Approximating Time to Extinction for Endemic Infection Models

Damian Clancy¹ · Elliott Tjia²

Received: 13 December 2016 /

Accepted: 7 February 2018 / Published online: 24 February 2018

© The Author(s) 2018. This article is an open access publication

Abstract Approximating the time to extinction of infection is an important problem in infection modelling. A variety of different approaches have been proposed in the literature. We study the performance of a number of such methods, and characterise their performance in terms of simplicity, accuracy, and generality. To this end, we consider first the classic stochastic susceptible-infected-susceptible (SIS) model, and then a multi-dimensional generalisation of this which allows for Erlang distributed infectious periods. We find that (i) for a below-threshold infection initiated by a small number of infected individuals, approximation via a linear branching process works well; (ii) for an above-threshold infection initiated at endemic equilibrium, methods from Hamiltonian statistical mechanics yield correct asymptotic behaviour as population size becomes large; (iii) the widely-used Ornstein-Uhlenbeck diffusion approximation gives a very poor approximation, but may retain some value for qualitative comparisons in certain cases; (iv) a more detailed diffusion approximation can give good numerical approximation in certain circumstances, but does not provide correct large population asymptotic behaviour, and cannot be relied upon without some form of external validation (eg simulation studies).

Keywords Stochastic epidemic models · Large deviations · Endemic fade-out · Persistence time

Mathematics Subject Classification (2010) 92D30 · 60J28

✉ 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

² University of Liverpool, Liverpool, England

1 Introduction

In modelling endemic infections, a quantity of particular interest is the persistence time until infection dies out of the population. For the simplest models, it is possible to compute the distribution of this random variable exactly, based on general Markov chain theory. In particular, the expected persistence time may be expressed as the solution of a linear system of equations. However, for large populations and for more complicated models, numerical computation of this exact solution can become very time-consuming, and may also suffer from numerical instability. Further, it is not straightforward to use the exact solution to address qualitative issues such as, for instance, the effect of greater variability in individual infectious periods upon the expected persistence time in the population. Hence it remains of great interest to find simple and accurate approximations.

Our aim is to assess a number of approximation techniques in terms of simplicity, accuracy, and generality. In order to do this, we focus upon relatively simple models, for which it is possible to compute exact values for comparison. Thus in Section 2 we begin by studying the classic susceptible-infected-susceptible (SIS) model of Weiss and Dishon (1971). Whilst a useful simple model to start from, this classic SIS model is not an entirely typical infection model. In particular, it possesses a natural 1-dimensional structure (only the number of infected individuals must be kept track of), whereas most infection models are multi-dimensional (eg one may need to keep track of susceptible, infected, and immune individuals, giving a 3-dimensional model). In Section 3 we therefore extend the SIS model to allow for Erlang-distributed infectious periods, generalising the exponentially distributed infectious periods assumed within the classic SIS model, and resulting in a model which is naturally multi-dimensional, but nevertheless still relatively simple. The models we study have the further advantage that in each case, rigorous asymptotic results are known giving the expected persistence time as population size becomes large. Results for the classic SIS model appear in Andersson and Djehiche (1998), and the extension to the SIS model with general infectious period distribution in Ball et al. (2016). These asymptotic results are derived by techniques which seem unlikely to generalise to broader classes of epidemic models, but provide a basis against which to compare the more general approximation methods that we are chiefly interested in.

We consider four general approximation techniques. First, we consider the below threshold case, when the basic reproduction number R_0 (the average number of new infections caused by a typical infected individual in an otherwise susceptible population) is less than one. Then epidemic threshold theory (Ball 2014), applicable to a wide class of models, tells us that the infectious process may be approximated by a linear branching process, and consequently the mean persistence time of the infection process may be approximated by the mean time to extinction of the branching process.

An alternative, applicable to both the below threshold ($R_0 \leq 1$) and above threshold ($R_0 > 1$) cases, is the approach often referred to as ‘the diffusion approximation’. Results such as those contained in section 11.3 of Ethier and Kurtz (1986) justify approximating the infection model (a Markov jump process) by a diffusion process over finite time intervals, in the limit as population size becomes large. However, in this large population limit, the mean persistence time generally tends to infinity, so that convergence over finite time intervals is not sufficient (see section 11.4 of Ethier and Kurtz (1986)). In spite of this lack of rigorous justification, a number of recent papers have nevertheless used the diffusion approximation to approximate mean persistence times for epidemic models, eg Doubova and Vellido (2016) and Wang et al. (2014).

In the case $R_0 > 1$, a simpler diffusion approximation is possible, in which the diffusion approximation referred to above is linearised about the equilibrium point of the deterministic approximation, leading to an Ornstein-Uhlenbeck diffusion process. This approximation is considerably simpler than the diffusion approximation described previously, and hence has seen widespread use, eg Andersson and Britton (2000), Britton and Neal (2010), Clancy (2005), Clancy and Pearce (2013), Lindholm and Britton (2007), Nåsell (1999), and Nåsell (2002).

Finally, a rather different approach valid for $R_0 > 1$ uses techniques from Hamiltonian statistical mechanics, and has seen much recent attention in the theoretical physics literature (Asaf and Meerson 2010; Dykman et al. 1994; Elgart and Kamenev 2004; Kamenev and Meerson 2008).

2 The Classic Susceptible-Infected-Susceptible (SIS) Model

2.1 Approximation Techniques

We consider first the classic stochastic SIS model introduced by Weiss and Dishon (1971). In this model, infection spreads between members of a closed population of size N consisting at time t of $S(t)$ susceptible and $I(t)$ infected individuals. It is sufficient to focus on $I(t)$, since $S(t) = N - I(t)$ for all $t \geq 0$. The process $\{I(t) : t \geq 0\}$ is a continuous-time Markov chain, with rates of transition given in Table 1, where $\beta, \gamma > 0$ are the infection and recovery rate parameters, respectively.

The basic reproduction number for this model is given by $R_0 = \beta/\gamma$. The state space consists of an absorbing state at $I = 0$ together with the transient communicating class $C = \{1, 2, \dots, N\}$. Consequently, absorption at state 0 will occur within finite time with probability 1, and our interest is in the random variable

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

We define

$$\tau_i = E[T | I(0) = i]$$

to be the expected time to extinction starting from i infected individuals, and write $\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_N)$. Then denoting by Q the transition rate matrix of the process, and by Q^C the rate matrix restricted to the transient states, we have (Norris (1997), theorem 3.3.3) that

$$Q^C \boldsymbol{\tau} = -\mathbf{1}. \quad (1)$$

Equation 1 may be solved numerically for $\boldsymbol{\tau}$. Alternatively, equation (5.28) of Norden (1982) gives the explicit solution

$$\tau_i = \frac{1}{\gamma} \left(\sum_{m=1}^i \sum_{j=m}^N \frac{1}{j} \left(\frac{R_0}{N} \right)^{j-m} \frac{(N-m)!}{(N-j)!} \right). \quad (2)$$

Table 1 Transition rates for the classic SIS model

Event	State transition	Transition rate
Infection of susceptible	$I \rightarrow I + 1$	$(\beta/N)I(N - I)$
Recovery of infected	$I \rightarrow I - 1$	γI

When the infection process is above threshold ($R_0 > 1$) then rather than specifying a fixed initial state $I(0)$, it is arguably of more interest to consider as initial condition a population in which infection has already settled to an endemic level. Since $\{I(t)\}$ is a process on a finite state-space with all non-absorbing states communicating, then (Darroch and Seneta 1967) there exists a unique quasi-stationary distribution $\mathbf{q} = (q_1, q_2, \dots, q_N)$ such that, whatever the initial state of the process,

$$q_i = \lim_{t \rightarrow \infty} \Pr(I(t) = i | I(t) > 0) \text{ for } i = 1, 2, \dots, N.$$

The quasi-stationary distribution $\mathbf{q}$ may be found as the unique solution of

$$\mathbf{q} Q^C = -\lambda \mathbf{q} \text{ with } q_1 + q_2 + \dots + q_N = 1, \quad (3)$$

where $-\lambda$ is the eigenvalue of Q^C with largest real part. The time to extinction from quasi-stationarity is then exponentially distributed with mean $\tau_{\mathbf{q}} = 1/\lambda = 1/\gamma q_1$. Note that the quasi-stationary distribution $\mathbf{q}$ and the mean extinction time $\tau_{\mathbf{q}}$ may both be evaluated for any $R_0 > 0$. However, for $R_0 < 1$ the process is likely to die out before settling to a quasi-stationary state, so that $\mathbf{q}$ and $\tau_{\mathbf{q}}$ are of little practical interest. In contrast, for $R_0 > 1$ the process is likely to settle around the quasi-stationary distribution for a considerable time prior to eventual extinction. Hence for $R_0 > 1$ it is natural to consider a population in which infection is endemic, and to study the time to extinction from quasi-stationarity.

The simplest possible approximation to the above model is the deterministic SIS model, defined as follows. For sufficiently large population size N , the scaled number of infectives $I(t)/N$ may be approximated by the deterministic process $y(t)$ satisfying

$$\frac{dy}{dt} = \beta y(1 - y) - \gamma y. \quad (4)$$

The deterministic system (4) has equilibria at $y = 0$ and $y^* = 1 - (1/R_0)$. For $R_0 < 1$, the former equilibrium (of extinction) is stable, and $y(t) \rightarrow 0$ as $t \rightarrow \infty$ for any initial state $y(0) \in [0, 1]$. For $R_0 > 1$, the endemic equilibrium y^* is stable, with $y(t) \rightarrow y^*$ as $t \rightarrow \infty$ for any initial state $y(0) \in (0, 1]$. In fact, Eq. 4 may be solved explicitly: for $\beta \neq \gamma$, starting from $y(0) = y_0$, we have

$$y(t) = \frac{(\gamma - \beta)y_0 e^{(\beta - \gamma)t}}{\gamma - \beta + \beta(1 - e^{(\beta - \gamma)t})y_0}. \quad (5)$$

Consequently, the process does not become extinct within finite time (for $y_0 > 0$), whatever the value of R_0 . For $R_0 > 1$, since $y(t) \rightarrow y^* > 0$ as $t \rightarrow \infty$, the deterministic process cannot provide any useful approximation for mean extinction time. For $R_0 < 1$, on the other hand, although $y(t) > 0$ for all t , it is at least the case that $y(t) \rightarrow 0$ as $t \rightarrow \infty$. Noting that we are approximating the discrete, integer-valued process $I(t)$ in terms of the continuous process $y(t)$, it is natural to employ a continuity correction; that is, for any initial state with $y_0 \in [1/2N, 1]$ we define

$$\tau^{\text{Det}}(Ny_0) = \inf \{t \geq 0 : Ny(t) \leq 0.5\},$$

and then for $i = 1, 2, \dots, N$ we may suggest $\tau^{\text{Det}}(i)$ as a crude first approximation to τ_i in the subcritical case $R_0 < 1$. Solving Eq. 5 for t with $y(t) = 1/2N$, we find

$$\tau^{\text{Det}}(y) = \frac{1}{\gamma} \left(\frac{1}{1 - R_0} \right) \left(\ln(2y) + \ln \left(\frac{1 - R_0 \left(1 - \frac{1}{2N}\right)}{1 - R_0 \left(1 - \frac{y}{N}\right)} \right) \right) \quad (6)$$

for $0.5 \leq y \leq N$.

The rather crude (first order) deterministic approximation may be improved by considering a (second order) diffusion approximation. That is, the process $I(t)$ may be approximated by the diffusion process $Y(t)$ satisfying

$$dY = \left(\frac{\beta}{N} Y(N - Y) - \gamma Y \right) dt + \sqrt{\frac{\beta}{N} Y(N - Y) + \gamma Y} dW \quad (7)$$

where $W(t)$ is a standard Brownian motion. Since we are again approximating a discrete process with a continuous process, we again apply a continuity correction, and approximate τ_i by $\tau^{\text{Diff}}(i)$, where for $0.5 \leq y \leq N$,

$$\tau^{\text{Diff}}(y) = E[\inf\{t \geq 0 : Y(t) \leq 0.5\} | Y(0) = y].$$

From Gardiner (2009), section 5.5.2, we have that $\tau^{\text{Diff}}(y)$ satisfies the Kolmogorov backward equation

$$\left(\frac{\beta}{N} y(N - y) - \gamma y \right) \frac{\partial \tau^{\text{Diff}}}{\partial y} + \frac{1}{2} \left(\frac{\beta}{N} y(N - y) + \gamma y \right) \frac{\partial^2 \tau^{\text{Diff}}}{\partial y^2} = -1 \quad (8)$$

for $y \in (0.5, N)$, with boundary conditions $\tau^{\text{Diff}}(0.5) = 0$, $\frac{\partial \tau^{\text{Diff}}}{\partial y} \Big|_{y=N} = 0$ (where we have imposed a reflecting boundary at $y = N$). The solution to Eq. 8 may be written explicitly (Gardiner (2009), equation 5.5.24) as

$$\begin{aligned} \tau^{\text{Diff}}(y) &= 2N \int_{1/2N}^{y/N} du \int_u^1 \frac{\exp\left(2N \int_u^z \frac{\beta v(1-v) - \gamma v}{\beta v(1-v) + \gamma v} dv\right)}{\beta z(1-z) + \gamma z} dz \\ &= \frac{2N}{\gamma} \int_{1/2N}^{y/N} du \int_u^1 \frac{\left(\frac{1+R_0(1-z)}{1+R_0(1-u)}\right)^{\frac{4N}{R_0}} e^{2N(z-u)}}{z(1+R_0(1-z))} dz. \end{aligned} \quad (9)$$

The above integral formula for $\tau^{\text{Diff}}(y)$ may be evaluated numerically using Matlab. In the case $R_0 > 1$, we will approximate the time to extinction from quasi-stationarity τ_q using $\tau^{\text{Diff}}(\lfloor Ny^* \rfloor)$, where $\lfloor Ny^* \rfloor$ denotes the integer part of the re-scaled deterministic endemic equilibrium point.

In the supercritical case ($R_0 > 1$), when studying fluctuations around endemic quasi-equilibrium, it is standard (see, for instance, Näsell (1999)) to further approximate the diffusion $Y(t)$ in terms of an Ornstein-Uhlenbeck process centred at the deterministic equilibrium point Ny^* . The drift and diffusion coefficients in Eq. 7 are approximated to leading order around $Y = Ny^*$, yielding the Ornstein-Uhlenbeck process $\tilde{Y}(t)$ satisfying

$$d\tilde{Y} = -(\beta - \gamma) (\tilde{Y} - Ny^*) dt + \sqrt{2N\gamma y^*} dW.$$

The stationary distribution of an Ornstein-Uhlenbeck process with drift coefficient J and local variance $\tilde{B}$, centred at Ny^* , is Gaussian with mean Ny^* and variance V satisfying

$$2JV = -\tilde{B}.$$

In this case we have $J = -(\beta - \gamma)$ and $\tilde{B} = 2N\gamma y^*$, so that $V = N/R_0$. The quasi-stationary distribution q may thus be approximated by a Gaussian distribution with mean Ny^* , variance N/R_0 . Recalling that the mean time to extinction is given by $\tau_q = 1/\gamma q_1$, we have the approximation

$$q_1 \approx \sqrt{\frac{R_0}{2\pi N}} \exp\left(-\frac{R_0}{2N} (Ny^* - 1)^2\right).$$

For large N , recalling that $y^* = 1 - (1/R_0)$ and making the slight further approximation $Ny^* - 1 \approx Ny^*$ gives

$$q_1 \approx \sqrt{\frac{R_0}{2\pi N}} \exp\left(-\frac{(R_0 - 1)^2}{2R_0} N\right).$$

For N large and $R_0 > 1$, the mean time to extinction from quasi-stationarity τ_q may thus be approximated by

$$\tau^{\text{OU}} = \frac{1}{\gamma} \sqrt{\frac{2\pi N}{R_0}} \exp\left(\frac{(R_0 - 1)^2}{2R_0} N\right). \quad (10)$$

The above argument can only be expected to give a reasonable approximation provided the Gaussian approximation to the quasi-stationary distribution assigns negligible probability to negative numbers of infected individuals. A rule of thumb suggested by Näsell (1999) is to require the coefficient of variation of the Gaussian distribution to be at most $1/3$, on the basis that a Gaussian distribution assigns negligible probability to values more than 3 standard deviations from the mean; that is, we require $N > 9R_0/(R_0 - 1)^2$.

Although a number of authors have employed diffusion and Ornstein-Uhlenbeck approximations in studying extinction time for infection models, such approximations are most appropriate in describing moderate deviations from the mean, whereas extinction arises as a result of a large deviation. Consequently, a more appropriate technique may be the methodology from Hamiltonian statistical mechanics described in, for example, Elgart and Kamenev (2004) and Kamenev and Meerson (2008). Defining $M(\theta; t) = E[e^{\theta I(t)}]$ to be the moment generating function of $I(t)$, then the Kolmogorov forward equations for the process may be written as

$$\frac{dM}{dt} = \beta(e^\theta - 1) \left(\frac{\partial M}{\partial \theta} - \frac{1}{N} \frac{\partial^2 M}{\partial \theta^2} \right) + \gamma(e^{-\theta} - 1) \frac{\partial M}{\partial \theta}. \quad (11)$$

In terms of the cumulant generating function $K(\theta; t) = \ln M(\theta; t)$, Eq. 11 is equivalent to

$$\frac{dK}{dt} = \beta(e^\theta - 1) \left(\frac{\partial K}{\partial \theta} - \frac{1}{N} \left(\frac{\partial^2 K}{\partial \theta^2} + \left(\frac{\partial K}{\partial \theta} \right)^2 \right) \right) + \gamma(e^{-\theta} - 1) \frac{\partial K}{\partial \theta}. \quad (12)$$

Supposing (*ansatz*) that $K(\theta; t) = NS(\theta; t) + o(N)$ for some function $S(\theta; t)$, then Eq. 12 becomes

$$\frac{dS}{dt} = \beta(e^\theta - 1) \frac{\partial S}{\partial \theta} \left(1 - \frac{\partial S}{\partial \theta} \right) + \gamma(e^{-\theta} - 1) \frac{\partial S}{\partial \theta} + o(1).$$

To leading order, this has the form of a Hamilton-Jacobi equation corresponding to the Hamiltonian

$$H(y, \theta; t) = \beta y(1 - y)(e^\theta - 1) + \gamma y(e^{-\theta} - 1). \quad (13)$$

In this formulation, y represents the scaled number of infected individuals, and θ represents (in the terminology of Hamiltonian mechanics) a conjugate momentum variable. Note that $\left. \frac{\partial S}{\partial \theta} \right|_{\theta=0} = E[I/N] + o(1)$.

References (Elgart and Kamenev 2004; Kamenev and Meerson 2008) formulate the problem in terms of the probability generating function; we find the above formulation in terms of the cumulant generating function to be more natural. An alternative approach is to work directly with the Kolmogorov forward equations in terms of the quasi-stationary probabilities q_i , see for example Dykman et al. (1994). The same Hamiltonian (13) is thus obtained.

The Hamiltonian describes a deterministic motion in (y, θ) space that follows the equations of motion

$$\frac{dy}{dt} = \frac{\partial H}{\partial \theta}, \quad \frac{d\theta}{dt} = -\frac{\partial H}{\partial y}.$$

For the specific Hamiltonian (13) the equations of motion are therefore

$$\left. \begin{aligned} \frac{dy}{dt} &= \beta y(1-y)e^\theta - \gamma y e^{-\theta}, \\ \frac{d\theta}{dt} &= -\beta(1-2y)(e^\theta - 1) - \gamma(e^{-\theta} - 1). \end{aligned} \right\} \quad (14)$$

Trajectories of the deterministic SIS model (4) may be obtained by taking $\theta = 0$ in Eq. 14. The system (14) possesses classical equilibrium points at $(y, \theta) = (0, 0)$ and $(y, \theta) = (y^*, 0)$, together with a non-classical disease-free equilibrium point at $(y, \theta) = (0, -\ln(R_0))$.

The expected time to extinction starting from the deterministic equilibrium point satisfies (see Elgart and Kamenev (2004)) $\ln(\tau_{[Ny^*]})/N \rightarrow A$ as $N \rightarrow \infty$, where A is known as the ‘action’. The value of A is found by integrating along the trajectory of Eq. 14 that goes from initial state $(y, \theta) = (y^*, 0)$ at time $t = -\infty$ to final state $(y, \theta) = (0, -\ln(R_0))$ at time $t = +\infty$, with

$$A = \int_{-\infty}^{\infty} \theta \frac{dy}{dt} dt.$$

Now $H(y, \theta)$ is a constant of the motion (sometimes referred to as the ‘energy’ of the system), and $H(y^*, 0) = 0$, so along this trajectory we have

$$\gamma y(e^{-\theta} - 1) = \beta y(1-y)(1 - e^\theta).$$

It follows that $y = 0$, or $\theta = 0$, or $\theta = -\ln(R_0(1-y))$. Along the relevant trajectory we do not have $y = 0$ or $\theta = 0$ except at the endpoints, and so

$$A = \int_{y^*}^0 \theta dy = \int_0^{y^*} \ln(R_0(1-y)) dy = (1/R_0) - 1 + \ln R_0.$$

For $R_0 > 1$ and N large, using $\tau_{[Ny^*]}$ to approximate τ_q , we therefore have that $\ln(\tau_q)/N$ may be approximated by $\ln(\tau_H)/N$, where

$$\tau_H = \exp(N((1/R_0) - 1 + \ln R_0)). \quad (15)$$

The methods outlined above are quite general, and can potentially be applied to any model for transmission of infection (or other population processes). More precise and rigorous results specific to the classic SIS model are presented in Kryscio and Lefèvre (1989), Andersson and Djehiche (1998), and Doering et al. (2005). The relevant results that we will use for comparison with the general approximation methods are as follows.

- (i) For $R_0 < 1$ and $I(0) = i$ fixed, then as $N \rightarrow \infty$ the mean extinction time τ_i converges to the extinction time of a linear birth-death process with birth and death rate parameters β, γ , respectively (Andersson and Djehiche 1998). That is, $\tau_i \rightarrow \tau_i^{\text{Lin}}$, where from equation (7.10), chapter 4, of Karlin and Taylor (1975), we find

$$\tau_i^{\text{Lin}} = \frac{1}{\gamma} \left(\frac{1}{1-R_0} \right) \left((1-R_0^{-i}) \ln(1-R_0) + \sum_{m=1}^{i-1} \frac{1-R_0^{m-i}}{m} \right). \quad (16)$$

- (ii) For $R_0 < 1$ and $i/N = I(0)/N \rightarrow y_0 \in (0, 1]$ as $N \rightarrow \infty$, then as $N \rightarrow \infty$ we have $\tau_i - \tau^{\text{DSS}}(Ny_0) = O(1)$, where (Doering et al. 2005)

$$\tau^{\text{DSS}}(y) = \frac{1}{\gamma} \left(\frac{1}{1 - R_0} \right) \ln y. \quad (17)$$

- (iii) For $R_0 > 1$, the expected time to extinction starting from quasi-stationarity satisfies $\tau_q \sim \tau^{\text{AD}}$, where (Andersson and Djehiche 1998)

$$\tau^{\text{AD}} = \frac{1}{\gamma} \sqrt{\frac{2\pi}{N}} \frac{R_0}{(R_0 - 1)^2} \exp(N((1/R_0) - 1 + \ln R_0)). \quad (18)$$

Result (i) for the case $R_0 < 1$ is based on the epidemic threshold theorem, and can be generalised to a wide range of infection models. In contrast, the derivation of result (iii) for the case $R_0 > 1$ relies upon the 1-dimensional structure of the model, so that whilst similar results may apply to more general birth-death processes, it does not seem straightforward to apply such methods to more general, multi-dimensional, infection models. Doering et al. (2005), by different arguments, confirmed the formula (18) for the case $R_0 > 1$. In the case $R_0 < 1$ with $I(0)/N \rightarrow y_0 \in (0, 1]$ as $N \rightarrow \infty$, however, formula (17) of Doering et al. (2005) is in disagreement with a corresponding result from Kryscio and

Fig. 1 Expected persistence time for the classic SIS model from any initial state $I(0)$, together with approximations. Parameter values $N = 100$, $\beta = 0.8$, $\gamma = 1$ (so that $R_0 = 0.8$). Exact values computed from formula (2); KL asymptotic formula (19); linear birth-death process approximation (16); deterministic approximation (6); diffusion approximation (9)

Lefèvre (1989), cited in Andersson and Djehiche (1998). For completeness, we include this alternative result, which states that $\tau_i - \tau^{\text{KL}}(Ny_0) \rightarrow 0$ with

$$\tau^{\text{KL}}(y) = \frac{1}{\gamma} \left(\frac{\ln(y) + \ln(1 - R_0(1 - (y/N))) + \gamma^E}{1 - R_0(1 - (y/N))} \right), \quad (19)$$

where γ^E denotes Euler's constant $\gamma^E \approx 0.577216$.

2.2 Numerical Results

2.2.1 Below Threshold Case $R_0 < 1$

For all numerical work we fix $\gamma = 1$, since varying γ for fixed R_0 simply amounts to a scaling of the time axis.

Figure 1 illustrates how our various approximations compare with the exact mean persistence time τ_i from any initial state $i \in \{1, 2, \dots, N\}$ for relatively small population size $N = 100$, with $R_0 = 0.8$. We see that the deterministic approximation τ^{Det} is, as one might expect, least accurate. The diffusion approximation τ^{Diff} seems to perform best of all our approximations here. The approximation τ^{Lin} performs very well for small initial infective numbers i , but much less well as i becomes large. Again this is to be expected, since this

Fig. 2 Expected persistence time for the classic SIS model starting with $I(0) = 1$ initial infective in a population of size N , together with approximations. Parameter values $\beta = 0.8$, $\gamma = 1$ (so that $R_0 = 0.8$). Exact values computed from formula (2); linear birth-death process approximation (16); deterministic approximation (6); diffusion approximation (9)

approximation is based on an assumption that the number of available susceptible individuals remains approximately equal to N at all times. The approximation τ^{KL} performs well in the range $5 \leq i \leq 20$, but increasingly badly as i increases, and actually produces negative answers for $i = 1, 2$. It is also notable that τ^{KL} , in contrast to the other approximations, is not monotonically increasing in i , whereas the exact solution (2) clearly is.

Next, we consider the performance of our approximations across a range of population sizes N , up to the moderately large size $N = 1000$. There are essentially two initial conditions of interest — either a fixed number of initial infectives, or a fixed fraction of the population initially infected. Figure 2 shows the performance of relevant approximations for the case of a fixed number of initial infectives, specifically $i = 1$. We see that the approximation τ^{Lin} , whose value is independent of N , performs well even for small population sizes N , and very well for larger N values. Neither τ^{Det} nor τ^{Diff} perform nearly as well, and they do not improve as N increases, which is not surprising, since both the deterministic and the diffusion process assume that the number of infective individuals is large enough to be treated as a continuous variable. For these parameter values, τ^{KL} takes negative values, and so is not shown.

For the case of a fixed initially infected fraction y_0 , Fig. 3 shows expected persistence time plotted against $\ln(N)$. The best approximation here is τ^{Diff} , with τ^{Det} , τ^{Lin} , τ^{KL} all

Fig. 3 Expected persistence time for the classic SIS model plotted against $\ln(N)$ starting with $I(0) = \lfloor 0.3N \rfloor$ initial infectives in a population of size N , together with approximations. Parameter values $\beta = 0.8$, $\gamma = 1$ (so that $R_0 = 0.8$). Exact values computed from formula (2); KL asymptotic formula (19); linear birth-death process approximation (16); deterministic approximation (6); diffusion approximation (9)

performing rather poorly. In terms of asymptotic (large N) behaviour, formula (17) of Doering et al. (2005) suggests that, for large N , we should observe a straight line of gradient $(1/\gamma)(1 - R_0)^{-1}$, whereas formula (19) of Kryscio and Lefèvre (1989) predicts a line of gradient $(1/\gamma)(1 - R_0(1 - y_0))^{-1}$. From Fig. 3 it is clear that formula (19), τ^{KL} , does not predict the correct asymptotic gradient. The approximations τ^{Det} , τ^{Diff} , τ^{Lin} all appear to display the correct asymptotic gradient, but with different offsets, with τ^{Diff} performing best. Note that the deterministic approximation (6) predicts the same asymptotic gradient as Doering et al. (2005).

Overall, our numerical results for the below threshold case suggest that: (i) the approximation τ^{KL} of Kryscio and Lefèvre (1989) gives the wrong asymptotic behaviour, and overall does not seem a reliable approximation; (ii) the formulae τ^{DSS} of Doering et al. (2005) and τ^{Det} give the correct first-order asymptotic behaviour, but τ^{DSS} differs from the exact answer by an unspecified offset, whilst τ^{Det} differs by a substantial constant offset, so that neither provides a useful numerical approximation; (iii) in the case of an initial small trace of infection, the best approximation is that derived from an approximating linear birth-death process, τ^{Lin} ; (iv) in the case of a substantial fraction of the population being initially infected, the best approximation is that derived from an approximating diffusion process, τ^{Diff} .

2.2.2 Above Threshold Case $R_0 > 1$

In the above threshold case, we focus upon the mean persistence time starting from endemicity; that is, we are interested in approximating τ_q .

Figure 4 shows the values of τ_q and relevant approximations for a system not far above threshold ($R_0 = 1.1$). The best approximation here is provided by the diffusion process, τ^{Diff} . Note that for the diffusion approximation, we take initial condition $Y(0) = Ny^*$, ie starting from the endemic equilibrium point of the deterministic system. The Ornstein-Uhlenbeck approximation is seen to perform very poorly, and to become even less accurate as N increases. The approximation τ^{AD} performs reasonably well for $N \geq 100$, although not as well as τ^{Diff} .

In Fig. 5 we consider a system well above threshold ($R_0 = 1.5$). We see that τ^{AD} gives the correct leading-order asymptotic behaviour; τ^{Diff} underestimates the mean persistence time; and τ^{OU} provides a very poor approximation. Figure 5 reproduces elements of figure 6 of Doering et al. (2005), although Doering et al. (2005) did not consider the Ornstein-Uhlenbeck approximation, and computed mean time to extinction from a fixed state close to the endemic equilibrium Ny^* , rather than from quasi-stationarity. There are also two differences in how the diffusion approximation is computed: firstly, we compute τ^{Diff} exactly using formula (9), whereas Doering et al. (2005) use an asymptotic approximation; secondly, our formula (9) employs a continuity correction, stopping the diffusion at $Y = 0.5$, whereas in Doering et al. (2005) the diffusion is simply absorbed at the boundary $Y = 0$. However, these refinements do not make any appreciable difference to the values plotted.

The asymptotic result derived by Doering et al. (2005) for the diffusion approximation (absorbed at $Y = 0$) is that, for $R_0 > 1$, the expected persistence time $\tau^{\text{Diff}}(y)$ from any initial state $y = O(N)$ satisfies, as $N \rightarrow \infty$,

$$\tau^{\text{Diff}}(y) \sim \frac{1}{\gamma} \sqrt{\frac{2\pi}{N}} \frac{R_0(R_0 + 1)}{2(R_0 - 1)^2 \sqrt{R_0}} \exp \left(N \left(2 \left(\frac{R_0 - 1}{R_0} \right) + \frac{4}{R_0} \ln \left(\frac{2}{R_0 + 1} \right) \right) \right). \quad (20)$$

Fig. 4 Expected persistence time for the classic SIS model starting from quasi-stationarity (τ_q) in a population of size N , together with approximations. Parameter values $\beta = 1.1$, $\gamma = 1$ (so that $R_0 = 1.1$ and $y^* \approx 0.0909$). Exact values computed from formula (3); AD asymptotic formula (18); diffusion approximation (9); Ornstein-Uhlenbeck approximation (10)

As pointed out by Doering et al. (2005), comparison of formulae (20) and (18) shows that the diffusion approximation does not give the correct asymptotic behaviour of $\ln(\tau)/N$, confirming what we see in Fig. 5. To look more closely at the failure of the diffusion approximation, in Fig. 6 we plot the exponential constants from formulae (18, 20). We see that as R_0 increases, leading order behaviour of τ^{Diff} becomes an increasingly poor approximation to the true leading order behaviour of τ_q .

It is similarly clear from formula (10) that τ^{OU} does not give correct leading-order asymptotic behaviour. On the other hand, the Hamiltonian approach, formula (15), does yield correct leading-order asymptotic behaviour.

In summary, for R_0 slightly above 1 and moderate population size N , the best approximation appears to be provided by the diffusion approximation. However, the diffusion approximation has the wrong leading-order asymptotic behaviour as $N \rightarrow \infty$, and gets worse as R_0 increases. Hence this approximation can only be used provided that N is not too large and R_0 is not too far above 1. This is in line with comments of Doering et al. (2005). Correct leading-order asymptotic behaviour may be found using the Hamiltonian approach. However, this approach does not provide a useful numerical approximation, since the result we obtain is that $\tau_q \sim C(N, R_0)\tau_H$ with $C(N, R_0) = o(e^N)$ as $N \rightarrow \infty$. Since the prefactor $C(N, R_0)$ is unknown, we cannot use this asymptotic result to evaluate a numerical approximation to τ_q . The Ornstein-Uhlenbeck approximation, though widely used, is seen to provide an extremely poor approximation to τ_q .

Fig. 5 $\ln(\tau_q)/N$, where τ_q is the expected persistence time for the classic SIS model starting from quasi-stationarity in a population of size N , together with approximations. Parameter values $\beta = 1.5$, $\gamma = 1$ (so that $R_0 = 1.5$ and $y^* \approx 0.3333$). Exact values computed from formula (3); AD asymptotic formula (18); diffusion approximation (9); Ornstein-Uhlenbeck approximation (10)

3 Generalising the Infectious Period Distribution

3.1 Approximation Techniques

The classic SIS model makes the implicit assumption that each individual's infectious period follows an exponential distribution. This is purely a mathematical convenience, not motivated by biological realism. We can improve the biological realism of the model by allowing infectious periods to follow an Erlang distribution, using the 'method of stages'. That is, when an individual becomes infected, it passes through k infectious stages, remaining in each stage for an exponentially distributed time of mean $1/k\gamma$, before returning to susceptibility. Thus the total infectious period follows an Erlang distribution, and the process can be modelled as a Markov chain $\{(I_1(t), I_2(t), \dots, I_k(t)) : t \geq 0\}$, where $I_m(t)$ is the number of individuals in infectious stage m at time t , and the number of susceptible individuals at time t is $S(t) = N - \sum_{m=1}^k I_m(t)$. Transition rates are given in Table 2.

The state space is $\{(i_1, i_2, \dots, i_k) \in (\mathbb{Z}_+)^k : i_1 + i_2 + \dots + i_k \leq N\}$, and consists of a single absorbing state at $\mathbf{I} = (I_1, I_2, \dots, I_k) = \mathbf{0}$ together with a transient communicating class C . The basic reproduction number is $R_0 = \beta/\gamma$ as before. The mean infectious period is $1/\gamma$ as before, but the variance of the infectious period is now $1/k\gamma^2$ and thus may be varied according to our choice of k .

We will consider only the case $R_0 > 1$, so that infection may become endemic in the population, and study the time to extinction from quasi-stationarity. Denoting by $\mathbf{q}^k$ the quasi-stationary distribution, then the mean extinction time is given by $\tau_{\mathbf{q}^k} = 1/\lambda_k$ where $-\lambda_k$ is the eigenvalue with largest real part of Q_k^C , the transition rate matrix restricted to the

Fig. 6 Exponential constants appearing in the asymptotic formula τ^{AD} of Andersson and Djehiche (1998) for the SIS model and in the asymptotic formula (20) of Doering et al. (2005) for the diffusion approximation

non-absorbing states. The state-space is now of size $\binom{N+k}{k}$, so that finding the mean time to extinction from Eq. 1 amounts to evaluating an eigenvalue of a sparse non-symmetric square matrix of dimension $\binom{N+k}{k} - 1$.

The process $\mathbf{I}(t)/N$ may be approximated in the large-population limit by the deterministic process $\mathbf{y} = (y_1, y_2, \dots, y_k)$ satisfying

$$\begin{aligned} \frac{dy_1}{dt} &= \beta \left(1 - \sum_{m=1}^k y_m \right) \left(\sum_{m=1}^k y_m \right) - k\gamma y_1, \\ \frac{dy_m}{dt} &= k\gamma y_{m-1} - k\gamma y_m \text{ for } m = 2, 3, \dots, k. \end{aligned}$$

Table 2 Transition rates for the SIS model with k infectious stages (Erlang distributed infectious periods)

Event	State transition	Transition rate
Infection of susceptible	$I_1 \rightarrow I_1 + 1$	$(\beta/N) \left(\sum_{m=1}^k I_m \right) \left(N - \sum_{m=1}^k I_m \right)$
Transition to next infectious stage	$(I_{m-1}, I_m) \rightarrow (I_{m-1} - 1, I_m + 1)$	$k\gamma I_{m-1}$ for $m = 2, 3, \dots, k$
Recovery of infected	$I_k \rightarrow I_k - 1$	$k\gamma I_k$

This system has two equilibria: the disease-free equilibrium $\mathbf{y} = \mathbf{0}$ and an endemic equilibrium $\mathbf{y}^* = (1 - (1/R_0))(\mathbf{1}/k)$, where $\mathbf{1}$ denotes the vector with all entries equal to 1.

The (second-order) diffusion approximation is the process $\mathbf{Y}(t)$ satisfying

$$\begin{aligned} dY_1 &= \left(\frac{\beta}{N} \left(\sum_{m=1}^k Y_m \right) \left(N - \sum_{m=1}^k Y_m \right) - k\gamma Y_1 \right) dt \\ &\quad + \sqrt{\frac{\beta}{N} \left(N - \sum_{m=1}^k Y_m \right) \left(\sum_{m=1}^k Y_m \right)} dW_0 - \sqrt{k\gamma Y_1} dW_1, \\ dY_m &= k\gamma (Y_{m-1} - Y_m) dt + \sqrt{k\gamma Y_{m-1}} dW_{m-1} - \sqrt{k\gamma Y_m} dW_m \text{ for } m = 2, 3, \dots, k, \end{aligned}$$

where $W_0, W_1, W_2, \dots, W_k$ are independent standard Brownian motions.

Defining

$$\tau^{\text{Diff}}(\mathbf{y}) = E \left[\inf \left\{ t \geq 0 : \max_{1 \leq m \leq k} \{Y_m(t)\} \leq 0.5 \right\} \middle| \mathbf{Y}(0) = \mathbf{y} \right]$$

then $\tau^{\text{Diff}}(\mathbf{y})$ satisfies the Kolmogorov backward equation

$$\sum_i A_i(\mathbf{y}) \frac{\partial \tau^{\text{Diff}}(\mathbf{y})}{\partial y_i} + \frac{1}{2} \sum_{i,j} B_{ij}(\mathbf{y}) \frac{\partial^2 \tau^{\text{Diff}}(\mathbf{y})}{\partial y_i \partial y_j} = -1 \text{ in } \Omega, \quad (21)$$

where the state-space is $\Omega = \{\mathbf{y} \in (\mathbb{R}_+)^k : \max \{y_1, y_2, \dots, y_k\} \geq 0.5, y_1 + y_2 + \dots + y_k \leq N\}$, and the coefficients are

$$\begin{aligned} A_1(\mathbf{y}) &= \frac{\beta}{N} \left(\sum_{m=1}^k y_m \right) \left(N - \sum_{m=1}^k y_m \right) - k\gamma y_1, \\ A_i(\mathbf{y}) &= k\gamma (y_{i-1} - y_i) & i = 2, 3, \dots, k, \\ B_{11}(\mathbf{y}) &= \frac{\beta}{N} \left(\sum_{m=1}^k y_m \right) \left(N - \sum_{m=1}^k y_m \right) + k\gamma y_1, \\ B_{ii}(\mathbf{y}) &= k\gamma (y_{i-1} + y_i) & i = 2, 3, \dots, k, \\ B_{i,i-1}(\mathbf{y}) &= B_{i-1,i}(\mathbf{y}) = -k\gamma y_{i-1} & i = 2, 3, \dots, k, \\ B_{ij} &= 0 & j \notin \{i-1, i, i+1\}. \end{aligned}$$

The process has an absorbing boundary at $S_a = \{\mathbf{y} \in \Omega : \max \{y_1, y_2, \dots, y_k\} = 0.5\}$, the remainder of the boundary, denoted S_r , being taken to be reflecting. Denoting by $\mathbf{n} = (n_1, n_2, \dots, n_k)$ a vector normal to the boundary S_r , the boundary conditions (Gardiner 2009, section 6.6) are

$$\begin{aligned} \tau^{\text{Diff}}(\mathbf{y}) &= 0 \text{ on } S_a, \\ \sum_{i,j} n_i B_{ij}(\mathbf{y}) \frac{\partial \tau^{\text{Diff}}(\mathbf{y})}{\partial y_j} &= 0 \text{ on } S_r. \end{aligned}$$

In this case, the boundary conditions on S_r reduce to

$$\begin{aligned}\frac{\partial \tau^{\text{Diff}}}{\partial y_1} &= 0 \text{ for } \mathbf{y} \in S_r \text{ with } y_1 = 0, \\ \frac{\partial \tau^{\text{Diff}}}{\partial y_i} - \frac{\partial \tau^{\text{Diff}}}{\partial y_{i-1}} &= 0 \text{ for } \mathbf{y} \in S_r \text{ with } y_i = 0 \quad (i = 2, 3, \dots, k), \\ \frac{\partial \tau^{\text{Diff}}}{\partial y_k} &= 0 \text{ for } \mathbf{y} \in S_r \text{ with } y_1 + y_2 + \dots + y_k = N.\end{aligned}$$

The above system is not uniformly elliptic on Ω , since the matrix $B(\mathbf{y})$ with entries $B_{ij}(\mathbf{y})$ has an eigenvalue of zero at points on the boundary. Consequently, we have no rigorous proof that a unique solution exists. Nevertheless, since our main interest is in approximations that work in practice, we will go ahead with investigating the performance of this diffusion approximation numerically. Specifically, we apply the finite element method (see, for example, Reddy 2006) using Freefem++ software (Freefem++ 2016).

In order to apply the finite element method, the partial differential equation must first be put into variational form, as follows. We start from the integral formulation

$$\int_{\Omega} \left(\sum_i A_i(\mathbf{y}) \frac{\partial \tau^{\text{Diff}}(\mathbf{y})}{\partial y_i} + \frac{1}{2} \sum_{i,j} B_{ij}(\mathbf{y}) \frac{\partial^2 \tau^{\text{Diff}}}{\partial y_i \partial y_j} \right) w(\mathbf{y}) d\Omega = - \int_{\Omega} w(\mathbf{y}) d\Omega$$

for appropriate test functions $w(\mathbf{y})$. Integrating by parts in order to eliminate second order terms (treating mixed derivatives $\frac{\partial^2}{\partial y_i \partial y_j}$ symmetrically), applying boundary conditions to the resulting boundary integral terms, and noting that the condition $\tau^{\text{Diff}}(\mathbf{y}) = 0$ on S_a implies that relevant test functions can be taken to satisfy $w(\mathbf{y}) = 0$ on S_a , then all boundary integral terms vanish, and we obtain the variational form

$$\begin{aligned}& \frac{1}{2} \int_{\Omega} \sum_{i,j} \left(B_{ij}(\mathbf{y}) \frac{\partial \tau^{\text{Diff}}}{\partial y_i} \frac{\partial w}{\partial y_j} + w(\mathbf{y}) \frac{\partial B_{ij}}{\partial y_j} \frac{\partial \tau^{\text{Diff}}}{\partial y_i} \right) d\Omega \\ & - \int_{\Omega} \sum_i A_i(\mathbf{y}) w(\mathbf{y}) \frac{\partial \tau^{\text{Diff}}}{\partial y_i} d\Omega = \int_{\Omega} w(\mathbf{y}) d\Omega.\end{aligned}\quad (22)$$

Freefem++ code to solve the above system in the case $k = 2$ is given in the [Appendix](#).

Approximating the drift and local variance matrices of the diffusion approximation to leading order around $N\mathbf{y}^*$ yields a multivariate Ornstein-Uhlenbeck approximation $\tilde{\mathbf{Y}}(t)$, centred at $N\mathbf{y}^*$, with drift matrix J and local variance matrix $\tilde{B}$ given by

$$J = \begin{pmatrix} 2\gamma - \beta - k\gamma & 2\gamma - \beta & 2\gamma - \beta & \cdots & 2\gamma - \beta \\ k\gamma & -k\gamma & 0 & \cdots & 0 \\ 0 & k\gamma & -k\gamma & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -k\gamma \end{pmatrix},$$

$$\tilde{B} = N\gamma \left(1 - \frac{1}{R_0} \right) \begin{pmatrix} 2 & -1 & 0 & 0 & \cdots & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 \\ 0 & -1 & 2 & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 2 \end{pmatrix}.$$

The stationary distribution of this Ornstein-Uhlenbeck process is multivariate Gaussian with mean $N\mathbf{y}^*$, variance matrix V satisfying

$$JV + V^T J = -\tilde{B}. \quad (23)$$

The solution to Eq. 23 is

$$V = \frac{N}{k} \left(1 - \frac{1}{R_0} \right) I_k + \frac{N}{k^2} \left(\frac{2}{R_0} - 1 \right) \mathbb{1}_k,$$

where I_k is the $k \times k$ identity matrix and $\mathbb{1}_k$ represents the $k \times k$ matrix with all entries equal to 1.

Under the Ornstein-Uhlenbeck approximation, the total number of infectives $\tilde{Y}_1 + \tilde{Y}_2 + \dots + \tilde{Y}_k$ has a Gaussian stationary distribution with mean $N\mathbf{y}^*$ and variance N/R_0 . That is, the Ornstein-Uhlenbeck approximation suggests that the quasi-stationary distribution of the total number of infected individuals does not depend upon k , the number of infected stages. Further, under the Ornstein-Uhlenbeck approximation $\tilde{Y}_1, \tilde{Y}_2, \dots, \tilde{Y}_k$ are exchangeable random variables, in stationarity. Hence we can approximate

$$q_{e_k}^k \approx \frac{q_1}{k}$$

where e_k denotes the unit vector with k th element equal to 1, $q_{e_k}^k$ is the quasi-stationary probability of being in state e_k , and q_1 is the quasi-stationary probability for the classic ($k = 1$) SIS model to be in state $I = 1$. The mean time to extinction τ_{q^k} is then approximated as

$$\tau_{q^k} = \frac{1}{k\gamma q_{e_k}^k} \approx \frac{1}{\gamma q_1} \approx \tau^{\text{OU}},$$

where τ^{OU} is given by equation (10), exactly as for the classic SIS model with exponentially distributed infectious periods.

Moving on to the Hamiltonian approach, our SIS model with Erlang distributed infectious periods has Hamiltonian

$$\begin{aligned} H_k(\mathbf{y}, \boldsymbol{\theta}) = & \beta \left(\sum_{m=1}^k y_m \right) \left(1 - \sum_{m=1}^k y_m \right) (e^{\theta_1} - 1) \\ & + k\gamma \sum_{m=1}^{k-1} y_m (e^{-\theta_m + \theta_{m+1}} - 1) + k\gamma y_k (e^{-\theta_k} - 1), \end{aligned}$$

describing a motion in $\mathbb{R}^{2k}$ that follows the equations of motion

$$\begin{aligned} \frac{dy_i}{dt} &= \frac{\partial H_k}{\partial \theta_i} \text{ for } i = 1, 2, \dots, k, \\ \frac{d\theta_i}{dt} &= -\frac{\partial H_k}{\partial y_i} \text{ for } i = 1, 2, \dots, k. \end{aligned}$$

Setting $\theta_0 = \theta_{k+1} = y_0 = 0$, these equations of motion may be written as

$$\left. \begin{aligned} \frac{dy_i}{dt} &= \beta \left(\sum_{m=1}^k y_m \right) \left(1 - \sum_{m=1}^k y_m \right) e^{\theta_1} \delta_{i1} + k\gamma (y_{i-1} e^{-\theta_{i-1} + \theta_i} - y_i e^{-\theta_i + \theta_{i+1}}) \\ &\quad \text{for } i = 1, 2, \dots, k, \\ \frac{d\theta_i}{dt} &= -\beta \left(1 - 2 \sum_{m=1}^k y_m \right) (e^{\theta_1} - 1) - k\gamma (e^{-\theta_i + \theta_{i+1}} - 1) \text{ for } i = 1, 2, \dots, k, \end{aligned} \right\} \quad (24)$$

where δ_{ij} is the Kronecker delta. This system has classical equilibrium points at $(\mathbf{y}, \boldsymbol{\theta}) = (\mathbf{0}, \mathbf{0})$ and $(\mathbf{y}^*, \mathbf{0})$, together with a non-classical disease-free equilibrium point at $(\mathbf{0}, \boldsymbol{\theta}^*)$, where $\boldsymbol{\theta}^* = (k, k-1, \dots, 3, 2, 1)\theta_k^*$ with θ_k^* satisfying

$$\frac{\beta}{k\gamma} \left(\sum_{m=1}^k e^{m\theta_k^*} \right) - 1 = 0. \quad (25)$$

The function $f(z) = \sum_{m=1}^k z^m$ is increasing, with $f(0) = 0$ and $f(z) \rightarrow \infty$ as $z \rightarrow \infty$, so there is a unique positive z^* satisfying $f(z^*) = k\gamma/\beta$, and then the unique real solution of Eq. 25 is $\theta_k^* = \ln z^*$. Also, for $R_0 > 1$ we have $f(z^*) < k$, and since $f(1) = k$ it follows that $z^* < 1$, so $\theta_k^* < 0$.

We may approximate $\ln(\tau_q)/N$ by $\ln(\tau_{H_k})/N$, where

$$\tau_{H_k} = \exp(N A_k) \quad (26)$$

with action A_k given by

$$A_k = \int_{-\infty}^{\infty} \sum_{m=1}^k \theta_k \frac{dy_k}{dt} dt, \quad (27)$$

the integration being along a trajectory going from $(\mathbf{y}^*, \mathbf{0})$ at $t = -\infty$ to $(\mathbf{0}, \boldsymbol{\theta}^*)$ at $t = +\infty$.

Solving the $2k$ -dimensional system of ordinary differential Eq. 24 numerically for the case $k = 2$, subject to boundary conditions at $t = -\infty$ and $t = +\infty$, and then using the solution $(\mathbf{y}(t), \boldsymbol{\theta}(t))$ thus obtained to evaluate A_k , it appears that $A_2 = A = (1/R_0) - 1 + \ln R_0$. Now formula (27) is equivalent to $A_k = -S_k(\boldsymbol{\theta}^*)$, where the function $S_k(\boldsymbol{\theta})$ satisfies the partial differential equation $H_k\left(\frac{\partial S_k}{\partial \boldsymbol{\theta}}, \boldsymbol{\theta}\right) = 0$ with $S_k(\mathbf{0}) = 0$. From the form of the equilibrium point $\boldsymbol{\theta}^*$, together with the conjecture that $A_k = A$, it is possible to guess the solution to be

$$S_k(\boldsymbol{\theta}) = \ln\left(\frac{\sum_{m=1}^k e^{\theta_m}}{k}\right) - \frac{\gamma}{\beta} \left(1 - \frac{k}{\sum_{m=1}^k e^{\theta_m}}\right). \quad (28)$$

It is now straightforward to check that $S_k(\boldsymbol{\theta})$ given by formula (28) does indeed satisfy $H_k\left(\frac{\partial S_k}{\partial \boldsymbol{\theta}}, \boldsymbol{\theta}\right) = 0$, and that $S_k(\boldsymbol{\theta}^*) = 1 - (1/R_0) - \ln R_0$. Hence $A_k = (1/R_0) - 1 + \ln R_0$ for $k = 1, 2, \dots$. We thus arrive at the same conclusion as suggested by the Ornstein-Uhlenbeck approximation: that, to leading order, the mean persistence time does not depend upon k .

For the SIS model with Erlang distributed infectious periods, a precise result corresponding to formula (18) is available from the recent paper (Ball et al. 2016). Consider an SIS epidemic model in which individual infectious periods are distributed as any non-negative random variable Q of finite variance, and suppose that $R_0 = \beta E[Q] > 1$. Denote by p_Q the asymptotic (large N) probability, starting from a single infected individual in an otherwise susceptible population, that only a minor outbreak occurs. Denoting $g(\theta) = E[e^{\theta Q}]$ then it is well-known that p_Q is the unique solution in $[0, 1)$ of

$$p_Q = g(-\beta(1 - p_Q)).$$

Lemmas 3.2 and 3.3 of Ball et al. (2016) together imply that $\tau \sim \tau^{\text{BBN}}$, where

$$\tau^{\text{BBN}} = E[Q] \sqrt{\frac{2\pi}{N}} \frac{1}{(R_0 - 1)(1 - p_Q)} \exp(N((1/R_0) - 1 + \ln R_0)). \quad (29)$$

In the case of the classic SIS model, then $p_Q = 1/R_0$ and formula (18) is recovered. When Q follows an Erlang distribution with mean $E[Q] = 1/\gamma$ and variance $1/k\gamma^2$, then $g(\theta) = (1 - (\theta/k\gamma))^{-k}$, so that p_Q is the unique solution in $[0, 1)$ of

$$p_Q (1 + R_0(1 - p_Q)/k)^k = 1. \quad (30)$$

3.2 Numerical Results

For numerical work, we take $k = 2$, and investigate how well our general approximation methods perform for a 2-dimensional problem.

Consider first a system not far above threshold, with $R_0 = 1.1$. Comparing Fig. 7 with Fig. 4, the pictures seem rather similar, but there are significant differences in how they are constructed. In both cases, the exact values are computed by solving Eq. 3. For the classic SIS model this requires the evaluation of an eigenvalue of a square matrix Q^C of dimension N ; for the case $k = 2$, the matrix in question is of dimension $N(N + 3)/2$.

The diffusion approximation for the classic SIS model is computed by evaluating the integral (9). For $k = 2$, no such explicit expression for τ^{Diff} is available, and we proceed via the Finite Element Method. Freefem++ code (see Appendix) is used, which gives a numerical approximation to $\tau^{\text{Diff}}(\mathbf{y})$ for a grid of $\mathbf{y}$ values within Ω . The user specifies

Fig. 7 Expected persistence time from quasi-stationarity (τ^q) for the SIS model with Erlang distributed infectious periods having shape parameter $k = 2$, in a population of size N , together with approximations. Parameter values $\beta = 1.1$, $\gamma = 1$ (so that $R_0 = 1.1$ and $y_1^* = y_2^* \approx 0.0455$). Exact values computed from formula (3); diffusion approximation (21); Ornstein-Uhlenbeck approximation (10). Diffusion approximation is initiated at the finite element meshpoint nearest to $\mathbf{Y}(0) = N(y_1^*, y_2^*)$

the set of grid points along the boundary $S_a \cup S_r$, and then grid points in the interior of Ω are selected automatically by the software. To approximate τ_q^k , we select the grid point at minimal Euclidean distance from $N\mathbf{y}^*$. Since the surface $\tau^{\text{Diff}}(\mathbf{y})$ is quite flat around $N\mathbf{y}^*$, our approximation is not overly sensitive to the precise value of $\mathbf{y}$ chosen.

Turning to the asymptotic result of Ball et al. (2016), with $k = 2$ then p_Q satisfying equation (30) may be found explicitly as

$$p_Q = \frac{4 + R_0 - \sqrt{R_0(8 + R_0)}}{2R_0}.$$

Formula (31) in this case reduces to

$$\tau^{\text{BBN}} = \frac{1}{\gamma} \sqrt{\frac{2\pi}{N}} \frac{2R_0 \exp(N((1/R_0) - 1 + \ln R_0))}{(R_0 - 1)(R_0 - 4 + \sqrt{R_0(8 + R_0)})}. \quad (31)$$

The performance of our approximations for a system well above threshold ($R_0 = 1.5$) is shown in Fig. 8. The picture here is very similar to that seen in Fig. 5 for the classic SIS model: τ^{BBN} gives the correct leading-order asymptotic behaviour; τ^{Diff} underestimates mean persistence time; and τ^{OU} provides a very poor approximation. Correct leading-order asymptotic behaviour, in agreement with the result of Ball et al. (2016), may be found using the Hamiltonian approach.

Fig. 8 $\ln(\tau_q)/N$ for the SIS model with Erlang distributed infectious periods having shape parameter $k = 2$ in a population of size N , together with approximations. Parameter values $\beta = 1.5$, $\gamma = 1$ (so that $R_0 = 1.5$ and $\mathbf{y}^* \approx (0.1667, 0.1667)$). Exact values computed from formula (3); diffusion approximation (21); Ornstein-Uhlenbeck approximation (10). Diffusion approximation is initiated at the finite element meshpoint nearest to $\mathbf{Y}(0) = N(\mathbf{y}_1^*, \mathbf{y}_2^*)$

4 Discussion

For the below-threshold case $R_0 < 1$, we have seen that a linear birth-death process approximation works well provided the initial number of infected individuals is small. Although we only looked at the classic SIS model, a wide range of infection models (Ball 2014) are known to satisfy an epidemic threshold theorem. The threshold theorem tells us that for $R_0 < 1$ and population size N large, infection is very likely to die away quickly without ever infecting a significant fraction of the population; and that throughout a typical short-lived outbreak the number of susceptible individuals $S(t)$ may be well approximated by the constant N , so that consequently the process of infected individuals $I(t)$ may be well-approximated by a (linear) branching process. Thus we have a quite general approximation technique, which is seen to perform well in practice provided the number of initially infected individuals is small (Figs. 1 and 2). A number of previous authors have also considered the case in which a significant fraction of the population is initially infected. This situation is of less practical interest, since for sufficiently large N the chance of a significant fraction of the population ever becoming infected is negligible. The only approximation that we have found to work well in this case is the ‘diffusion approximation’ (Fig. 3), and the value of this approximation seems doubtful, since it appears no more straightforward to compute than the exact answer (indeed, if anything it seems more difficult to work with).

We turn now to the more interesting case $R_0 > 1$, so that it is possible for infection to sustain itself in the population over the long term. In this case, we studied time to extinction from an endemic state. First of all, consider the ‘diffusion approximation’. We saw that for moderate population sizes and R_0 not too far above 1, the approximation performs rather well (Figs. 4 and 7). On the other hand, it is known that this approximation gives the wrong asymptotic (large N) behaviour for the classic SIS model (Doering et al. 2005). We have seen (Figs. 5 and 8) that for the SIS model with Erlang-distributed infectious periods, exactly as for the classic SIS model, the diffusion approximation gives the wrong asymptotic behaviour. Hence we conclude that, firstly, the diffusion approximation does not seem trustworthy in cases where it is not possible to compute the exact answer for comparison; and secondly, the circumstances in which the diffusion approximation performs well seem to be precisely those circumstances in which it is straightforward to numerically evaluate the exact solution (1). Thus the value of this approximation method seems at best doubtful. The recent paper (Wang et al. 2014) uses a diffusion approximation to study time to extinction of *E. coli* O157:H7 infection from a herd of cattle. The diffusion is approximating a 2-dimensional Markov jump process, with population size of the order of 100 cattle and 10^{11} colony forming units (cfu) of *E. coli*. The state-space of the Markov jump process is therefore far too large to compute the mean extinction time exactly from Eq. 1, and some approximation technique is required. For different scenarios considered in Wang et al. (2014), mean extinction times computed from the diffusion approximation are from around 1 month up to around 40 years. This suggests that the infection process is not very far above threshold, and that the approximation may perform well. On the other hand, a population size of 10^{11} cfu is very large, which may lead the approximation to break down. Although it is certainly possible in these circumstances that the diffusion may give a good approximation to mean extinction time, our results suggest that one cannot have any real confidence of this. Similarly, in Doubova and Vadillo (2016), diffusion processes are used to compute mean extinction times for a Lotka-Volterra predator-prey model, for two variants of the SIS epidemic model, and for a two-pathogen epidemic model; there is no attempt to validate the approximations with reference to the underlying Markov chains.

Secondly, consider the Ornstein-Uhlenbeck approximation. This approach has been widely used in the literature due to its great simplicity. A formula such as Eq. 10 is much more analytically tractable than the exact solution (2) or the diffusion approximation (9). However, we have seen (Figs. 4, 5, 7 and 8) that this approximation performs extremely poorly. Nevertheless, we would argue that this approach retains some value in limited circumstances: specifically, for qualitative comparison between two infection models which share the same deterministic endemic prevalence level. This is because, for N large and R_0 well above 1, this approach gives a good approximation to the body of the quasi-stationary distribution — see, for example, Clancy (2005) and Nåsell (1999). The approximation fails in the tails of the distribution, hence the poor performance of our numerical approximation, which relies upon approximating the quasi-stationary probability that exactly one individual is infected. For qualitative comparison of two models, provided that both models have the same mean number of infected individuals in quasi-stationarity (approximated by the deterministic endemic level), it seems reasonable to suppose that the model whose quasi-stationary distribution has higher variance will assign greater probability to the state with one individual infected. The Ornstein-Uhlenbeck approach gives a good approximation for the variance of the quasi-stationary distribution, provided N is large and $R_0 \gg 1$. It is therefore reasonable to suppose that for two infection models sharing a common deterministic endemic prevalence level, the model under which the variance of the stationary distribution of the Ornstein-Uhlenbeck approximating process is greater will have the lower mean persistence time.

Finally, we turn to the Hamiltonian approach. Whereas diffusion approximations deal with moderate deviations from a deterministic mean, the Hamiltonian approach is suited to dealing with large deviations, and so would be expected to perform better when studying time to disease extinction. Our results (Figs. 5 and 8) confirm that this approach does indeed give the correct leading-order asymptotic behaviour for large N . This approach therefore seems the most promising overall, but a number of difficulties remain. Firstly, we have not obtained a useful numerical approximation for finite N , due to the unknown pre-factor C in the relation $\tau_q \sim C(N, R_0) \exp(NA)$. For 1-dimensional systems, including the classic SIS model, it is possible to evaluate the pre-factor C by retaining higher order information when approximating $K(\theta, t)$, see Doering et al. (2005) and Asaf and Meerson (2010). It does not seem straightforward to generalise this to higher dimensional systems, although some progress has been made for one particular 2-dimensional infection model in van Herwaarden and Grasman (1995), where it is shown that the pre-factor is of the form $C_0/\sqrt{N}$ for some (unknown) C_0 that does not depend upon N . Secondly, to evaluate the constant A in the dominant exponential term will, in general, require the solution of a boundary value problem for a $2k$ dimensional ordinary differential equation system, where k is the dimensionality of the original model. Recent work on numerical approaches to this problem includes (Schwartz et al. 2011; Lindley et al. 2014), where systems in $k = 2, 3$ dimensions are analysed. An alternative is to seek approximations valid within certain regions of parameter space, such as the ‘adiabatic approximation’ of Dykman et al. (2008), where (for their model of interest) a simple explicit expression for A valid for R_0 close to 1 is derived. Our results for the SIS model with Erlang-distributed infectious periods suggest another possibility: that for systems with sufficient symmetry, it may be possible to guess an explicit solution to the partial differential equation satisfied by $S(\theta)$, as we did to obtain formula (28), which leads directly to an explicit formula for A . In fact, it is not necessary to solve for $S(\theta)$ for all θ , but only to evaluate $S(\theta^*) - S(\mathbf{0})$. To illustrate this, consider an SEIS model, with a latent (‘exposed’) period between being infected and becoming infectious.

Suppose the latent period follows an Erlang distribution corresponding to j latent stages, each of mean $1/j\nu$, whilst the infectious period remains Erlang with k stages each of mean $1/k\gamma$. The Hamiltonian for this system is

$$\begin{aligned} H_{j,k}(\mathbf{y}, \boldsymbol{\theta}) = & \beta \left(\sum_{m=j+1}^{j+k} y_m \right) \left(1 - \sum_{m=1}^{j+k} y_m \right) (e^{\theta_1} - 1) \\ & + j\nu \sum_{m=1}^j y_m (e^{-\theta_m + \theta_{m+1}} - 1) \\ & + k\gamma \sum_{m=j+1}^{j+k-1} y_m (e^{-\theta_m + \theta_{m+1}} - 1) + k\gamma y_{j+k} (e^{-\theta_{j+k}} - 1) \end{aligned}$$

where $y_1, y_2, \dots, y_j$ correspond to latent stages and $y_{j+1}, y_{j+2}, \dots, y_{j+k}$ to infectious stages, and similarly for the components of $\boldsymbol{\theta}$. The corresponding equations of motion have a disease-free equilibrium point with $\boldsymbol{\theta}^* = (k, k, \dots, k, k-1, k-2, \dots, 3, 2, 1)\boldsymbol{\theta}_k^*$, where $\boldsymbol{\theta}_k^*$ is given by Eq. 25. On the hyperplane $\theta_1 = \theta_2 = \dots = \theta_j = \theta_{j+1}$, the function $S_k(\boldsymbol{\theta})$ given by formula (28) (with components $\theta_1, \theta_2, \dots, \theta_k$ re-labelled as $\theta_{j+1}, \theta_{j+2}, \dots, \theta_{j+k}$) is readily seen to satisfy the relevant Hamilton-Jacobi equation, and it follows that the action is given by $A_{j,k} = (1/R_0) - 1 + \ln R_0$ as before. Thus we see that the existence of a latent period has no effect upon the leading-order term in the expected time to extinction of infection. The above argument can be straightforwardly extended to allow for a latent period distributed according to any phase-type distribution.

Acknowledgments Elliott Tjia was supported by a studentship from the Engineering and Physical Sciences Research Council. The authors would like to thank Bernd Schroers, Robert Weston and Des Johnston for helpful discussions regarding the Hamiltonian approach.

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

Freefem++ code (Freefem++ 2016) to numerically solve the Kolmogorov backward equation, Eq. 22, in the case $k = 2$.

```
// set parameter values
real gamma=1., beta=2.0, n=260., epsilon=0.5;

// define the boundary
border B1a(t=0,epsilon){x=t; y=epsilon;}
border B1b(t=0,epsilon){x=epsilon; y=epsilon-t;}
border B2(t=epsilon,n){x=t; y=0;}
border B3(t=0,n){x=n-t; y=t;}
border B4(t=0,n-epsilon){x=0;y=n-t;}
```

```

// define mesh
mesh Th = buildmesh (B1a(10)+B1b(10)+B2(100)+B3(100)+B4(100));
// define finite element space
fespace Vh(Th,P2);
Vh tau,w;

// solve PDE
solve Backward(tau,w,solver=LU) =
int2d(Th) ( - (beta/(2*n)) * (x+y) * (n-x-y) + gamma*x * dx(tau) * dx(w)
-gamma*(x+y) * dy(tau) * dy(w)
+gamma*x*dy(tau) * dx(w)
+gamma*x*dx(tau) * dy(w) )
+int2d(Th) (2*gamma*(x-y)*w*dy(tau)
+ ((beta/n) * (x+y) * (n-x-y) - 2*gamma*x - (beta/(2*n)) * (n-2*x-2*y)
-gamma) * w * dx(tau) )
-int2d(Th) (-w)
+ on(B1a,tau=0) + on(B1b,tau=0) ;

// Output results, for processing in Matlab
{ ofstream ff("SISk.txt");
for (int i=0;i<Th.nt;i++)
{ for (int j=0; j <3; j++)
ff<<Th[i][j].x << " "<< Th[i][j].y<< " "<<tau[] [Vh(i,j)]<<endl;
ff<<Th[i][0].x << " "<< Th[i][0].y<< " "<<tau[] [Vh(i,0)]<< "
\n\n\n";
}
}

```

References

- Andersson H, Britton T (2000) Stochastic epidemics in dynamic populations: quasi-stationarity and extinction. *J Math Biol* 41:559–580
- Andersson H, Djehiche B (1998) A threshold limit theorem for the stochastic logistic epidemic. *J Appl Probab* 35:662–670
- Asaf M, Meerson B (2010) Extinction of metastable stochastic populations. *Phys Rev E* 81:021116
- Ball FG (2014) Epidemic thresholds. *Wiley StatsRef: Statistics Reference Online*. Published online 29. September 2014. Accessed 26 Sept 2016. <https://doi.org/10.1002/9781118445112.stat05273>
- 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, Neal P (2010) The time to extinction for a stochastic SIS-household-epidemic model. *J Math Biol* 61:763–779
- Clancy D (2005) A stochastic SIS infection model incorporating indirect transmission. *J Appl Probab* 42:726–737
- Clancy D, Pearce CJ (2013) The effect of population heterogeneities upon spread of infection. *J Math Biol* 67:963–987
- 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

- Dobova A, Vadillo F (2016) Extinction-time for stochastic population models. *J Comput Appl Math* 295:159–169
- Dykman MI, Mori E, Ross J, Hunt PM (1994) Large fluctuations and optimal paths in chemical kinetics. *J Chem Phys* 100:5735–5750
- Dykman MI, Schwartz IB, Landsman AS (2008) Disease extinction in the presence of random vaccination. *Phys Rev Lett* 101:078101
- Elgart V, Kamenev A (2004) Rare event statistics in reaction-diffusion systems. *Phys Rev E* 70:041106
- Ethier SN, Kurtz TG (1986) Markov processes: characterization and convergence. Wiley, New York
- Freefem++ (2016). <http://www.freefem.org/>. Accessed 26 Sept 2016
- Gardiner C (2009) Stochastic methods: A handbook for the natural and social sciences. 4th edition. Springer, Berlin
- Kamenev A, Meerson B (2008) Extinction of an infectious disease: a large fluctuation in a nonequilibrium system. *Phys Rev E* 77:061107
- Karlin S, Taylor HM (1975) A First Course in Stochastic Processes, 2nd edition. Academic Press, Cambridge
- Krysio RJ, Lefèvre C (1989) On the extinction of the S-I-S stochastic logistic epidemic. *J Appl Probab* 26:685–694
- Lindholm M, Britton T (2007) Endemic persistence or disease extinction: the effect of separation into sub-communities. *Theo Popul Biol* 72:253–263
- Lindley BS, Shaw LB, Schwartz IB (2014) Rare-event extinction on stochastic networks. *EPL* 108:58008
- Nåsell I (1999) On the time to extinction in recurrent epidemics. *J Roy Statist Soc B* 61:309–330
- Nåsell I (2002) Stochastic models of some endemic infections. *Math Biosci* 179:1–19
- Norden RH (1982) On the distribution of the time to extinction in the stochastic logistic population model. *Adv Appl Probab* 14:687–708
- Norris JR (1997) Markov Chains. Cambridge University Press, Cambridge
- Reddy JN (2006) An introduction to the finite element method. 3rd edition. McGraw-Hill, New York
- Schwartz IB, Forgoston E, Bianco S, Shaw LB (2011) Converging towards the optimal path to extinction. *J Roy Soc Interface* 8:1699–1707
- van Herwaarden OA, Grasman J (1995) Stochastic epidemics: major outbreaks and the duration of the endemic period. *J Math Biol* 33:581–601
- Wang X, Gautam R, Pinedo PJ, Allen LJS, Ivanek R (2014) A stochastic model for transmission, extinction and outbreak of *Escherichia coli* O157:H7 in cattle as affected by ambient temperature and cleaning practices. *J Math Biol* 69:501–532
- Weiss GH, Dishon M (1971) On the asymptotic behavior of the stochastic and deterministic models of an epidemic. *Math Biosci* 11:261–265