Preface

Mathematical modeling in population biology aims to give insight into the dynamics of both single populations and multiple interacting populations. Fundamental questions are of qualitative nature. They concern the long-term survival or extinction of any given population or subpopulation. Work in this area has for a long time followed two different mathematical branches: deterministic and stochastic. The deterministic models have given powerful results by responding to fundamental qualitative questions concerning survival and extinction. At the same time, it is well known that the deterministic models suffer from a serious weakness: the state-space used is continuous, while counts of actual populations are always given by nonnegative integers. This weakness is avoided by using stochastic models with discrete state space. The question therefore arises if the qualitative results that hold in the deterministic setting have counterparts in the more realistic stochastic version of the model. Another way of formulating this question is to ask whether the threshold results that hold in the deterministic setting have counterparts in the stochastic world.

There is no easy or general answer to this question. We approach it by presenting an analysis of the so-called SIS model that is used to study the spread of infection without immunity in a constant population. This is one of the simplest models where this investigation is meaningful. Our methods may be useful in more complicated and more realistic models. The deterministic version of the SIS model has a threshold where a parameter denoted by R_0 takes the value 1. This means that an endemic infection level will establish itself if $R_0 > 1$ and the initial proportion of infected individuals is positive, while any infection will ultimately disappear if $R_0 < 1$. The stochastic version of the SIS model takes the form of a birth–death process with an absorbing state at the origin. Here, any infection will ultimately disappear for all values of R_0. We investigate the long-term behavior of this model by studying both the time to extinction and the so-called quasi-stationary distribution, which is a stationary distribution, conditional on not being absorbed at the origin. Its analysis is more intricate than the analysis of an ordinary

stationary distribution of a birth–death process. In particular, explicit solutions are not available. Therefore, the emphasis is on deriving approximations.

Several approximation steps are required to derive satisfactory approximations of the quasi-stationary distribution $\mathbf{q}$ of the SIS model. These approximation steps have been developed gradually over time from 1978 till 2011. The first step was taken by Cavender (1978) [16]. He introduced an auxiliary process as a related birth–death process without an absorbing state. This auxiliary process has a stationary distribution $\mathbf{p}^{(0)}$ that can be determined explicitly. Cavender showed that it serves as a lower bound of $\mathbf{q}$ in the sense of stochastic ordering: $\mathbf{p}^{(0)} \prec_{ST} \mathbf{q}$.

The next step was taken by Kryscio and Lefèvre (1989) [39]. They introduced an additional auxiliary process as a second related birth–death process without an absorbing state. Again, its stationary distribution $\mathbf{p}^{(1)}$ can be determined explicitly. Also, Kryscio and Lefèvre made the important conjecture that this stationary distribution provides an upper bound of the quasi-stationary distribution $\mathbf{q}$ in the sense of stochastic ordering: $\mathbf{q} \prec_{ST} \mathbf{p}^{(1)}$.

After this, Nåsell (1996) [44] proceeded to derive approximations of the stationary distributions $\mathbf{p}^{(0)}$ and $\mathbf{p}^{(1)}$. He searched for asymptotic approximations as the total population size N became large. He showed that qualitatively different results could be established in three different parameter regions. Two of these were counterparts to the two regions $R_0 < 1$ and $R_0 > 1$ that hold in the deterministic case, while a third region appeared as a transition region near $R_0 = 1$. Nåsell also identified a map Ψ that could be used to determine $\mathbf{q}$ numerically by iteration. He used this map to derive explicit approximations of $\mathbf{q}$ by applying one iteration step to the approximation $\mathbf{p}^{(1)}$. The resulting approximation is ad hoc, since it was not then clear in what sense an approximation of $\Psi\left(\mathbf{p}^{(1)}\right)$ is an approximation of the quasi-stationary distribution $\mathbf{q}$ itself.

The map Ψ was also used by Ferrari, Kesten, Martínez, and Picco (1995) [27], although this publication was not available when the 1996 paper by Nåsell was written. Ferrari et al. showed that the map Ψ was important for the quasi-stationary distribution $\mathbf{q}$ in the sense that $\lim_{i\to\infty}\Psi^i(\mathbf{p}) = \mathbf{q}$ holds for arbitrary distributions $\mathbf{p}$, and also that $\Psi(\mathbf{q}) = \mathbf{q}$.

The conjecture by Kryscio and Lefèvre that $\mathbf{p}^{(1)}$ is an upper bound of $\mathbf{q}$ in the sense of stochastic ordering remained an open problem until it was settled by Clancy and Pollett (2003) [18]. In their proof of this conjecture, they made use of the map Ψ. They also established an important theorem that showed that the map Ψ preserves a certain ordering between probability vectors.

In this monograph, we derive approximations of some of the distributions $\Psi\left(\mathbf{p}^{(1)}\right)$, $\Psi\left(\mathbf{p}^{(0)}\right)$, and $\Psi^i\left(\mathbf{p}^{(1)}\right)$ in the three parameter regions. Using the theorem established by Clancy and Pollett, we are able to show that these approximations actually give approximations of the quasi-stationary distribution $\mathbf{q}$ itself.

The final step in the derivation of approximations of the quasi-stationary distribution is also taken in this monograph, where we derive approximations that are uniformly valid across all three of the parameter regions.

We also give approximations of the time to extinction for two specific initial distributions. The first case is used to study extinction time for an established infection, and the second one is used to study establishment of an infection. For the first case, we let the initial distribution be equal to the quasi-stationary distribution. The time to extinction then has an exponential distribution whose expectation is determined by the quasi-stationary distribution. Its approximation is therefore found from the approximation of the quasi-stationary distribution described above. For the second case, we consider the case when initially one infective individual is present. The expected time to extinction can then be determined from the distribution $\mathbf{p}^{(0)}$, for which an approximation has been derived, as described above.

The monograph is written for a reader who has a good working knowledge about birth–death processes. An introduction is given to the less well-known concept of quasi-stationarity, and to its relation with extinction times. Heavy use is made in the monograph of ideas from the area of asymptotic approximation. An effort is made to present these ideas in a simple way, since it is an area of applied mathematics that is less well known to the stochastic community.

I am grateful to Joshua Ross and to three anonymous referees for careful reading of the manuscript and for making valuable comments that have improved the presentation of the results considerably.

Stockholm   *Ingemar Nåsell*
March 2011