

NUMERICAL METHODS IN THE CONTEXT OF COMPARTMENTAL MODELS IN EPIDEMIOLOGY

PETER KRATZ¹, ETIENNE PARDOUX¹ AND BRICE SAMEGNI KEPGNOU¹

Abstract. We consider compartmental models in epidemiology. For the study of the divergence of the stochastic model from its corresponding deterministic limit (i.e., the solution of an ODE) for long time horizon, a large deviations principle suggests a thorough numerical analysis of the two models. The aim of this paper is to present three such motivated numerical works. We first compute the solution of the ODE model by means of a non-standard finite difference scheme. Next we solve a constraint optimization problem via discrete-time dynamic programming: this enables us to compute the leading term in the large deviations principle of the time of extinction of a given disease. Finally, we apply the τ -leaping algorithm to the stochastic model in order to simulate its solution efficiently. We illustrate these numerical methods by applying them to two examples.

Résumé. On considère des modèles comportementaux en épidémiologie. Afin d'étudier l'écart en temps long entre le modèle stochastique et sa limite loi des grands nombres (qui est la solution d'une EDO), on se base sur un principe des grandes déviations, qui nous conduit à mener une étude numérique des deux modèles, sur trois aspects différents. Tout d'abord, nous calculons une solution approchée de l'EDO à l'aide d'une méthode numérique dite "non-standard". Ensuite nous résolvons numériquement un problème de contrôle sous contrainte, afin de calculer le terme principal des grandes déviations du temps de sortie d'une situation endémique. Enfin nous mettons en oeuvre l'algorithme du " τ -leaping" pour simuler efficacement la solution du système stochastique. Nous illustrons ces simulations numériques en les appliquant à deux exemples.

INTRODUCTION

It is well-known that deterministic ODE models of population dynamics (such as the evolution of diseases) are not appropriate for small population sizes N (i.e., $N < 10^3, 10^6$). Recently, it has been shown (see [1]) that in some cases the ODE models can diverge from the corresponding stochastic models even for larger population sizes ($N > 10^6$).

We consider the "natural" stochastic extensions for many compartmental ODE models in epidemiology: individual based Poisson driven models. The ODE models can be recovered from these models as the law of large numbers limit as the population size N approaches infinity. Hence, the deterministic model can be considered an appropriate first approximation of the stochastic model for large N . While the dynamics of many of the underlying deterministic models are well-understood, the analysis of the corresponding stochastic models is often difficult.

We are interested in the long-term behavior of the epidemic process. For many deterministic models, the disease will finally either die out or become endemic, depending on the parameter choice and/or the initial

¹ Aix-Marseille Université, CNRS, Centrale Marseille, I2M, UMR 7373 13453 Marseille, France;
e-mail: kratz@mathematik.hu-berlin.de; etienne.pardoux@univ-amu.fr; skbrice@yahoo.fr

sizes of the compartments, i.e., the solution of the ODE converges to an asymptotically stable equilibrium. The theory of Large Deviations (also called the large deviations principle (LDP)) allows us to quantify the time taken by the stochastic model to diverge significantly from its deterministic limit. In particular, one can compute asymptotically, for large N , the expected time which is needed to transfer the epidemic process from the domain of attraction of one equilibrium to the one of another equilibrium. For the simplest models, this helps, e.g., to answer the question of when an epidemic becomes extinct although it should be endemic according to the deterministic approximation. Particularly interesting cases are epidemiological models where the ODE has several stable equilibria (usually a disease-free equilibrium and an endemic equilibrium as, e.g., in [2, 3]).

The remainder of this article is structured as follows. In section 1, we formulate stochastic and deterministic compartmental models and introduce two specific models which we analyze throughout this article by means of numerical methods (section 1.1). We furthermore describe how the time of exit from the domain of attraction of a stable equilibrium can be computed (section 1.2). This motivates three different numerical projects which we outline in section 1.3. We address these questions in sections 2 - 4 as follows. First we apply a non-standard finite difference scheme due to [4] in order to compute the solution of the deterministic model numerically. Second, in section 3, by solving a control problem via dynamic programming, we compute the leading coefficient in the large deviations analysis of the time of exit τ^N by the solution of the SDE from the domain of attraction of a locally stable equilibrium of the ODE. Finally, we simulate the stochastic process using the so-called τ -leaping algorithm.

1. COMPARTMENTAL MODELS IN EPIDEMIOLOGY

1.1. Stochastic and deterministic models

We consider diseases in large populations of (initially) N individuals which are split up into different groups (or compartments) according to the disease status of the individuals. Such groups are, e.g., the group of individuals susceptible to the disease, the group of infectious individuals and the group of immune individuals. In order to better compare the population dynamics for different values of N , we normalize the compartment sizes by dividing by N ; instead of considering the number of individuals in the different groups, we hence rather consider their “proportions” with respect to the total size of the population (note that for models with non-constant population size, these are only initially truly proportions).

For a closed set $A \subset \mathbb{R}^d$, we define the d -dimensional jump process Z^N via a set of k jump directions $h_j \in \mathbb{Z}^d$ and respective Poisson rates $\beta_j(z)$, $j = 1, \dots, k$; $Z_i^N(t)$ denotes the “proportion” of individuals in compartment $i = 1, \dots, d$ at time $t \geq 0$:

$$\begin{aligned} Z^N(t) &:= Z^{N,x}(t) := x + \frac{1}{N} \sum_{j=1}^k h_j P_j \left(\int_0^t N \beta_j(Z^N(s)) ds \right) \\ &= x + \int_0^t \sum_{j=1}^k h_j \beta_j(Z^N(s)) ds + \frac{1}{N} \sum_{j=1}^k h_j M_j \left(\int_0^t N \beta_j(Z^N(s)) ds \right); \end{aligned} \quad (1)$$

here $x \in A$, $(P_j)_{j=1, \dots, k}$ are i.i.d. standard Poisson processes and $M_j(t) = P_j(t) - t$, $j = 1, \dots, k$, are the associated martingales. The jump directions h_j and their respective rates are chosen in such a way that $Z^N(t) \in A$ almost surely for all t .

The dynamics of the deterministic model corresponding to the stochastic process in (1) are given by the solution of the following ODE:

$$Z(t) := Z^x(t) = x + \int_0^t \sum_{j=1}^k h_j \beta_j(Z(s)) ds. \quad (2)$$

The two models are connected via a Law of Large Numbers (LLN) (see [5]): if the rates β_j are Lipschitz continuous, we have

$$\lim_{N \rightarrow \infty} Z^{N,x}(t) = Z^x(t) \quad \text{a.s. uniformly on compact time intervals } [0, T]. \tag{3}$$

Note also that in the case of Lipschitz continuous rates, equation (2) admits a unique solution on $[0, T]$.

Example 1.1. SIS model. We first consider a simple SIS model without demography ($S(t)$ being the number of susceptible individuals and $I(t) = N - S(t)$ the number of infectious individuals at time t). For $\beta > 0$, we assume that the rate of infections is $\beta S(t)I(t)/N$.¹ In a similar way, we assume for $\gamma > 0$ that the rate of recoveries is $\gamma I(t)$. We ignore immunity due to earlier infection and assume that an infectious individual becomes susceptible after recovery. As population size is constant, we can reduce the dimension of the model by solely considering the proportion of infectious at time t , that is $d = 1$ and $Z_t = I_t$. Using the notation of equations (1) and (2), we have

$$A = [0, 1], \quad h_1 = 1, \quad h_2 = -1, \quad \beta_1(z) = \beta z(1 - z), \quad \beta_2(z) = \gamma z.$$

It is easy to see that the ODE (2) has a disease free equilibrium $\bar{x} = 0$. This equilibrium is asymptotically stable if $R_0 = \beta/\gamma < 1$.² If $R_0 > 1$, \bar{x} is unstable and there exists a second, endemic equilibrium $x^* = 1 - \gamma/\beta$ which is asymptotically stable. While in the deterministic model the proportion of infectious individuals converges to the endemic equilibrium x^* , the disease goes ultimately extinct in the stochastic model.

Example 1.2. A model with vaccination. We consider a (deterministic) model with vaccination and demography described in [2] and its stochastic counterpart. We illustrate the different transitions in Figure 1; here, S and I denote the number of susceptible respectively infectious individuals as before and V denotes the number of vaccinated individuals. We assume that individuals are vaccinated at a certain rate but can lose their protection again; the vaccine is not perfect but decreases the rate of infection by a factor $\sigma \in [0, 1]$: if $\sigma = 1$, the vaccine is useless, if $\sigma = 0$ it protects the individuals perfectly. Furthermore, we assume that individuals are born susceptible and die (at the same rate) independently of their disease status. This ensures that the population size remains constant in the deterministic model. For simplicity, we ensure constant population size by synchronizing births and deaths in the stochastic model. Hence, we can again reduce the dimension of the models; the first coordinate of the process denotes the proportion of infective individuals, the second coordinate denotes the proportion of vaccinated individuals. Using the notations of equations (1) and (2), we obtain

$$\begin{aligned} A &= \{z \in \mathbb{R}_+^2 \mid 0 \leq z_1 + z_2 \leq 1\}, \\ h_1 &= (1, 0)^\top, \quad \beta_1(z) = \beta z_1(1 - z_1 - z_2), \quad h_2 = (1, -1)^\top, \quad \beta_2(z) = \sigma \beta z_1 z_2, \\ h_3 &= (-1, 0)^\top, \quad \beta_3(z) = \gamma z_1, \quad h_4 = (0, 1)^\top, \quad \beta_4(z) = \theta z_2, \quad h_5 = (0, -1)^\top, \quad \beta_5(z) = \eta(1 - z_1 - z_2), \\ h_6 &= (-1, 0)^\top, \quad \beta_6(z) = \mu z_1, \quad h_7 = (0, -1)^\top, \quad \beta_7(z) = \mu z_2. \end{aligned}$$

The basic reproduction number of this model is given by

$$R_0 = \frac{\beta}{\mu + \gamma} \frac{\mu + \theta + \sigma \eta}{\mu + \theta + \eta}, \quad \tilde{R}_0 := \frac{\beta}{\mu + \gamma};$$

here, \tilde{R}_0 denotes the basic reproduction number without vaccination, i.e., $V(0) = 0$ and $\phi = 0$ (see [2]). There exists a disease-free equilibrium \bar{x} ($\bar{x}_1 = 0$) which is asymptotically stable for $R_0 < 1$ and unstable for $\tilde{R}_0 > 1$.

¹The reasoning behind this is the following. Each individual (in particular each infected individual) meets other individuals at rate α . The probability that the encounter is with a susceptible is $S(t)/N$. Denote by p the probability that such an event yields a new infection. Hence the total rate of new infections is $\beta S(t)I(t)/N$, if $\beta = p\alpha$.

² R_0 is the so-called *basic reproduction number*. It denotes the average number of secondary cases infected by one primary case during its infectious period, at the start of the epidemic.

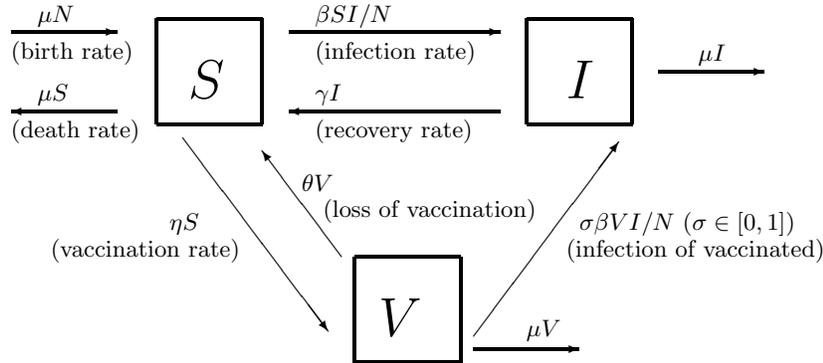


FIGURE 1. Schematic representation of the disease transmission for the model with vaccination by [2].

For $R_0 < 1 < \tilde{R}_0$ (and additional assumptions on the parameters, see [2] again), two endemic equilibria x^* and \tilde{x} ($x_1^*, \tilde{x}_1 > 0$) exist: one (say x^*) is locally asymptotically stable and one is unstable; the disease-free equilibrium is locally asymptotically stable.

1.2. Large deviations

In order to quantify the deviation of the stochastic model from the deterministic model, an LDP is desirable. For the models we have in mind (cf. Examples 1.1 and 1.2), some of the rates $\beta_j : A \rightarrow \mathbb{R}_+$ tend to zero as x approaches the boundary of A (this ensures that the process does not leave the domain A). For this reason, general results for LDPs for Markov jump processes (see [6–8]) do not apply. [9] provides a large deviations principle which allows for such diminishing rates; however, their assumptions do often not apply to the more complicated models in epidemiology.³ [10] provides an LDP on the space $D([0, T]; A)$ ⁴ for a large class of epidemiological models with compact domain A by generalizing the results from [9]; the rate function $I_{T,x}$ of the LDP is given as follows:⁵

$$I_{T,x}(\phi) := \begin{cases} \int_0^T L(\phi(t), \phi'(t)) dt & \text{if } \phi \text{ is absolutely continuous and } \phi(0) = x \\ \infty & \text{else,} \end{cases}$$

where L denotes the Legendre-Fenchel transform:

$$L(x, y) := \sup_{p \in \mathbb{R}^d} \ell(p, x, y) \tag{4}$$

for

$$\ell(p, x, y) = \langle p, y \rangle - \sum_{j=1}^k \beta_j(x) (e^{\langle p, h_j \rangle} - 1).$$

By [9], we have that

$$I_{T,x}(\phi) = 0 \quad \text{if and only if } \phi \text{ solves the ODE (2).} \tag{5}$$

We interpret $I_{T,x}(\phi)$ as the *action* required to follow a trajectory ϕ , or in other words the action required to deviate from the solution of the ODE.

³The assumptions in [9] apply to the SIS model (Example 1.1) but not to the model with vaccination in Example 1.2.

⁴Here, $D([0, T]; A)$ denotes the space of càdlàg functions on $[0, T]$ with domain A equipped with an appropriate metric such that the resulting space is Polish (see, e.g., [11]).

⁵See, e.g., [12] for an introduction to the theory of large deviations and a definition of the rate function.

We are particularly interested in the Freidlin-Wentzell theory (see, e.g., [12,13]). First, we want to compute the *time* $\tau^{N,x}$ at which the process $Z^{N,x}$ leaves the domain of attraction of an asymptotically stable equilibrium O , i.e.,

$$\tau^{N,x} := \inf\{t > 0 \mid Z^{N,x}(t) \in A \setminus O\}.$$

In the case of the SIS model (and $R_0 > 1$) we consider the set $O = (0, 1]$; for the model with vaccination (and appropriate parameter choice, in particular $R_0 < 1 < \tilde{R}_0$), we consider the set

$$O = \{z \in A \mid \lim_{t \rightarrow \infty} Z^z(t) = x^*\}.$$
⁶

In other words, we are interested in time of extinction of the disease⁷ which should eventually become endemic according to the deterministic model.

If A is compact, it can be deduced from the LDP that for all $\delta > 0$,

$$\lim_{N \rightarrow \infty} \mathbb{P} \left[e^{N(\bar{V}-\delta)} < \tau^{N,x} < e^{N(\bar{V}+\delta)} \right] = 1, \tag{6}$$

where \bar{V} is given by the solution of the following optimization problem:

$$\bar{V} := \inf_{z \in A \setminus O} V(x^*, z) \tag{7}$$

$$\text{for } V(x^*, z) := \inf_{T>0, \phi \in D([0,T];A): \phi(0)=x^*, \phi(T)=z} I_{T,x^*}(\phi)$$

(see [10]).

Second, we are interested in the *place*, where Z^N leaves the domain O . Although existing Freidlin-Wentzell theory cannot be applied to the epidemiological models in Examples 1.1 and 1.2, we strongly believe that the following holds (see [12, 13] again for corresponding results for other processes): for any compact set $C \subset \partial O \cap (A \setminus O)$ with $\bar{V} < \inf_{z \in C} V(x^*, z)$, $x \in O$

$$\lim_{N \rightarrow \infty} \mathbb{P} [Z^{N,x}(\tau^{N,x}) \in C] = 0.$$

In particular, if there exists a $z^* \in \partial O \cap (A \setminus O)$ with $V(x^*, z^*) < V(x^*, z)$ for all other $z \in \partial O \cap (A \setminus O)$, then for all $\delta > 0$,

$$\lim_{N \rightarrow \infty} \mathbb{P} [|Z^{N,x}(\tau^{N,x}) - z^*| > \delta] = 0. \tag{8}$$

1.3. Research questions

In order to better understand the differences between deterministic and stochastic compartmental models in epidemiology, the theoretical results above suggest a thorough numerical analysis of the specific models of Examples 1.1 and 1.2. In the remainder of this article, we address the following questions.

1. In order to analyze the deviation of the stochastic model from its Law of Large Numbers limit, we first have to compute the solution of the deterministic models in a reliable way in section 2. In order to avoid instabilities of the solution (such as components becoming negative), we apply a *non-standard finite difference scheme* developed in [4].

⁶We want to remark here that we cannot consider the exit from the domain of attraction of the disease-free equilibrium \bar{x} , i.e.,

$$O = \{z \in A \mid \lim_{t \rightarrow \infty} Z^z(t) = \bar{x}\}.$$

The reason for this is that the absorbing set $\tilde{A} = \{z \in A \mid z_1 = 0\}$ (where z_1 is the first coordinate of z) cannot be left both by the solution of the ODE and by the solution of the SDE.

⁷Respectively in the time when the process is no longer attracted by the endemic equilibrium and hence the probability of extinction is significantly increased, cf. the LLN, equation (3).

2. The computation of the time (and place) of exit from domain depends on the solution of the optimization problem (7). We use dynamic programming to solve the optimization problem numerically and hence deduce the results about the time of exit according to (6) in section 3.
3. Finally, we simulate the process Z^N in section 4. It is not difficult to simulate the process in an exact way by using Gillespie's stochastic simulation algorithm (SSA) (see, e.g., [14, 15]). However, the simulation is rather slow as soon as N becomes large, and we hence adapt the τ -leaping algorithm (see, e.g., [16]).

2. COMPUTATION OF THE SOLUTION OF THE ODE

The large deviations theory will tell us about the deviations between the long-term behaviors of the solution of the SDE and of its ODE LLN limit. Hence we are interested in the long-term behavior of the deterministic epidemiological models described by equation (2). Consequently we need a numerical scheme to compute the solution of the ODE (2) in a "reliable" way. By reliable, we mean in particular that the fixed points of the solution of the (discrete) numerical scheme are the same as the equilibria of the ODE (with the same local stability), which forces the long time behavior of the numerical scheme to resemble that of the solution of the ODE. Explicit finite difference schemes for the solution of ODEs can have undesirable instabilities. In section 2.2 below, we consider the SIS model from Example 1.1 and illustrate that the solution of such a scheme can be oscillating around zero (and in particular become negative) if the parameters (particularly the step-size h) are not chosen with great care. For this reason, we apply a non-standard finite difference (NSFD) method for models in epidemiology which respects the qualitative behavior of the solution of the ODE (see [4]). This NSFD method follows two basic rules for NSFD methods (for a detailed description of the NSFD method see, e.g., [17, 18]). First, quadratic terms of the ODE are approximated in a non-local way (by mixed explicit/implicit terms). Second, non-trivial denominator functions are used for the discrete derivatives.

The remainder of this section is structured as follows. In section 2.1, we describe the specific NSFD method for epidemiological models from [4] and outline in which sense the qualitative long-term behavior of the ODE is respected. In section 2.2, we apply the method to the SIS model (Example 1.1) and compare it to the corresponding explicit scheme. Finally, we apply the method to the model with vaccination (Example 1.2) and compute the boundary separating the domains of attraction of the stable endemic equilibrium and the disease-free equilibrium.

2.1. A NSFD method for epidemiological models

We consider deterministic models for which the ODE (2) can be rewritten in the following form

$$\frac{dZ}{dt} = A(Z)Z + f, \quad Z(0) = x, \quad (9)$$

where $A(Z) = (a_{i,j}(Z))_{i,j=1,\dots,d} \in \mathbb{R}^{d \times d}$ is a non-linear Metzler matrix⁸ and $f \in \mathbb{R}^d$; usually, f represents the rates of birth and immigration into the different compartments (if they exist, cf. Example 1.2).

We follow [4] and discretize time by fixing the time step $h > 0$; we write $t_m = mh$ for $m \in \mathbb{N}$ and approximate the solution Z of the ODE (9) at the discrete time points t_m by the solution \tilde{Z} of the following implicit NSFD scheme (i.e., $\tilde{Z}^m \approx Z(t_m)$):

$$\frac{\tilde{Z}_i^{m+1} - \tilde{Z}_i^m}{\psi(h)} = \sum_{j=1}^d a_{i,j}(\tilde{Z}^m) \tilde{Z}_j^{m+1} + f_i \quad (i = 1, \dots, d), \quad \tilde{Z}^0 = x;$$

the denominator function ψ is specified below and satisfies $\psi(h) = h + o(h^2)$. We first note that this scheme can equivalently be written as

$$(I - \psi(h)A(\tilde{Z}^m))\tilde{Z}^{m+1} = \tilde{Z}^m + \psi(h)f, \quad \tilde{Z}^0 = x; \quad (10)$$

⁸A Metzler matrix is a matrix $A = (a_{i,j})_{i,j=1,\dots,d}$ which satisfies $a_{i,j} \geq 0$ for $i \neq j$.

$(I - \psi(h)A(\tilde{Z}^m))$ is an M -matrix⁹ and hence invertible (see [4]). For the definition of the denominator function, we assume that all equilibria of the ODE (9) are hyperbolic. With $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\varphi(h) = h + o(h^2)$ and $0 < \varphi(z) < 1$ for $z > 0$ (e.g., $\varphi(z) = 1 - \exp(-z)$ or $\varphi(z) = z/(1 + z^2)$), ψ is defined as

$$\psi(h) := \frac{\varphi(Qh)}{Q} \quad \text{for} \quad Q \geq \max \left\{ \frac{|\lambda|^2}{2|\operatorname{Re}(\lambda)|} \right\}, \tag{11}$$

where the maximum is taken over the eigenvalues λ of the Jacobian of the right hand side of (9) at the equilibrium points. In particular, we have $\psi(h) = h + o(h^2)$ as desired.

The NSFD scheme (10) is *elementary stable* in the sense that it has exactly the same fixed points as the continuous ODE (2) it approximates, with the same local stability for all values of h (see [4]). In addition, it can be ensured that the components of the solution of the scheme remain positive by choosing h and hence $\phi(h)$ small enough (see [4] again). For the examples we consider, the solution of the scheme always stays in the “natural domain” of the solution of the ODE (2).

2.2. The SIS model

For the SIS model of Example 1.1, the ODE (2) becomes

$$Z(t) = x + \int_0^t (\beta Z(s)(1 - Z(s)) - \gamma Z(s)) ds = x + \int_0^t (-\beta Z(s)^2 + (\beta - \gamma)Z(s)) ds; \tag{12}$$

recall that in this case $Z(t)$ denotes the number of infectious individuals at time t . This is a Riccati differential equation with constant coefficients which admits an explicit solution:

$$Z(t) = \begin{cases} \frac{(\beta - \gamma)x \exp((\beta - \gamma)t)}{(\beta - \gamma) + \beta x (\exp((\beta - \gamma)t) - 1)} & \text{if } \beta \neq \gamma \\ \frac{x}{1 - \beta x t} & \text{else.} \end{cases} \tag{13}$$

We nevertheless compute the solution of the ODE (12) numerically by a standard (explicit) difference scheme and by the NSFD scheme introduced in section 2.1 in order to illustrate possible instabilities of the explicit scheme.

Let us first consider the following explicit scheme:

$$\frac{\tilde{Z}^{m+1} - \tilde{Z}^m}{h} = \beta \tilde{Z}^m (1 - \tilde{Z}^m) - \gamma \tilde{Z}^m, \quad \tilde{Z}^m = x \tag{14}$$

or equivalently

$$\tilde{Z}^{m+1} = \tilde{Z}^m (1 - \gamma h + \beta h) - h\beta (\tilde{Z}^m)^2, \quad \tilde{Z}^m = x.$$

We note that if, e.g., $\beta = 1/h$ and $\gamma = 2/h$, $\tilde{Z}^m < 0$ for all $m \geq 1$ which is obviously undesirable. If, e.g., $\beta = 1/h$ and $\gamma = 4/h$, the solution is oscillating around zero with $\tilde{Z}^m < 0$ for m odd and $\tilde{Z}^m > 0$ for m even. We illustrate the solution for the latter case in the left picture of Figure 2.

Let us now apply the NSFD scheme to the model. We represent the ODE in the form (9) by setting $A(Z) = -\beta Z + \beta - \gamma$ and define the denominator function ψ according to equation (11) (for $\varphi(z) = 1 - \exp(-z)$). The solution does not suffer from the same instabilities as the solution of the explicit scheme as illustrated in the middle picture of Figure 2. We compare the solutions of the two schemes with the exact solution of the ODE as given in equation (13) (right picture of Figure 2). We want to remark here that the NSFD scheme is of course by no means the only possible scheme which avoids instabilities in this simple case. However, [4] provides us with a theoretical basis which ensures that instabilities do not occur even for much more complicated models.

⁹An M -matrix is a matrix A with $a_{i,j} \leq 0$ for $i \neq j$ which is invertible such that all entries of its inverse are non-negative.

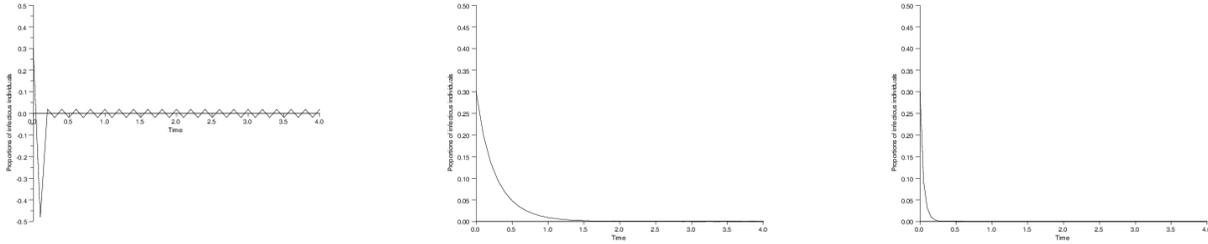


FIGURE 2. The left picture denotes the solution of the explicit finite difference scheme (14). The middle picture denotes the solution of the NSFD scheme. The right picture denotes the exact solution of the ODE according to equation (13). $Z(0) = 0.3$, $T = 4$, $h = 0.1$, $\beta = 4/h = 40$, $\gamma = 2/h = 20$. We choose here a large step size h on purpose, to show how the two schemes behave for large step size. Of course, both improve if we reduce h .

2.3. The model with vaccination

We now consider the model with vaccination from Example 1.2. We denote the proportions of susceptible, vaccinated and infectious individuals at time t by $Z_1(t)$, $Z_2(t)$ and $Z_3(t) = 1 - Z_1(t) - Z_2(t)$, respectively. Note that we cannot reduce the dimension from three to two if we want to represent the ODE in the form (9) (with appropriate Metzler matrix A). We set

$$A(Z) = \begin{pmatrix} -\beta Z_3 - \mu - \eta & \theta & \gamma \\ \eta & -\sigma\beta Z_3 & 0 \\ \beta Z_3 & \sigma\beta Z_3 & -\mu - \gamma \end{pmatrix} \quad \text{and} \quad f = (\mu, 0, 0)^\top$$

and choose the denominator function ψ according to equation (11) for $\varphi(z) = 1 - \exp(-z)$. We readily observe that the relevant assumptions for the NSFD scheme are satisfied and apply the scheme for time step $h = 0.1$ to different initial values $Z(0)$ and a set of parameters (cf. Figure 3) such that there exists a stable disease-free equilibrium ($\bar{x} = (0.14, 0.86, 0)^\top$), a stable endemic equilibrium ($x^* = (0.24, 0.45, 0.31)^\top$) and an unstable endemic equilibrium ($\tilde{x} = (0.23, 0.59, 0.18)^\top$); cf. the corresponding discussion in Example 1.2). We illustrate the evolution of the proportions of vaccinated and infectious individuals for different initial values in Figure 3. While the disease eventually dies out in the left picture, it becomes endemic in the right picture.

The NSFD scheme also allows us to approximately compute the *characteristic boundary*, i.e., the boundary separating the domains of attractions of the two stable equilibria \bar{x} and x^* (in other words: the set of points for which the solution of the ODE converges to the unstable equilibrium \tilde{x}). We illustrate the characteristic boundary in Figure 4.

3. COMPUTATION OF \bar{V}

In order to answer the question of when the disease dies out, we have to compute the quantity \bar{V} (cf. equations (6) and (7)). In other words, we have to solve a control problem. We first fix the time horizon $T > 0$. We recall that we only have to consider absolutely continuous trajectories (as else $I_{T,x}(\phi) = \infty$) and hence we control the trajectory ϕ via its “derivative”; for $t \in [0, T]$, $s \in [t, T]$ and an integrable function α , we write

$$\phi^\alpha(s) := \phi(s) = \phi(t) + \int_t^s \alpha(r) dr;$$

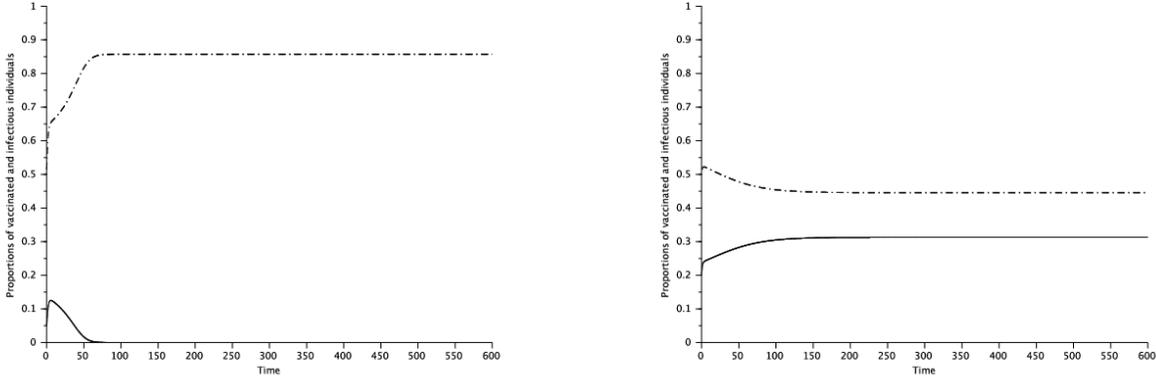


FIGURE 3. Evolution of the proportions of vaccinated individuals (dashed lines) and infectious individuals (solid lines) for initial values $Z_2(0)=0.5, Z_3(0) = 0.05$ (left picture) respectively $Z_2(0)=0.5, Z_3(0) = 0.2$ (right picture). $T = 600, \beta = 3.6, \gamma = 1, \eta = 0.3, \theta = 0.02, \mu = 0.03, \sigma = 0.1$.

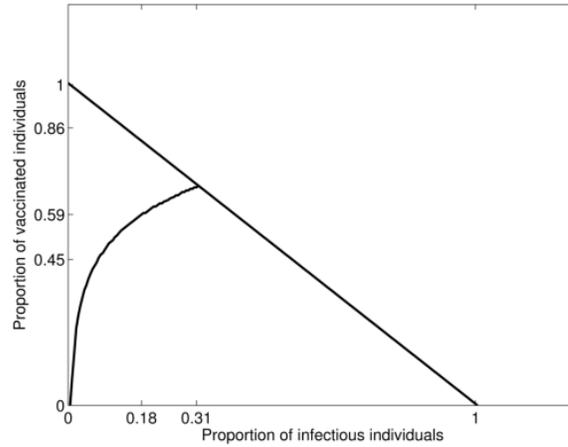


FIGURE 4. The characteristic boundary separating the domains of attractions of $\bar{x} = (0, 0.86)^\top$ (on the left) and $x^* = (0.31, 0.45)^\top$ (on the right). The unstable equilibrium $\tilde{x} = (0.18, 0.59)^\top$ is on the characteristic boundary. Note the slight abuse of notation due to the reduction from dimension three to two. The parameters are as in Figure 3.

for $x \in O$, we define the set of admissible controls starting from (t, x) by

$$\mathcal{A}^T(t, x) := \left\{ \alpha : [t, T] \rightarrow \mathbb{R}^d \text{ integrable} \mid x + \int_t^s \alpha(r) dr \in \bar{O} \forall s \in [t, T] \text{ and } x + \int_t^T \alpha(r) dr \in A \setminus O \right\}. \quad (15)$$

The cost of an admissible control α starting from (t, x) is given by

$$J^T(t, x, \alpha) = \int_t^T L(\phi^\alpha(s), \alpha(s)) ds,$$

and thus the value function of the optimization problem becomes ($t \in [0, T)$)

$$v^T(t, x) = \inf_{\alpha \in \mathcal{A}^T(t, x)} J(t, x, \alpha) \quad (16)$$

and

$$v^T(T, x) = \begin{cases} 0 & \text{if } x \notin O \\ \infty & \text{else.} \end{cases}$$

The value function thus has a singularity at the terminal time T which is due to the constraint $\phi(T) \notin O$. For the models we have in mind (cf. Examples 1.1 and 1.2), it is possible to move on $\partial O \cap (A \setminus O)$ by following the solution of the ODE (2) (which is free of cost, cf. (5)).¹⁰ For this reason, we have $v^T(t, x) \leq v^{\tilde{T}}(t, x)$ for $\tilde{T} < T$ and therefore

$$\lim_{T \rightarrow \infty} v^T(0, x^*) = \bar{V}.$$

We want to remark here that any admissible control α with $\phi^\alpha(T) \notin (A \setminus O) \cap \partial O$ (cf. (5) again) cannot be optimal. Therefore, it would be equivalent to require the weaker restriction $x + \int_0^s \alpha(r) dr \in A$ for $s \in [0, T]$ in equation (15).

In the following, we describe a discrete dynamic programming algorithm for the numerical approximation of \bar{V} (section 3.1) and apply this algorithm to the models from Example 1.1 (section 3.2) and Example 1.2 (section 3.3).

3.1. Approximation by discrete-time dynamic programming

In order to numerically compute the solution of the optimization problem (16), we apply discrete-time dynamic programming. To this end, we discretize time and consider $n + 1$ ($n \in \mathbb{N}$) equidistant time points

$$0 = t_0 < \dots < t_n = T, \quad \text{i.e.,} \quad t_m^n = t_m = m\Delta t \text{ for } \Delta t = \frac{T}{n}.$$

For $j = 0, \dots, n - 1$ let $\tilde{\alpha} = (\tilde{\alpha}(t_m))_{m=j, \dots, n-1}$, where $\tilde{\alpha}(t_m) \in \mathbb{R}^d$. We consider piecewise constant trajectories and define recursively for $x \in O$,

$$\tilde{\phi}^{\tilde{\alpha}}(t_j) = \tilde{\phi}(t_j) = x, \quad \tilde{\phi}(t_{m+1}) = \tilde{\phi}(t_m) + \tilde{\alpha}(t_m)\Delta t \quad (m = j, \dots, n - 1).$$

We define the set of admissible controls by

$$\tilde{\mathcal{A}}^{T, n}(t_j, x) := \left\{ \tilde{\alpha} \mid \tilde{\phi}^{\tilde{\alpha}}(t_m) \in A \forall m = j, \dots, n - 1 \text{ and } \tilde{\phi}^{\tilde{\alpha}}(t_n) \in (A \setminus O) \cap \partial O \right\}.$$

The cost of an admissible control $\tilde{\alpha} \in \tilde{\mathcal{A}}^{T, n}(t_j, x)$ is given by

$$\tilde{J}^{T, n}(t_j, x, \tilde{\alpha}) := \Delta t \sum_{m=j}^{n-1} L(\tilde{\phi}^{\tilde{\alpha}}(t_m), \tilde{\alpha}(t_m))$$

and the value function is given by

$$\tilde{v}^{T, n}(t_j, x) := \inf_{\tilde{\alpha} \in \tilde{\mathcal{A}}^{T, n}(t_j, x)} \tilde{J}^{T, n}(t_j, x, \tilde{\alpha}). \quad (17)$$

¹⁰For the SIS model, we have $\partial O \cap (A \setminus O) = \{0\}$. For the model with vaccination, we have

$$\partial O \cap (A \setminus O) = \{z \in A \mid \lim_{t \rightarrow \infty} Z^z(t) = \tilde{x}\}$$

(recall that \tilde{x} is the unstable endemic equilibrium).

The constraint implies that

$$\tilde{v}^{T,n}(t_{n-1}, x) = \Delta t \inf_{\tilde{\alpha} \in \mathbb{R}^d: x + \tilde{\alpha} \Delta t \in (A \setminus O) \cap \partial O} L(x, \tilde{\alpha}). \quad (18)$$

Furthermore, for $j < n - 1$, the dynamic programming principle implies the following Bellman equation:

$$\tilde{v}^{T,n}(t_j, x) = \inf_{\tilde{\alpha} \in \mathbb{R}^d: x + \tilde{\alpha} \Delta t \in A} \left\{ L(x, \tilde{\alpha}) \Delta t + \tilde{v}^{T,n}(t_{j+1}, x + \tilde{\alpha} \Delta t) \right\}.$$

Hence the solution of the optimization problem can be computed backward in time.

We need next to discretize space. Note also that in the SIS model, the functional L is explicitly known, while this is not the case in the vaccination model. For that reason we now distinguish between the two cases.

3.2. The SIS model

We discretize space via a grid of the set $A = [0, 1]$ with fineness $\Delta x = 1/\bar{n}$, for some new integer \bar{n} . Hence the grid points are given by

$$x^i := i \Delta x \quad \text{for } i = 0, \dots, \bar{n}.$$

Given a function f which is only defined on the grid points, we define values of the function for arbitrary $x \in A$ by its linear interpolation f_{int} . In other words, we obtain for $x \in A$,

$$f_{\text{int}}(x) := \bar{n} \sum_{i=0}^{\bar{n}} \mathbb{1}_{\{x \in [x^i, x^{i+1}]\}} \left\{ f(x^i)(x^{i+1} - x) + f(x^{i+1})(x - x^i) \right\}.$$

We obtain the following (backward in time) algorithm for the solution of the optimization problem (17).

1. For x on the grid,

$$\bar{v}^{T,n,\bar{n}}(t_{n-1}, x) = \Delta t \inf_{\tilde{\alpha} \in \mathbb{R}^d: x + \tilde{\alpha} \Delta t \in (A \setminus O) \cap \partial O} L(x, \tilde{\alpha}).$$

We denote the solution of this minimization problem by $\tilde{\alpha}^*(t_{n-1}) = \tilde{\alpha}^*(t_{n-1}, x)$.

2. For $m < n - 1$, x on the grid,

$$\bar{v}^{T,n,\bar{n}}(t_m, x) = \inf_{\tilde{\alpha} \in \mathbb{R}^d: x + \tilde{\alpha} \Delta t \in A} \left\{ L(x, \tilde{\alpha}) \Delta t + \bar{v}_{\text{int}}^{T,n,\bar{n}}(t_{m+1}, x + \tilde{\alpha} \Delta t) \right\}.$$

We denote the solution of this minimization problem by $\tilde{\alpha}^*(t_m) = \tilde{\alpha}^*(t_m, x)$.

Note that this algorithm involves the solution of $n \times \bar{n}$ minimization problems. We recover an approximation $\tilde{\phi}^* = (\phi^*(t_m))_{m=0, \dots, n}$ of the optimal trajectory starting from the stable equilibrium x^* , by the following procedure (including another linear interpolation).

3. For $m = 0, \dots, n - 1$,

$$\tilde{\phi}^*(t_0) = x^*, \quad \tilde{\phi}^*(t_{m+1}) = \phi^*(t_m) + \tilde{\alpha}_{\text{int}}^*(t_m, \phi^*(t_m)) \Delta t.$$

It is well-known that under appropriate assumptions, the solution of this procedure converges to the solution of the optimization problem (16) (see, e.g., [19, 20]); the optimal coupling of time and space discretization is $\Delta t = \Delta x$. We want to remark here that standard assumptions do not apply to our model. For instance, the function L is not Lipschitz continuous, even $L(x, y) \rightarrow \infty$ for $x \rightarrow \partial A$ is possible.¹¹ Furthermore, the terminal

¹¹This is e.g. the case for $y < 0$ and $x \rightarrow 0$. Note however that the action required to approach the boundary, e.g., at constant speed one is finite and we have $\int_0^t L(t-s, -1) ds \rightarrow 0$ as $t \rightarrow 0$.

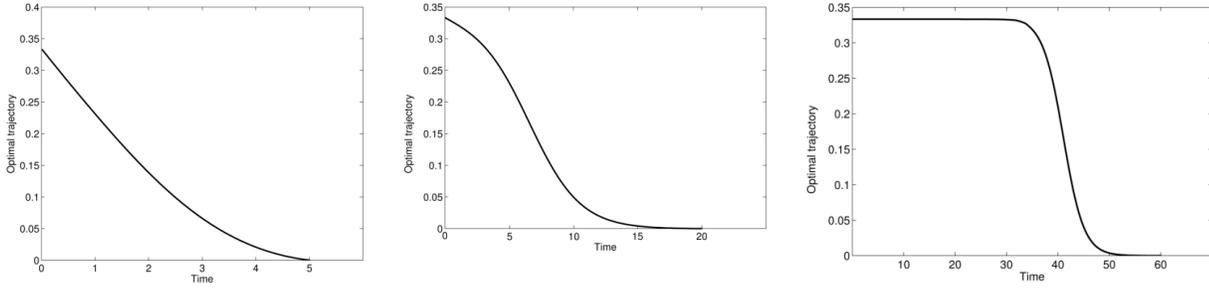


FIGURE 5. Optimal trajectory $\tilde{\phi}^*$ for $T = 5$ (left picture), $T = 20$ (middle picture) and $T = 60$ (right picture); $\Delta t = \Delta x = \frac{1}{100}$, $\beta = 1.5$ and $\gamma = 1$.

constraint results in a singularity of the value function at terminal time T . We are not aware of any analytical results including our set-up. We nevertheless have reasons to believe that

$$\bar{V} \approx \bar{v}_{\text{int}}^{T,n,\bar{n}}(t_0, x^*) \quad \text{for large enough } T.$$

It is not in the scope of this article to provide a thorough analytical proof of the convergence of the algorithm.

We readily observe that for the SIS model, the Legendre-Fenchel transform is given explicitly by

$$L(x, y) = y \log \tilde{\theta} - \beta x(1-x)(\tilde{\theta} - 1) - \gamma x \left(\frac{1}{\tilde{\theta}} - 1 \right),$$

where

$$\tilde{\theta} := \frac{y + \sqrt{y^2 + 4\beta\gamma x^2(1-x)}}{2\beta x(1-x)}.$$

Furthermore, the boundary consists of the single point $\{0\}$; hence equation (18) becomes

$$\tilde{v}^{T,n}(t_{n-1}, x) = \Delta t L(x, -x/\Delta t).$$

We consider the following parameters: $\beta = 1.5$, $\gamma = 1$; in particular, we have $R_0 = 1.5 > 1$ and the endemic equilibrium $x^* = 1/3$ is stable. We apply the dynamic programming algorithm described in section 3.1 with $\Delta t = \Delta x = 1/100$ and $T = 5, 10, 20, 40, 60$ (i.e., $\bar{n} = 100$ and $n = 500, 1000, 2000, 4000, 6000$). We obtain

$$\begin{aligned} v_{\text{int}}^{5,n,\bar{n}}(t_0, x^*) &= 0.0953, & v_{\text{int}}^{10,n,\bar{n}}(t_0, x^*) &= 0.0757, & v_{\text{int}}^{20,n,\bar{n}}(t_0, x^*) &= 0.0705, \\ v_{\text{int}}^{40,n,\bar{n}}(t_0, x^*) &= 0.0702, & v_{\text{int}}^{60,n,\bar{n}}(t_0, x^*) &= 0.0702. \end{aligned}$$

We hence deduce that $\bar{V} \approx 0.0702$ and thus $\mathbb{E}[\tau^{N,x}] \approx \exp(0.0702N)$ for large N . We illustrate the solution of the optimization problem in Figure 5 for $T = 5, 20, 60$. We observe that it is the cheapest to move slowly near the stable equilibria x^* and the unstable equilibrium $\bar{x} = 0$.

3.3. The model with vaccination

Unlike in section 3.2, the place of exit from O is a priori not clear for the model with vaccination (cf. Example 1.2). It seems plausible that it is cheapest to exit O at the unstable endemic equilibrium \tilde{x} . Indeed, if T is chosen large enough, an optimal (or near optimal) trajectory ϕ can move on the boundary towards \tilde{x} after the exit time without generating any extra cost (cf. (5)); in this case we have $\lim_{t \rightarrow \infty} \phi(t) = \tilde{x}$. Hence, we can assume that \tilde{x} is indeed the place of exit for the numerical computation of \bar{V} . Therefore, equation (18) becomes

$$\tilde{v}^{T,n}(t_{n-1}, x) = \Delta t L(x, (\tilde{x} - x)/\Delta t).$$

However, the Legendre-Fenchel transform L cannot be computed explicitly. Here we use the following definition of L , which is equivalent to (4) (see [10]):

$$L(x, y) := \begin{cases} \inf_{\mu \in B_{x,y}} \tilde{\ell}(\mu, x) & \text{if } B_{x,y} \neq \emptyset \\ \infty & \text{if not,} \end{cases}$$

where for $\mu \in B_{x,y}$,

$$\tilde{\ell}(\mu, x) := \sum_j \left\{ \beta_j(x) - \mu_j + \mu_j \log \left(\frac{\mu_j}{\beta_j(x)} \right) \right\};$$

and

$$B_{x,y} := \left\{ \mu \in \mathbb{R}_+^k \mid \mu_j > 0 \text{ only if } \beta_j(x) > 0 \text{ and } y = \sum_j \mu_j h_j \right\}.$$

Since we seek to minimize

$$\bar{V} := \inf_{T, \phi \text{ s.t. } \phi(0)=x^*, \phi(T)=\tilde{x}} \int_0^T L(\phi(t), \phi'(t)) dt,$$

our discretized dynamic programming equation reads: for each $x \in \mathcal{G}$ (the space grid),

$$\begin{aligned} v(t_k, x) &= \inf_{x' \in \mathcal{G}} \left\{ L \left(x, \frac{x' - x}{\Delta t} \right) \Delta t + v(t_{k+1}, x') \right\} \\ &= \inf_{x' \in \mathcal{G}} \left\{ \tilde{\ell}(\hat{\mu}(x, x'), x) \Delta t + v(t_{k+1}, x') \right\}, \end{aligned}$$

where $\hat{\mu}(x, x')$ is the argument which realizes the following minimum

$$\inf_{\mu, x + \Delta \sum_j \mu_j h_j = x'} \tilde{\ell}(\mu, x).$$

Note that here we avoid to consider interpolation between points on the grid. We have implemented this algorithm on a grid whose one of the axis is the line passing through x^* and \tilde{x} , with the following values of the parameters : $\beta = 3.6$, $\theta = 0.02$, $\mu = 0.03$, $\eta = 0.3$, $\sigma = 0.1$ and $\gamma = 1$.

The approximation of \bar{V} which we have found is 0.3891, with $\Delta t = 0.05$, and the $\Delta x_1 = \Delta x_2 = 0.005$. This means that $\mathbb{E}[\tau^{N, x^*}] \approx \exp(0.3891N)$ for large N . Note that if Δt is not small enough compared to Δx , since we restrict the above minimization to $x' \in \mathcal{G}$, we cannot allow small speed, and we get larger values of \bar{V} . We expect that more accurate results could be obtained, by doing the minimization over x' on a finer grid, which necessitates to interpolate $v(t_{k+1}, \cdot)$ between points on the grid. We have not been able to implement this refinement, by lack of time. This will be done in the future. It should improve the numerical estimate of \bar{V} . The above quantity should be considered as an upper bound of the true value.

4. SIMULATION OF Z^N

In this section, we describe known exact and approximate simulation methods for the simulation of the solution Z^N of equation (1). To this end, we fix N ; with the notations $\nu_j = \frac{h_j}{N}$ and $a_j(z) = N\beta_j(z)$, equation (1) can be rewritten as

$$Z^N(t) = x + \sum_{j=1}^k \nu_j P_j \left(\int_0^t a_j(Z^N(s)) ds \right). \quad (19)$$

Note that our ultimate goal would be to simulate the value of τ^N which, from the above Large Deviations results, is huge, as soon as N is large. This may be seen as impossible. However, this is similar to the problem of rare event simulation, for which there exist specific techniques, which we intend to implement in the future. See also the discussion below in section 4.7.2. In the framework of this paper, we present some preliminary numerical work, which aims only at simulating the stochastic process Z^N over a given finite time horizon.

In the following, we first introduce the so-called stochastic simulation algorithm (SSA) proposed in [14, 15] (section 4.1). The SSA allows for the exact simulation of the process Z^N . Unfortunately, this algorithm is rather slow in some circumstances, typically when there are too many jumps of the Poisson processes on a fixed time interval. Therefore, we discuss faster approximate simulation methods in the subsequent sections: in section 4.2, we introduce the explicit τ -leaping algorithm (see [16]) in the form of [21]; in section 4.3, we discuss modifications of the τ -leaping algorithm: implicit and mid-point τ -leaping (section 4.3.1) and the algorithm from [22] which ensures that the simulated process remains in its domain, e.g., by preventing the components of the process from becoming negative (section 4.3.2). We apply the SSA and the modified τ -leaping algorithm from section 4.3.1 to the specific models from Examples 1.1 and 1.2 in sections 4.4 and 4.5, respectively. We conclude by addressing open research questions concerning the simulation of the process Z^N in section 4.7.

4.1. The stochastic simulation algorithm

A mathematically exact procedure for simulating the evolution of the process Z^N given by equation (19) is the stochastic simulation algorithm (SSA) first proposed in [14, 15]. The SSA changes the value of the solution at each jump time of one of the Poisson processes. The simplest implementation of the SSA is the so-called *direct method*:

Whenever $t < T$:

1. While in state z at time t , evaluate all functions $a_j(z)$ and their sum $a_0(z) := \sum_{j=1}^k a_j(z)$.
2. Generate the time increment τ as the realization of an exponential random variable with parameter $a_0(z)$.
3. Choose an index $j \in \{1, \dots, k\}$ according to the probability distribution $\frac{a_j(z)}{a_0(z)}$ ($j = 1, \dots, k$).
4. Update
 - i. $t \leftarrow t + \tau$,
 - ii. $z \leftarrow z + \nu_j$.
5. Record (t, z) . Return to 1. if $t + \tau < T$; else stop.

Carrying out step 3. is mathematically straightforward: we draw one realization r of a uniform random variable on $[0, 1]$; then, j is the smallest positive integer for which

$$\sum_{i=1}^j a_i(z) > r a_0(z).$$

The SSA is mathematically exact. However, the task of explicitly simulating each jump makes the SSA too slow for practical implementation in some circumstances. The reason is that for large N , there are too many jumps on any fixed time interval.

4.2. The explicit Poisson τ -leaping algorithm

A faster but approximate stochastic simulation procedure is the explicit Poisson τ -leaping algorithm (see [16]). The basic idea of this procedure is to advance the system by a *preselected* time increment τ (in contrast to the *randomly generated* time increment τ in the SSA). On the one hand, τ should be chosen large enough in order to ensure that “many” jumps occur during this length of time (as else, the savings in computation time are not significant); on the other hand, it should be small enough such that the values of the functions a_j are unlikely to change “significantly” (as else, the procedure is not accurate enough). The latter restriction is called the *leap condition*. A strategy due to [21] for satisfying this condition is to require that the changes of the functions a_j

during a leap τ are “likely” to be bounded by $\epsilon a_0(z)$; here, ϵ ($0 < \epsilon \ll 1$) is the error control parameter. In other words,

$$|a_j(Z^N(t + \tau)) - a_j(Z^N(t))| \leq \epsilon a_0(Z^N(t)) \quad \text{for } j = 1, \dots, k \text{ with high probability.}$$

We estimate the largest value of τ that meets this particular requirement as follows. First, we compute

$$f_{jj'}(z) := \sum_{i=1}^d \frac{\partial a_j(z)}{\partial z_i} \nu_{ij'} \quad \text{for } j, j' = 1, \dots, k; \tag{20}$$

here ν_{ij} denotes the i -th component of the vector ν_j . Subsequently, we calculate

$$\mu_j(z) := \sum_{j'=1}^k f_{jj'}(z) a_j(z) \quad \text{and} \quad \sigma_j^2(z) := \sum_{j'=1}^k f_{jj'}^2(z) a_j(z) \quad \text{for } j = 1, \dots, k \tag{21}$$

and set

$$\tau = \tau(z) := \min_{j=1, \dots, k} \left\{ \frac{\epsilon a_0(z)}{|\mu_j(z)|}, \frac{\epsilon^2 a_0^2(z)}{\sigma_j^2(z)} \right\}. \tag{22}$$

Note that for a standard Poisson process P ,

$$\begin{aligned} P_j \left(\int_0^{t+\tau} a_j(Z^N(s)) ds \right) - P_j \left(\int_0^t a_j(Z^N(s)) ds \right) &\stackrel{d}{=} P \left(\int_t^{t+\tau} a_j(Z^N(s)) ds \right) \\ &\approx P(a_j(Z^N(t))\tau) = P(a_j(z)\tau) =: p_j \end{aligned}$$

and thus

$$Z^N(t + \tau) \stackrel{d}{\approx} Z^N(t) + \sum_{j=1}^k p_j \nu_j = z + \sum_{j=1}^k p_j \nu_j.$$

This procedure for the selection of τ yields the following τ -leaping algorithm.

Whenever $t < T$:

1. While in state z at time t , evaluate the functions a_j and their sum $a_0(z) = \sum_{j=1}^k a_j(z)$.
2. Compute τ via the formulas (20)-(22).
3. Fix a small integer n (typically, e.g., $n = 10$) and a moderate integer \bar{n} (typically, e.g., $\bar{n} = 100$).
 - i. If $\tau < \frac{n}{a_0(z)}$, reject τ and execute \bar{n} iterations of the SSA. Then, return to step 1.
 - ii. If $\tau \geq \frac{n}{a_0(z)}$, proceed to step 4.
4. For $j = 1, \dots, k$, generate p_j as a Poisson random variable with mean $a_j(z)\tau$.
5. Update
 - i. $t \leftarrow t + \tau$,
 - ii. $z \leftarrow z + \sum_{j=1}^k p_j \nu_j$.
6. Record (t, z) . Return to 1. if $t + \tau < T$; else stop.

Let us shortly comment on this algorithm. p_j represents an approximation of the number of jumps of type j during the time interval $[t, t + \tau]$. We remark again that the approximation by p_j in step 4. is justified as long as $a_j(z)$ remains more or less constant over the subsequent time interval of length τ . The reason for step 3. i. is the following. $\frac{1}{a_0(z)}$ is the mean time step to the next jump when applying the SSA. If τ is not significantly larger than $\frac{1}{a_0(z)}$, the computation of τ (via equations (20)-(22)) is no more efficient than applying the SSA directly (which is also of course more accurate); as this property is likely to persist for a while, it is more appropriate to apply the SSA for some time before returning to the τ -selection procedure again.

The explicit τ -leaping algorithm above has been shown to provide accurate simulation results; furthermore, it is substantially faster than the SSA for many “not-too-stiff” systems¹² (see [22]; see also [24] for theoretical evidence for this). We illustrate the gain of computation time by numerical experiments in sections 4.4 and 4.5 below for a modification of this algorithm (cf. section 4.3.2).

4.3. Modifications of the τ -leaping algorithm

There exist several variants of the above described explicit τ -leaping algorithm. We now describe some of these.

4.3.1. Implicit and mid-point τ -leaping algorithms

One may want to freeze the rates not at the value of the solution at the beginning of the discretization interval, but rather at an approximate value of the position at the end of that interval, or else at the mid-point.

The “implicit” τ -leaping algorithm is obtained by replacing $a_j(Z^N(t))\tau$ with

$$a_j(Z^N(t) + \tau b(Z^N(t)))\tau,$$

and the mid-point τ -leaping algorithm consists in replacing the same quantities by

$$a_j(Z^N(t) + \tau b(Z^N(t))/2)\tau.$$

In [24] it is proved that the error in the mid-point τ -leaping algorithm is less than the one produced by the explicit τ -leaping algorithm.

4.3.2. A modified τ -leaping algorithm

As most numerical schemes, the τ -leaping algorithm has the disadvantage that the approximate solution might very well jump over the boundaries which the solution of the stochastic equation (19) can never cross. More precisely, the solution of (1) satisfies certain constraints. For example, in the case of constant total population, we have that $Z_i^N(t) \geq 0$, $1 \leq i \leq d$, and $Z_0^N(t) := 1 - \sum_{i=1}^d Z_i^N(t) \geq 0$. We would like that the approximate solution satisfies these constraints as well. However, this might not be the case with a standard τ -leaping algorithm. In particular, if for some $0 \leq i \leq d$, $Z_i^N(t)$ is very small, then some of the Poisson driving processes may push it over to a negative value during the next time interval of length τ . Because of that, at the beginning t of each time step, we define

$$L_j := \min_{1 \leq i \leq d, \nu_{ij} < 0} \lfloor N Z_i^N(t) \rfloor$$

and choose a critical value n_c . The critical directions will be chosen by comparing L_j to n_c .

The modified Poisson τ -leaping procedure requires that a critical Poisson process cannot jump more than once in a time interval of length τ . That makes it impossible for a critical Poisson process to drive any component of Z^N to a negative value. The theoretical justification for each step in this modified τ -leaping procedure is given in [22].

Whenever $t < T$:

1. While in state z at time t , evaluate the functions $a_j(z)$ and their sum $a_0(z) = \sum_{j=1}^k a_j(z)$.
2. Identify the currently critical Poisson processes as those P_j for which

$$a_j(z) > 0 \quad \text{and} \quad L_j < n_c.$$

3. Modify the τ -selection procedure from equations (20)-(22) by only taking into account those jump directions j' corresponding to the *non-critical* Poisson processes. Denote the resulting time step by τ' . If there is no non-critical Poisson process, set $\tau' = \infty$.

¹²I.e., systems for which the difference between the smallest and the largest jump rates is not too large (see, e.g., [22]).

4. Fix a small integer n (typically, e.g., $n = 10$) and moderate integer \bar{n} (typically, e.g., $\bar{n} = 100$).
 - i. If $\tau' < \frac{n}{a_0(z)}$, reject τ' and execute \bar{n} iterations of the SSA. Then, return to step 1.
 - ii. If $\tau' \geq \frac{n}{a_0(z)}$ proceed to step 5.
5. Compute the sum $a_0^c(z)$ of the functions $a_j(z)$ of the critical Poisson processes. Generate τ'' as a realization of the exponential random variable with parameter $a_0^c(z)$.
6.
 - i. If $\tau' < \tau''$, set $\tau = \tau'$. For all the critical Poisson processes P_j , set $p_j = 0$. For all the non-critical Poisson processes P_j , generate p_j as a realization of the Poisson random variable with mean $a_j(z)\tau$.
 - ii. If $\tau'' \leq \tau'$, set $\tau = \tau''$. Generate an integer j_c according to the probability distribution $a_j(z)/a_0^c(z)$ (i.e., j runs over the index values of the critical Poisson processes). Set $p_{j_c} = 1$, and for all the other critical Poisson process, set $p_j = 0$. For all the non-critical Poisson processes P_j , generate p_j as a realization of the Poisson random variable with mean $a_j(z)\tau$.
7. Update
 - i. $t \leftarrow t + \tau$,
 - ii. $z \leftarrow z + \sum_{j=1}^d p_j \nu_j$.
8. If z is not in its domain, replace $\tau' \leftarrow \tau'/2$, and return to step 6.
9. Record (t, z) . Return to 1 if $t + \tau < T$; else stop.

Remark 4.1. It may seem strange to use a uniform critical value n_c , independent of the value of the rate of each Poisson process. It would be natural to compare each L_j with the corresponding $a_j(Z^N(t))\tau$, which however would require to compute a first value of τ . We intend to search for improvements in that direction in the future.

Remark 4.2. Although the modified τ -leaping algorithm is designed so as to reduce the chance of a negative value of any of the $Z_i^N(t)$, it cannot completely rule them out. This is the reason for step 8. of the algorithm. We suspect that this step might introduce a bias in the simulation, in that it tends to delete simulations with exceptionally large values of some of the Poisson process increments, and also to reduce the chance of approaching or hitting any of the boundaries, which is specially annoying for our purposes.

We intend to check two other possible approaches. The first one consists in projecting the computed point, if it is outside the domain. The second consists in replacing τ by $\tau/2$, but without resimulating a new independent value of the Poisson process increment, but rather by simulating the new increments according to the conditional law of $P_j(t + \tau) - P_j(t)$, given the previously simulated value of $P_j(t + \tau) - P_j(t)$, which, if $P_j(t + \tau) - P_j(t) = k$, follows a binomial law with parameters k and τ'/τ .

4.4. The SIS model

In this section, we apply the simulation algorithms from sections 4.1 and 4.3.2 to the SIS model (Example 1.1). We consider an endemic situation ($\beta = 1.5$, $\gamma = 1$, hence $R_0 = 1.5 > 1$ and $x^* = 1/3$) and compare the performance of the SSA to the modified τ -leaping algorithm. If the population size is relatively small, $N = 2000$, we observe that the computation times are similar for both algorithms (even smaller for the SSA). A closer look at the realized trajectory of the τ -leaping procedure confirms that almost all jumps are of size $1/2000$. This suggests that in most cases, the selected time step τ' is rejected and the SSA is applied (cf. steps 3. and 4. i. of the algorithm). Hence, there is no significant gain in computation time. If we let $N = 20000$, the number of jumps (and hence the computation time) for the SSA increases approximately by a factor ten. The τ -leaping algorithm is significantly faster which suggests that the time step τ' is now rarely refused. While the SSA becomes inefficient for very large time horizons (e.g., $T = 600$) and population sizes (e.g., $N = 200000$), the process can still be simulated efficiently by the τ -leaping algorithm.

We illustrate realized trajectories of the SSA and the modified τ -leaping algorithm in Figures 6 and 7, respectively (left pictures: $N = 2000$, right pictures: $N = 20000$). We observe that they approach the solution of the ODE as N increases for both algorithms as predicted by the LLN (cf. (3); see [5]).

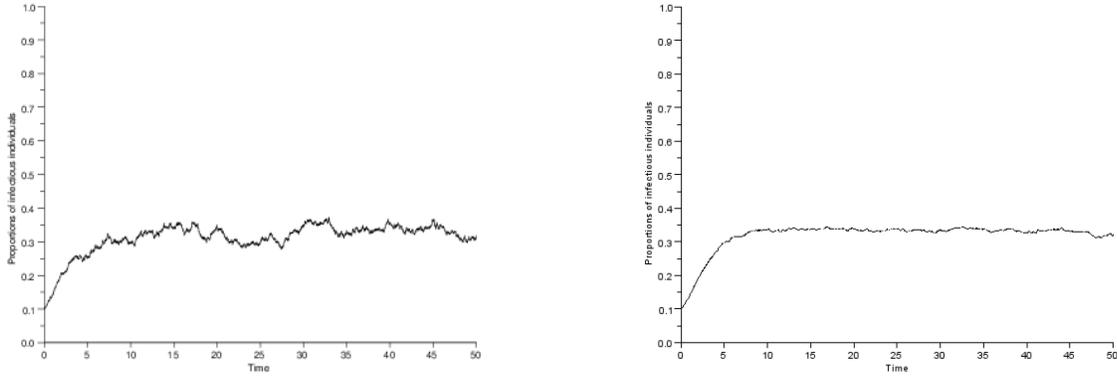


FIGURE 6. Realized trajectories of the SSA for $N = 2000$ (left picture) and $N = 20000$ (right picture). $Z(0) = 0.1$, $T = 50$, $\beta = 1.5$, $\gamma = 1$.

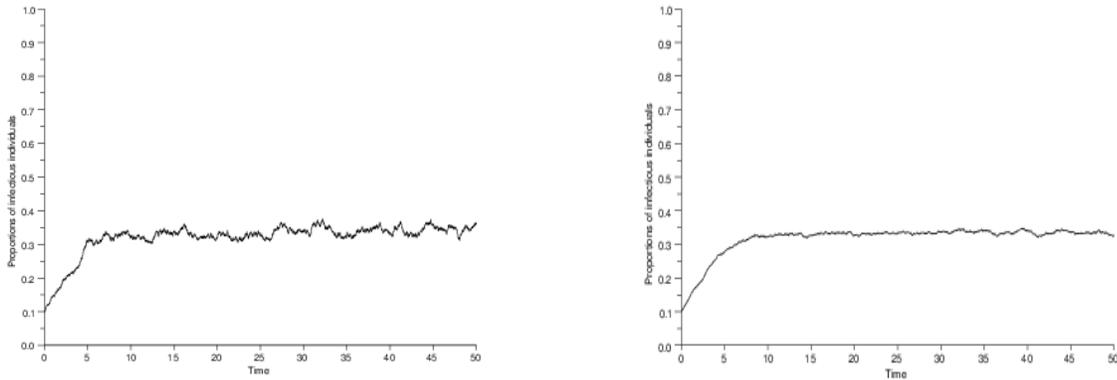


FIGURE 7. Realized trajectories of the modified τ -leaping procedure for $N = 2000$ (left picture) and $N = 20000$ (right picture). The parameters are as in Figure 6.

4.5. The model with vaccination

We now apply the same algorithms to the model with vaccination (Example 1.2). We consider the bistable situation from section 2.3 (cf. Figure 3 for the parameters) and observe the same qualitative properties concerning computation times as for the SIS model in section 4.4. In Figures 8 and 9, we illustrate the realized trajectories for the SSA and the modified τ -leaping algorithm, respectively (left pictures: $N = 2000$, right pictures: $N = 20000$). For simplicity, we only display the proportions of infectious individuals. As in section 4.4, we observe that the trajectories approach the solution of the ODE as $N \rightarrow \infty$.

It is tempting to compare those figures, in order to obtain the error induced by the tau-leaping scheme, compared to the exact SSA scheme. However, this would require that the two schemes are computed using the same realization of the Poisson processes. This is not what we have done here. We intend to do this in a near future, in order to confront simulations with the predictions of the theory concerning the error induced by the tau-leaping algorithm, see e.g. [23], [24].

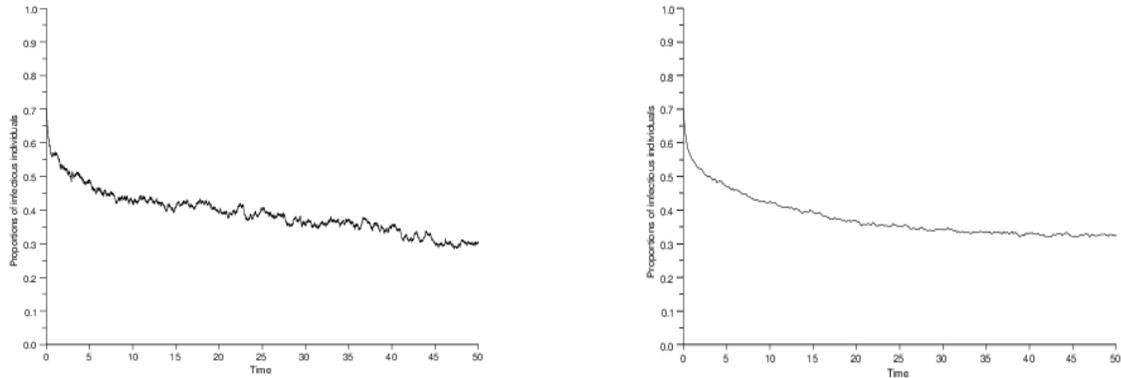


FIGURE 8. Realized trajectories for the proportion of infectious individuals simulated by the SSA (left picture: $N = 2000$, right picture: $N = 20000$). $Z_2(0) = 0.2$, $Z_3(0) = 0.7$, $T = 50$. All other parameters are as in Figure 3.

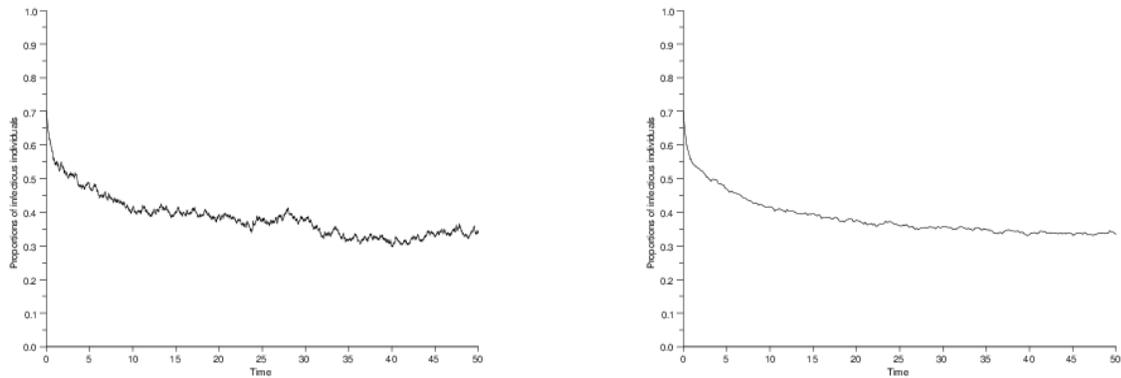


FIGURE 9. Realized trajectories for the proportion of infectious individuals simulated by the τ -leaping algorithm (left picture: $N = 2000$, right picture: $N = 20000$). All parameters are as in Figure 8.

4.6. Computation time

In the next table we present the computation times (in seconds) of the SSA and modified tau-leaping for different population sizes in the SIS model. All programs were written and executed in Matlab. We remark that from $N = 20000$ on, the computation time of the tau-leaping scheme is spectacularly smaller than that of the SSA scheme. We also note that the computation time of the tau-leaping scheme seems to be a decreasing function of N . The reason is probably that the larger N , the less frequently the scheme uses SSA steps.

SSA	Population size N	Time horizon T	Computation time
	2000	50	2.8
	20000	50	24.66
	200000	50	225.44
Modified tau-leaping			
	2000	50	2.3
	20000	50	3.02
	200000	50	1.05

The next table presents the computation times of the SSA and modified tau-leaping for different population sizes in the SIV model. The same comments as for the SIS model apply.

SSA	Population size N	Time horizon T	Computation time
	2000	50	4.51
	20000	50	36.4
	200000	50	352.64
Modified tau-leaping			
	2000	50	2.95
	20000	50	1.04
	200000	50	0.67

4.7. Open questions

We conclude by describing two further open questions concerning the simulation of the process Z^N . We discuss briefly the simulation of models where some jump types are significantly more frequent than others (section 4.7.1). Finally, we discuss the simulation of rare events with the help of large deviations theory (section 4.7.2).

4.7.1. Models with different time scales

The simulation algorithms described in sections 4.1-4.3 tend to become inefficient or inaccurate if some of the jump rates are much larger than others (cf. also Footnote 12). In epidemiology, this is typically the case for so-called host-vector diseases (e.g., malaria), where the disease is transmitted from vectors (mosquitoes) to hosts (humans) and vice versa. In the particular case of malaria, the jumps concerning only the mosquito population have a much shorter time scale than those also concerning the human population (see, e.g., the model in [3]). It is hence time consuming to simulate a significant number of jumps within the human population.

In order to overcome this problem, we simulate only the jumps with small rates via τ -leaping. Within the interval $[t, t + \tau]$ ¹³, the jumps with large rates are approximated by the solution of the ODE (cf. section 2).¹⁴ We intend to analyze the efficiency and accuracy of this procedure. In particular, we are interested in whether the large deviations behavior (cf. section 1.2) is changed by this partial ODE approximation significantly.

4.7.2. Simulation of rare events

We are particularly interested in sampling the probability of the event that the process Z^N exits the domain of attraction of a stable equilibrium before a given time T (cf. section 1.2). If N is large and T is relatively small, the probability of this event is typically extremely small. Applying a naïve Monte Carlo (MC) approach is therefore inefficient and we rely on rare events simulation techniques. In [25], the application of Importance Sampling is explained in the context of diffusion processes. We intend to adapt this method in the following way. Via a change of measure, we transfer the process Z^N to a jump process \tilde{Z} (see, e.g., [6] for a Girsanov theorem for processes of the type of Z^N) for which the event of interest is much more probable. Then, we apply MC simulation to the process \tilde{Z} and associate a weight with each realization (hence the average in the MC is

¹³Here, τ is the time step chosen by the τ -selection procedure applied to only the small rates.

¹⁴The initial value z for this computation can either be chosen explicitly, i.e., $z = Z^N(t)$ or implicitly/at the mid-point of the interval (cf. the somehow related issue in section 4.3.1).

replaced by a weighted average). In order to choose the process \tilde{Z} appropriately (i.e., such that the probability of exit is “maximized”), we use the theory of large deviations: the “cheapest” way to exit the domain is given by the trajectory ϕ^* (cf. sections 1.2 and 3); hence, we change the rates of the process in such a way that the corresponding rate function $I(\phi^*)$ becomes small (ideally, ϕ^* is approximately the solution of the ODE for the modified rates).

Acknowledgements: This research was supported by the ANR project MANEGE, the DAAD, the Labex Archimède and CEMRACS.

We thank Nicolas Champagnat and Tony Lelièvre for useful discussions concerning the τ -leaping algorithm.

REFERENCES

- [1] F. Campillo and C. Lobry. Effect of population size in a predator-prey model. *Ecological Modelling*, 246:1–10, 2012.
- [2] Christopher M. Kribs-Zaleta and Velasco-Hernández. A simple vaccination model with multiple endemic states. *Math. Biosci.*, 164(2):183–201, 2000.
- [3] Nakul Chitnis, J. M. Cushing, and J. M. Hyman. Bifurcation analysis of a mathematical model for malaria transmission. *Siam Journal of Applied mathematics*, 67(1):22–46, 2006.
- [4] Roumen Anguelov, Yves Dumont, Jean Lubuma, and M. Shillor. Dynamically consistent nonstandard finite difference schemes for epidemiological models. *Journal of Computational and Applied Mathematics*, 255:161–182, 2014.
- [5] Thomas G. Kurtz. Strong approximation theorems for density dependent Markov chains. *Stochastic Processes and their Applications*, 6(3):223–240, 1978.
- [6] Adam Shwartz and Alan Weiss. *Large Deviations for Performance Analysis*. Chapman Hall, London, 1995.
- [7] P. Dupuis and R. S. Ellis. *A weak convergence approach to the theory of large deviations*. Wiley series in probability and statistics. Wiley, New York, 1997.
- [8] Jin Feng and Thomas G. Kurtz. *Large deviations for stochastic processes*, volume 131 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, 2006.
- [9] Adam Shwartz and Alan Weiss. Large deviations with diminishing rates. *Mathematics of Operations Research*, 30(2):281–310, 2005.
- [10] Peter Kratz and Etienne Pardoux. Large deviations in epidemiology. *Preprint*, 2015.
- [11] Patrick Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1999.
- [12] Amir Dembo and Ofer Zeitouni. *Large deviations techniques and applications*. Springer-Verlag, New York, 2009.
- [13] M. I. Freidlin and A. D. Wentzell. *Random perturbations of dynamical systems*. Springer, New York, 1998.
- [14] Daniel T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.
- [15] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of chemical physics*, 81(25):2340–2361, 1977.
- [16] Daniel T. Gillespie. Approximate accelerated stochastic simulation of chemically systems. *Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- [17] Ronald E. Mickens, editor. *Applications of nonstandard finite difference schemes*. World Scientific, Singapore, 2000.
- [18] Ronald E. Mickens, editor. *Advances in the applications of nonstandard finite difference schemes*. World Scientific, Singapore, 2005.
- [19] Maurizio Falcone and Tiziana Giorgi. An approximation scheme for evolutive Hamilton–Jacobi equations. In *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, pages 288–303. Birkhäuser, 1994.
- [20] Maurizio Falcone and Roberto Ferretti. *Semi-Lagrangian Approximation Schemes for Linear and Hamilton-Jacobi Equations*. SIAM, Philadelphia, 2013.
- [21] Yang Cao, Daniel T. Gillespie, and Linda R. Petzold. Improved leap-selection for accelerated stochastic simulation. *The journal of chemical physics*, 119(16):8229–8234, 2003.
- [22] Yang Cao, Daniel T. Gillespie, and Linda R. Petzold. Avoiding negative populations in explicit Poisson tau-leaping. *The journal of chemical physics*, 8(2):1–8, 2005.
- [23] Yucheng Hu and Tiejun Li. Highly Accurate Tau-leaping Methods with Random Corrections. *J. Chem. Phys.*, 130(12), 2009.
- [24] David F. Anderson, Arnab Ganguly, and Thomas G. Kurtz. Error analysis of tau-leap simulation methods. *The Annals of Applied Probability*, 21(6):2226–2262, 2011.
- [25] Gilles Wainrib. *Some numerical methods for rare events simulation and analysis*. Lecture Notes in Mathematics series (Biosciences subseries). Springer, 2012.