ is sufficiently large for the right hand side of Eq. (31) to be neglected, then Eq. (31) reduces to

$$\sum_{l \in L} \left(e^{\theta^T l} - 1 \right) W_l \left(\frac{\partial U}{\partial \theta} \right) = 0.$$

That is, $H \left(\frac{\partial U}{\partial \theta}, \theta \right) = 0$.

Note that the Hamilton–Jacobi equations (7) thus have boundary conditions $U(\mathbf{0}) = V(\mathbf{y}^*) = 0$. However, we can add an arbitrary constant to any solution of either Eq. (7) and still have a solution, and any such additive constant will cancel out when we come to estimate $(\ln \tau)/N$ using Eq. (8). We therefore ignore the boundary conditions and instead choose additive constants so as to keep our presented solutions for $U(\theta)$, $V(\mathbf{y})$ as simple as possible.

References

- Andersson H, Djehiche B (1998) A threshold limit theorem for the stochastic logistic epidemic. *J Appl Probab* 35:662–670
- Assaf M, Meerson B (2010) Extinction of metastable stochastic populations. *Phys Rev E* 81:021116
- Assaf M, Meerson B (2017) WKB theory of large deviations in stochastic populations. *J Phys A: Math Theor* 50:263001
- Ball FG, Britton T, Neal P (2016) On expected durations of birth-death processes, with applications to branching processes and SIS epidemics. *J Appl Probab* 53:203–215
- Britton T, Janson S, Martin-Löf A (2007) Graphs with specified degree distributions, simple epidemics, and local vaccination strategies. *Adv Appl Probab* 39:922–948
- Chen N, Olvera-Cravioto M (2013) Directed random graphs with given degree distributions. *Stoch Syst* 3:147–186
- Clancy D (2015) Generality of endemic prevalence formulae. *Math Biosci* 269:30–36
- Clancy D, Pearce CJ (2013) The effect of population heterogeneities upon spread of infection. *J Math Biol* 67:963–987
- Clancy D, Tjia E (2018) Approximating time to extinction for endemic infection models. *Methodol Comput Appl Probab*. <https://doi.org/10.1007/s11009-018-9621-8>
- Darroch JN, Seneta E (1967) On quasi-stationary distributions in absorbing continuous-time finite Markov chains. *J Appl Probab* 4:192–196
- Doering CR, Sargsyan KV, Sander LM (2005) Extinctions times for birth-death processes: exact results, continuum asymptotics, and the failure of the Fokker–Planck approximation. *Multiscale Model Simul* 3:283–299
- Dorogovtsev SN, Goltsev AV, Mendes JFF (2008) Critical phenomena in complex networks. *Rev Mod Phys* 80:1275–1335
- Dykman MI, Mori E, Ross J, Hunt PM (1994) Large fluctuations and optimal paths in chemical kinetics. *J Chem Phys* 100:5735–5750
- Elgart V, Kamenev A (2004) Rare event statistics in reaction–diffusion systems. *Phys Rev E* 70:041106
- Ethier SN, Kurtz TG (2005) *Markov processes: characterization and convergence*. Wiley, New York
- Harris TE (1976) On a class of set-valued Markov processes. *Ann Probab* 4:175–194

-
- Hindes J, Schwartz IB (2016) Epidemic extinction and control in heterogeneous networks. *Phys Rev Lett* 117:028302
- Holley A, Liggett T (1975) Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann Probab* 3:643–663
- Kamenev A, Meerson B (2008) Extinction of an infectious disease: a large fluctuation in a nonequilibrium system. *Phys Rev E* 77:061107
- Lajmanovich A, Yorke JA (1976) A deterministic model for gonorrhea in a nonhomogeneous population. *Math Biosci* 28:221–236
- Lindley BS, Shaw LB, Schwartz IR (2014) Rare-event extinction on stochastic networks. *Europhys Lett* 108:58008
- Marshall AW, Olkin I, Arnold BC (2011) *Inequalities: theory of majorization and its applications*. Springer, Berlin
- Masoero D (2014) A Laplace’s method for series and the semiclassical analysis of epidemiological models. [arXiv:1403.5532](https://arxiv.org/abs/1403.5532)
- Nold A (1980) Heterogeneity in disease-transmission modelling. *Math Biosci* 52:227–240
- Shaked M, Shanthikumar JG (2007) *Stochastic orders*. Springer, Berlin
- Pricer K, Britton T (2015) The configuration model for partially directed graphs. *J Stat Phys* 161:965–985
- van Doorn EA, Pollett PK (1993) Quasi-stationary distributions for discrete-state models. *Eur J Oper Res* 230:1–14
- van Herwaarden OA, Grasman J (1995) Stochastic epidemics: major outbreaks and the duration of the endemic period. *J Math Biol* 6:581–601
- Weiss GH, Dishon M (1971) On the asymptotic behavior of the stochastic and deterministic models of an epidemic. *Math Biosci* 11:261–265
- Wilkinson RR, Sharkey KJ (2013) An exact relationship between invasion probability and endemic prevalence for Markovian SIS dynamics on networks. *PLoS ONE* 8:e69028