

Models for predicting extinction times: shall we dance (or walk or jump)?

^{1,2}Cairns, B.J., ¹J.V. Ross and ¹T. Taimre

¹*School of Physical Sciences, The University of Queensland, St Lucia QLD 4072, Australia.*

²*School of Biological Sciences, University of Bristol, Bristol BS8 1UG, U.K.*

E-Mail: jvr@maths.uq.edu.au

Keywords: Extinction; population model; estimation; prediction; catastrophes.

EXTENDED ABSTRACT

Measures of the risk of extinction, such as the expected time to extinction, are often used in population viability analysis (PVA) and in subsequent decision-making procedures, to gauge the impact of various management actions. Since PVA results are often used in the allocation of conservation funding, which in turn has a major impact on the future persistence of populations, the calculation of quantities such as the expected time to extinction is of great importance. These quantities are typically calculated by assuming a specific model for changes in the population over time. In particular, diffusion models are often used because they are simple to analyse and often give rise to explicit formulæ for most quantities of interest. However, while they are widely used, they frequently lead to inaccurate predictions of critical quantities such as the expected time to extinction. Hence, management decisions based on these predictions may be similarly flawed.

Often, a more appropriate model for describing the behaviour of the population in question is a discrete-state Markov process describing the actual number of individuals in the population. The most commonly used such models are birth-death processes or extensions thereof which allow for catastrophic events. Unfortunately, whilst these may be more appropriate for modelling the dynamics of the population in question, they are usually more difficult to work with, from both analytical and computational points of view. For these reasons, it is important to find some balance between accuracy of predictions based on the models and tractability of the method of prediction. Advances can be made by considering the limiting processes that correspond to these discrete-state models; in particular, Ornstein-Uhlenbeck processes and piecewise-deterministic processes with stochastic jumps. These models may still provide inaccurate predictions of extinction times, but should show improvement over the simplest Brownian motion approximation.

We consider populations that have density dependent demographic rates (in a specially-defined sense), and which may also be subject to environmental catastrophes. In particular, we assume that these populations may be modelled by continuous-time Markov chains — the stochastic SIS logistic model with or without binomial catastrophes occurring at a constant rate — and compare the accuracy of several approximations to the expected time to extinction. We contrast the various advantages of several methods for predicting extinction times for the above mentioned models, and compare these predictions for simulated data and a population of Bay checkerspot butterflies using model parameters estimated from data. We pay particular attention to the question of whether the extra analytical and computational effort required for the more complex models is necessary to inform decision-making in a conservation context.

We find that a variety of different models may give comparable results for measures of the risk of extinction (such as the expected time to extinction). This is true even in the situations we examine where catastrophes are known to play a role in population dynamics, but are not modelled when analysing the data. One model we consider in detail, the stochastic SIS logistic model and its Ornstein-Uhlenbeck diffusion approximation, is particularly robust in allowing for either strong or weak limiting of populations by their carrying capacities, and in adjusting for catastrophic events. A simpler geometric Brownian motion approximation may also provide reasonable results, but is less reliable due to shortcomings in its estimation of the population ceiling. Finally, we determine empirically that heuristic approximation methods for the stochastic SIS logistic model subject to catastrophes can provide accurate values for the expected time to extinction when the true parameters are known. The present lack of a suitable estimation procedure for these models would preclude their wider use, but fortunately other models, such as the Ornstein-Uhlenbeck approximation, can provide reasonable estimates of the expected time to extinction.

1 INTRODUCTION

One of the main challenges in population modelling, as with all mathematical modelling, is to find a balance between tractability and accuracy. While complex models involving many parameters might more accurately reflect reality, they may be difficult to analyse or difficult to fit from data. Conversely, simple models may fail to capture essential aspects of the population dynamics in question. In many circumstances, a useful approach is to develop a model that attempts to capture all of the relevant behaviour of the process, and then to find an approximation to this model that is more tractable. One such example is the use of diffusion approximations in population viability analysis (PVA). In performing such analyses, the scientist does not suppose that population size is actually a continuous variable, but rather that the population size is adequately approximated by a diffusion process. In many cases, such assumptions about the suitability of the approximating process can be justified mathematically; however, care must still be taken with quantities (such as extinction times) derived from such approximations (Pollett (2001)).

In this paper we will consider populations that have density dependent demographic rates (in a specially-defined sense), and which may also be subject to environmental catastrophes. We will assume that these populations may be modelled by continuous-time Markov chains—the stochastic SIS logistic model (Weiss and Dishon (1971)) with or without binomial catastrophes occurring at a constant rate—and compare the accuracy of several approximations to the expected time to extinction, one measure of the risk of extinction. Since the risk of extinction obtained in a PVA is often used in the allocation of conservation funding, which in turn can have a major impact on the future persistence of populations, these calculations are of great importance. We will pay particular attention to the question of whether the extra analytical and computational effort required for the more complex models is necessary to inform decision-making in a conservation context.

2 THE MODELS

We will consider a number of models for a biological population with per capita birth and death rates that are functions of the population density, rather than of the population size. That is, we assume that the functional relationships between the birth and death rates and the population size x have the form $nf(x/n)$ where n might be equal to a carrying capacity K or population ceiling N . We will call population processes that possess such rates *density dependent*.¹ Density dependent processes of this

¹In its common usage in ecology, ‘density dependence’ refers to a strong tendency of some populations to decline when above some carrying capacity; the carrying capacity is simply the population size above which this tendency begins. The definition we give here

form are amongst the most important in modelling biological populations of all types. Notable examples include most of the classical models for epidemics, the Levins model for metapopulations (Levins (1969)), and many early models for human populations (e.g. Verhulst (1838)). In this paper, our goal is to compare estimates of the time to extinction obtained from continuous-time Markov chain (CTMC) models with those obtained from approximations to these models involving either central limit-type results or heuristic approximation schemes.

2.1 Continuous-time Markov chains

For our purposes in this paper, we use only the transition rate matrix Q of the processes discussed in this subsection, which has off-diagonal elements q_{ij} giving the rate of jumps from state i to state j , and diagonal elements $q_{ii} = -q_i$ where q_i represents the rate of jumps from state i ($q_i = \sum_{j \neq i} q_{ij}$). For further details on the extensive theory of CTMCs see, for example, Norris (1997).

The stochastic SIS logistic model. We will begin by assuming that the underlying populations change over time according to a stochastic SIS logistic model (henceforth, simply the SIS model). Such models are named for their use in the study of epidemics in which susceptible individuals may become infective, then recover to become susceptible again (Susceptible-Infective-Susceptible); however, they are broadly applicable as models for density dependent populations. The SIS model has transition rates, for i and j in $\{0, 1, \dots, N\}$,

$$b(i) = q_{i,i+1} = \lambda i (1 - i/N), \quad (1)$$

$$d(i) = q_{i,i-1} = \mu i, \quad (2)$$

for births and deaths, respectively. The per capita birth and death rate parameters are respectively λ and μ . N is the maximum size of the population or population ceiling. For both parameter estimation and calculation of expected times to extinction, the SIS model may be approximated by an Ornstein-Uhlenbeck diffusion, as detailed in §2.2 and Appendix A.

A birth-and-catastrophe model. In some cases the main drivers of mortality in a population may be catastrophic events causing mass, rather than individual death. A number of authors have explicitly considered Markov chains in which catastrophic mortality plays an important role, including models for populations of Crabeater seals, *Lobodon carcinophagus* (Wilcox and Elder (2003)), and of California spotted owls, *Strix occidentalis occidentalis* (Andersen and Mahato (1995)). In this birth-and-catastrophe model, we replace individual death entirely by death due to catastrophic events, occurring at a constant rate ν , which affect each

is a more literal ‘dependence on the population density’, of any strength, in the given form.

and every individual in the population independently, killing them with a certain probability p . In this situation, we have $b(i)$ as in (1) above, but instead of individual death at rate $d(i)$, take the rate of jumps down from i to j ($< i$) as

$$c(i, j) = q_{ij} = \nu \binom{i}{j} (1-p)^j p^{i-j}. \quad (3)$$

That is, catastrophes occur at a constant rate ν and have size given by a binomial distribution. We will approximate the expected extinction times for this model using a piecewise-deterministic Markov process (§2.3).

2.2 Diffusion and deterministic approximations

A number of approximation methods exist for the SIS model (and other density dependent CTMCs). One such method is the deterministic approximation for density dependent processes, which is obtained by scaling the process from $\{0, 1, \dots, N\}$ to $[0, 1]$ by dividing through by N , and letting $N \rightarrow \infty$ to obtain a functional law of large numbers (Kurtz (1970)). For example, the SIS model has a deterministic approximation given by the solution to the ordinary differential equation $dx/dt = \lambda x(1-x) - \mu x$, for $x(0) \in [0, 1]$.

Deterministic approximations may not be useful in obtaining measures of extinction risk, since they may predict that the population reaches a steady state at 0 or some non-zero fixed point. Pollett (2001) compares two methods for approximating the expected time to extinction when the population ceiling N is large: an asymptotic formula, and the numerical solution to the hitting times of a diffusion approximation. By removing the trend in the process given by the deterministic approximation described above, and scaling by \sqrt{N} rather than N , an Ornstein-Uhlenbeck (OU) approximation around a stable fixed point in the deterministic dynamics may be found (see Pollett (2001)). This OU process is stationary, Gaussian, Markovian, and has parameters that are again given explicitly in terms of the original parameters λ and μ .

2.3 Piecewise-deterministic Markov processes

The birth-and-catastrophe process is not density dependent in the sense required for a diffusion or deterministic approximation of the type discussed in §2.2. However, it is density dependent ‘between catastrophes’, and the size of a catastrophe—conditional on its occurrence and the population size just before the catastrophe—converges to a constant proportion of this population size as $N \rightarrow \infty$. This suggests an approximation for the birth-and-catastrophe process that is somewhat analogous to the deterministic, law of large numbers, approximation for density dependent processes.

In this approximation, between catastrophes the population dynamics are described by an ordinary

differential equation, given by the solution to $dx/dt = \lambda x(1-x)$. Then, catastrophes occur according to a Poisson process with rate parameter ν , and when they occur they instantly reduce the size of the population to a proportion $(1-p)$ of the population size just before a catastrophe. That is, if S is the time of a catastrophe, then $x(S) = (1-p)(\lim_{t \uparrow S} x(t))$. Since this process is Markovian and exhibits deterministic behaviour between stochastic jumps, it is known as a *piecewise-deterministic Markov process* (PDMP). The class of PDMPs includes almost every Markovian process that is not a diffusion (Davis (1984)), including CTMCs, but is relatively little-studied in general. We propose a simple but robust procedure for numerical calculation of the expected time to extinction of this PDMP (Appendix B).

2.4 Geometric Brownian motion

Foley (1994) considers a simple procedure for estimating extinction risk to populations using geometric Brownian motion (GBM). Under this model, the population changes according to a stochastic differential equation, such that, if the natural log of the population size is $Y(t) = \log X(t)$ at time t , then $Y(t)$ is assumed to satisfy the stochastic differential equation $dY(t) = r_d dt + \sqrt{v_r} dW$, where W is the white noise process. The parameters are the mean and variance of the natural log of the growth rate of the population, r_d and v_r respectively. The advantage of this model is that expected times to extinction can be found analytically in all cases; however, it explicitly leaves out density dependence in the population growth, and is not proposed as an approximation to any particular discrete-state population model (however, it may be constructed as the limit of a random walk). Foley uses simulation methods to explore the possibility of density dependence in a range of ecological time series, but reports very little difference in the results.

3 RESULTS

We first consider the time series of the Jasper Ridge ‘JRH’ population of Bay checkerspot butterfly, *Euphydryas editha bayensis* sampled yearly between 1960 and 1986 (Harrison *et al.* 1991).² The data from this population was analysed by Foley (1994) in the context of the GBM model described in §2.4. Harrison *et al.* (1991) found that the observed data are consistent with the assumption of density independence or only weak density dependence, where they use this term in the ecological sense noted in §2. We obtained estimates of birth and death rates (λ and μ respectively) and of the population ceiling N from the OU approximation to the SIS model, using a numerical maximum likelihood procedure. We used

²As we discuss in the next section, this population is now extinct. Although data exists up to the year of extinction, it is of interest here to compare approaches for prediction of extinction risk using data and analyses from a time when the population persisted.

Foley's estimates of parameters r_d and v_r for GBM or found our own (see Appendix A). The best estimates obtained using these procedures are given in Table 1.

	Estimate	JRH	Sim.
(a)	N	1998*	299
	r_d	-0.052*	-0.0011
	v_r	0.840*	0.0571
(b)	N	1998*	400
	λ	4.8585	1.6554
	μ	3.9506	1.0989
(c)	N	297417	3200
	λ	240.4360	3.4413
	μ	240.0446	3.2811

Table 1. Estimates obtained from the JRH butterfly data and the simulated birth-and-catastrophe data, using Appendix A: (a) for GBM; and for the SIS model: (b) assuming N is as in (a), and (c) assuming N is unknown. Starred values (*) are due to Foley (1994).

Figure 1 shows expected times to extinction calculated from the GBM model, the SIS model and its OU approximation, using the estimates from Table 1 JRH(a) & (c), and the methods detailed in Appendix B. The population was considered extinct when it dropped below 1 individual.




Figure 1. Estimated mean time to extinction for the JRH population of Bay checkerspot butterflies. Foley's (1994) approximation (dash-dot), an SIS model (solid) and its OU approximation (dashed) assuming N is unknown.

For the SIS model, when $N = 1998$ was considered known we used the estimates from Table 1 JRH(b), and found that the mean time to extinction was approximately 4.14×10^{13} years across the entire range of starting values; considerably longer than most estimates for the current age of the universe. We omit it from Figure 1 in order to better see the relationship between the other results. The values for the OU approximation for $N = 1998$ are assumed to

be similarly large, but consequently we were unable to compute them accurately.

We used simulated data to examine the role of catastrophes in complicating the prediction of extinction risk. We constructed a birth-and-catastrophe CTMC model as detailed in §2.1, with parameters $N = 400$, $\lambda = 0.25$, $\nu = 2/3$ and $p = 0.25$. This represents a growing population with relatively frequent catastrophes which, when they occur, kill, on average, a quarter of the population; note, however, that the mean per capita death rate νp is only two thirds of λ , so when the population density is low, the population tends to increase. A hundred years of simulated data is given in Figure 2.




Figure 2. One hundred years of population sizes obtained by simulating a birth-and-catastrophe process.

We again computed estimates for λ , μ and N using the OU approximation, and for r_d and v_r for the GBM model, assuming either that N was 400, as used in the actual simulation, or that N was unknown (see Table 1 Sim.(a)–(c)). From these estimates, expected times to quasi-extinction below 40 individuals were computed for the various models (Figure 3).³ Again, the SIS model and OU approximation from the Sim.(b) estimates produced large values, 4.66×10^6 and 1.60×10^7 respectively, and are omitted from Figure 3. We compared the expected times to quasi-extinction from each of the fitted models to the values computed from the underlying birth-and-catastrophe process and its PDMP approximation using the true parameter values; the width of the plotted 'ribbon' gives the precision of the numerical procedure for the PDMP. (We also computed an additional piecewise-deterministic approximation by taking the estimates of λ , μ and N obtained from the Sim.(c) estimates, while assuming catastrophes with killing probability $p = 0.25$. The mean times to quasi-extinction found in this way were an order of magnitude smaller than the plotted values. This indicates that, although this procedure might be a natural first guess for scientists

³This level of quasi-extinction is the size below which our hypothetical population is considered functionally extinct or at least severely threatened.




Figure 3. Expected times to quasi-extinction below 40 individuals for: the SIS (black, solid) and OU approximation (black, dashed) using estimates with N unknown; the underlying birth-and-catastrophe process used in the simulation (dark grey solid line), the piecewise-deterministic approximation to this (dark grey ribbon), assuming $N = 400$, and the GBM approximation (light grey dashed).

attempting to extract information about catastrophes from estimates for individual mortality, it is unlikely to produce accurate results). The SIS and OU results were reasonably accurate despite overestimating the expected time to extinction.

4 DISCUSSION

Our analysis of the JRH data provides some interesting insights into the relationship between the various models presented here. First, the estimates for λ and μ differ markedly depending on whether or not N is assumed to be known (Table 1). The N -unknown estimates support the conclusion of Harrison *et al.* (1991): since the estimates for λ and μ are so close, the population tends to be small, relative to N , and hence $b(i) \approx \lambda i$ over a wide range of population sizes. Consequently, the population does not exhibit a strong downward trend in size following deviations from its long-term mean size. This result indicates a degree of flexibility in the SIS model and the estimation procedure that we propose when N is assumed to be unknown, because, while the model allows for a tendency to decline at large population sizes, it does not require this tendency to be strong. As a wide variety of density dependent processes may be analysed using this approach, it appears to provide a robust method for assessing whether populations are strongly limited by their carrying capacity.

The SIS, OU and GBM models for the butterfly population showed somewhat different predictions for the expected time to extinction, although the SIS model and its OU approximation gave very similar values (Figure 1). The GBM model produced predictions that were similar to those obtained from the SIS and OU models assuming an unknown N ,

unlike those found when assuming $N = 1998$. This suggests that both the SIS/OU and GBM approximations can give comparable estimates despite the difference in their structure and estimates for N . However, because the estimate of N for the GBM model is biased by the quantity of data, the degree of compatibility between these results may vary. A comparison of the results of these models may serve as a useful test of the reliability of predictions of the time to extinction, especially when considering density dependence.

The Bay checkerspot butterfly is now extinct on Jasper Ridge, with the JRH population thought to have gone extinct in 1998 due to a combination of extensive habitat loss and a period of increased climatic variability beginning in 1972 (McLaughlin *et al.* (2002)). It is interesting to note that, despite the simplicity of the OU approximation, its expected time to extinction is quite close to the observed time (over the recorded population range), 12 years after the last measurements we used.

The birth-and-catastrophe and PDMP models tell a rather different story. Because there is at present no suitable method for estimating parameters for a birth-and-catastrophe model from real-world data, we constructed data (Figure 2) simulated from a CTMC birth-and-catastrophe model in order to assess the performance of the OU and GBM approximations in dealing with the presence of catastrophic mortality. Again, both these and the SIS model derived from the OU estimates of λ , μ and N , gave results that are roughly similar to the exact results for the underlying birth-and-catastrophe process (Figure 3). This suggests that these models are able to adjust for the effects of catastrophes. Finally, the PDMP model appears to give very good approximations to the mean times to (quasi-) extinction for birth-and-catastrophe processes (Figure 3); these were consistently overestimates, but followed the pattern of the true values quite closely. However, the mathematical links between the birth-and-catastrophe process and the PDMP approximation remain to be properly described, and in the absence of a procedure for estimating birth and catastrophe rates and killing probabilities, the utility of this model in dealing with real-world data will remain limited. Fortunately, with careful use, models such as the OU approximation are adequate in at least some circumstances.

Shall we dance with a diffusion, walk with a Markov chain, or jump with a piecewise-deterministic process? It is important to capture as much of the dynamics as possible when constructing a model of a population, and so this question is an important one. Although this is only a preliminary study, the evidence presented herein suggests that the SIS model, with parameter estimates provided by its OU approximation, can provide reasonable approximations to measures of extinction risk, such as

the expected time to extinction. The GBM model may also be capable of providing reasonable predictions, although it may be prone to overestimate the expected time to extinction. This may even be true when catastrophic declines in population are a feature of the process, although care should be taken to examine the reliability of these estimates, for example by simulation methods such as we have used here.

5 ACKNOWLEDGEMENT

This work was supported by PhD scholarships to each of the authors from the Australian Research Council Centre of Excellence for Mathematics and Statistics of Complex Systems. The authors would like to thank Phil Pollett for providing code to compute the mean time to extinction for the OU approximation, and the referees for their helpful suggestions.

6 REFERENCES

- Andersen, M.C. and D. Mahato (1995), Demographic models and reserve designs for the California spotted owl. *Ecol. Appl.* 5, 639–647.
- Davis, M.H.A. (1984), Piecewise Deterministic Markov Processes: a general class of non-diffusion stochastic models. *J.R.S.S. (B)* 4, 353–388.
- Foley, P. (1994), Predicting extinction times from environmental stochasticity and carrying capacity. *Cons. Biol.* 8, 124–137.
- Gihman, I.I. and A.V. Skorohod (1972), *Stochastic Differential Equations*. Springer-Verlag, Berlin.
- Hansen, E. and G.W. Walster (2004), *Global Optimization Using Interval Analysis*, 2nd edition. Marcel Dekker, New York.
- Hanson, F.B. and H.C. Tuckwell (1981), Logistic growth with random density independent disasters. *J. Math. Biol.* 19, 1–18.
- Hanson, F.B. and H.C. Tuckwell (1997), Population growth with randomly distributed jumps. *Theor. Pop. Biol.* 36, 169–187.
- Harrison, S., J.F. Quinn, J.F. Baughman, D.D. Murphy and P.R. Ehrlich (1991), Estimating the effects of scientific study on two butterfly populations. *Am. Nat.* 137, 227–243.
- Kurtz, T.G. (1970), Solutions of ordinary differential equations as limits of pure jump processes. *J. Appl. Prob.* 7, 49–58.
- Levins, R. (1969), Some demographic and genetic consequences of environmental heterogeneity for biological control. *Bull. Entom. Soc. Amer.* 15, 237–240.
- McLaughlin, J.F., J.J. Hellmann, C.L. Boggs and P.R. Ehrlich (2002), Climate change hastens population extinctions. *Proc. Nat. Acad. Sci.* 99, 6070–6074.

Pollett, P.K. (2001), Diffusion approximations for ecological models. *Proceedings of the International Congress on Modelling and Simulation*, Vol. 2, Modelling and Simulation Society of Australia and New Zealand, Canberra, Australia, 843–848.

Rubinstein, R.Y. and D.P. Kroese (2004), *The Cross-Entropy Method: A unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer-Verlag.

Verhulst, P.F. (1838), Notice sur la loi que la population suit dans son accroissement. *Corresp. Math. et Phys.* X, 113–121.

Weiss, G.H. and M. Dishon (1971), On the asymptotic behaviour of the stochastic and deterministic models of an epidemic. *Math. Biosci.* 11, 261–265.

Wilcox, C. and B. Elder (2003), The effect of density-dependence catastrophes on population persistence time. *J. App. Ecol.* 40, 859–871.

APPENDIX A. PARAMETER ESTIMATION

The SIS model

Our parameter estimation method for the SIS model uses the OU approximation as follows. Firstly, the OU process is strongly stationary if we start it in equilibrium: $Z(0) \sim N(0, \sigma^2)$, where $\sigma^2 = \mu/\lambda = \rho (< 1)$ (see Pollett (2001)). Therefore, we have that

$$X_N(0) \sim N\left(x^*, \frac{\sigma^2}{N}\right),$$

where $x^* = 1 - \rho$ is the stable fixed point of the density dependent deterministic approximation, and N is the population ceiling (see Pollett (2001)). We also have that the covariance of the state of the OU process at successive times is given by

$$\begin{aligned} c(t) &:= \text{Cov}(X_N(s), X_N(s+t)) \\ &= c(0) \exp(B|t|), \end{aligned} \quad (4)$$

where $c(0) = \text{Var}(X_N(s)) = \sigma^2/N$ and $B = -(\lambda - \mu)$ (see Pollett (2001)). Hence, we know explicitly the correlation structure of the Gaussian vector $(X_N(t_1), X_N(t_2), \dots, X_N(t_n))$, where n is the number of observations. Thus, we know its likelihood function:

$$f(x) = \frac{1}{\sqrt{(2\pi)^n |V|}} \exp\left[-\frac{1}{2}(x - \bar{x})V^{-1}(x - \bar{x})'\right], \quad (5)$$

where $\bar{x} = (\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n)$, $\bar{x}_i = x^*$ for all $i = 1, 2, \dots, n$, and V is the covariance matrix with elements $v_{i,j}$, where $v_{i,i} = \sigma^2/N$ and $v_{i,i+s} = (\sigma^2/N) \exp(B|t_i - t_{i+s}|)$.

We may therefore evaluate the (joint) maximum likelihood estimators, which are the values of λ and μ that maximise (5). Explicit calculation of these is

not feasible if the sample size is large. Hence, we use the Cross-Entropy Method of numerical optimisation (see Rubinstein and Kroese (2004)) to find the parameters which maximise the likelihood function, but other numerical optimisation procedures should be similarly effective. When the population ceiling N is unknown, we estimate it by using the approximation $NX_N(0) \sim N(x^*N, \sigma^2N)$ and then calculate the joint maximum likelihood estimators with the addition of N as an unknown parameter.

Note that because the OU approximation is achieved by letting the maximum population size tend to infinity, this procedure is best for *large* population sizes. However, unequally spaced sampling of the process does not present difficulties, as can be seen from the covariance structure (4).

Geometric Brownian motion

Foley (1994) proposes a range of tools for estimating the parameters r_d , v_r and N using the GBM approximation. For comparison, we employ the simplest such approach. If $r = Y(t) - Y(t-1)$, where $Y(t)$ is the natural logarithm of the population size at time t , then r_d is just the arithmetic mean of r , and v_r is just the variance of r , as computed using the standard unbiased estimators. The population ceiling N is chosen in this instance as one more than the largest observed population size.

APPENDIX B. MEAN TIMES TO EXTINCTION

Continuous-time Markov chains

When dealing with CTMCs, the expected time to extinction is relatively easy to calculate in many cases, using the transition rate matrix Q (e.g. constructed from (1)–(3)). First, we restrict Q to the non-extinct population sizes; we remove all those rows and columns with entries corresponding to (quasi-) extinct population sizes, and call this restricted matrix M . Then, the expected time to extinction is the minimal, non-negative solution τ to $M\tau = -1$. This system of linear equations can easily be solved using virtually any numerical computing package.

Ornstein-Uhlenbeck processes

We follow Pollett (2001) and solve the partial differential equations corresponding to the OU approximation to the SIS model, using the `bvp4c` procedure in MATLAB on the appropriate boundary value problem. We choose slightly different bounds between which to evaluate the solution, to allow for quasi-extinction and to make the upper bound equivalent to N . See Pollett (2001) for further details.

Geometric Brownian motion

Foley (1994) gives the derivation of the expected time to extinction for a population under the assumption

that the population size changes approximately according to a GBM. The expected time to extinction $T_e(x_0)$, when the log of the initial population size is x_0 , is given by

$$T_e(x_0) = 2 \frac{x_0}{v_r} \left(n - \frac{x_0}{2} \right), \quad r_d = 0,$$

where r_d is the mean of the natural log of the population growth rate, v_r is the variance in the log growth rate, and $n = \log N$ is the log of the population ceiling, or,

$$T_e(x_0) = \frac{1}{2sr_d} (e^{2sn}(1 - e^{-2sx_0}) - 2sx_0),$$

if $r_d \neq 0$, where $s = r_d/v_r$. Because of the symmetry of the process in the log of the population size, we may adjust for quasi-extinction by subtracting the log of the quasi-extinction level from both x_0 and n (in the above, we have assumed the extinction level is 1).

Piecewise-deterministic Markov processes

Expected times to extinction for PDMPs have been considered analytically by Hanson and Tuckwell (1981, 1997). We employ a rather different, unsophisticated, yet very robust approach to obtain numerical solutions. The expected times to extinction $\tau(u)$, for $u \in [0, 1]$, satisfy a delay differential equation that may be solved implicitly, by use of an integrating factor and careful selection of boundary values, to obtain a (delay) integral equation: $\tau(u)$ then also solves

$$\tau(u(t; u_e)) = \frac{1}{\nu} + e^{\nu t} \int_t^\infty \nu e^{-\nu x} \tau((1-p)u(t; u_e)) dx. \quad (6)$$

Here, as earlier, λ is a growth rate constant, ν is the rate of occurrence of catastrophes, p is the killing probability. The constant $u_e \in (0, 1)$ is the quasi-extinction level, and $u(t; u_e)$ is the solution to $du/dt = \lambda u(1-u)$, with $u(0) = u_e$. In order to proceed, we note that (6) has the form $\tau = K\tau$, where K is the functional given by the right-hand side of the equation, and hence the expected time to extinction is the minimal, non-negative fixed point of K .

We exploit the form of (6) by approximating the integral by its lower and upper Riemann sums, respectively \underline{K} and \overline{K} . Then, it is easy to show that, because $\underline{K}h \geq \underline{K}g$ and $\overline{K}h \geq \overline{K}g$ if $h(u) \geq g(u)$ for all u (the functionals are monotone increasing), the minimal non-negative fixed points of \underline{K} and \overline{K} bound τ from below and above, respectively (see Lemma 2.11.1 of Gihman and Skorohod, 1972). These fixed points can be found using numerical linear algebra. Then, by refining the mesh on which the Riemann sums are calculated, the bounds on τ can be tightened, and the use of interval analysis (see, e.g., Hansen and Walster (2004)) can ensure that numerical errors in the calculation of these bounds are also accounted for.