



Automated importance sampling via optimal control for stochastic reaction networks: A Markovian projection–based approach

Chiheb Ben Hammouda^a, Nadhir Ben Rached^b, Raúl Tempone^{c,d}, Sophia Wiechert^{e,*}

^a Utrecht University, Mathematical Institute, 3584 CD Utrecht, The Netherlands

^b University of Leeds, School of Mathematics, Woodhouse, Leeds LS2 9JT, UK

^c King Abdullah University of Science and Technology (KAUST), Computer, Electrical and Mathematical Sciences & Engineering Division (CEMSE), Thuwal 23955-6900, Saudi Arabia

^d RWTH Aachen University, Alexander von Humboldt Professor in Mathematics for Uncertainty Quantification, 52062 Aachen, Germany

^e RWTH Aachen University, Chair of Mathematics for Uncertainty Quantification, Pontdriesch 14-16, 52062 Aachen, Germany

ARTICLE INFO

Keywords:

Stochastic reaction networks

Tau-leap

Importance sampling

Stochastic optimal control

Markovian projection

Rare event

ABSTRACT

We propose a novel alternative approach to our previous work (Ben Hammouda et al., 2023) to improve the efficiency of Monte Carlo (MC) estimators for rare event probabilities for stochastic reaction networks (SRNs). In the same spirit of Ben Hammouda et al. (2023), an efficient path-dependent measure change is derived based on a connection between determining optimal importance sampling (IS) parameters within a class of probability measures and a stochastic optimal control formulation, corresponding to solving a variance minimization problem. In this work, we propose a novel approach to address the encountered curse of dimensionality by mapping the problem to a significantly lower-dimensional space via a Markovian projection (MP) idea. The output of this model reduction technique is a low-dimensional SRN (potentially even one dimensional) that preserves the marginal distribution of the original high-dimensional SRN system. The dynamics of the projected process are obtained by solving a related optimization problem via a discrete L^2 regression. By solving the resulting projected Hamilton–Jacobi–Bellman (HJB) equations for the reduced-dimensional SRN, we obtain projected IS parameters, which are then mapped back to the original full-dimensional SRN system, resulting in an efficient IS-MC estimator for rare events probabilities of the full-dimensional SRN. Our analysis and numerical experiments reveal that the proposed MP-HJB-IS approach substantially reduces the MC estimator variance, resulting in a lower computational complexity in the rare event regime than standard MC estimators.

1. Introduction

This paper proposes an efficient estimator for rare event probabilities for a particular class of continuous-time Markov processes, stochastic reaction networks (SRNs). We design an automated importance sampling (IS) approach based on the approximate explicit tau-leap (TL) scheme to build a Monte Carlo (MC) estimator for rare event probabilities of SRNs. The used IS change of measure was introduced in [1], wherein the optimal IS controls were determined via a stochastic optimal control (SOC) formulation. In that same work, we also presented a learning-based approach to avoid the curse of dimensionality. Building on that work, we propose an alternative method for high-dimensional SRNs that leverages dimension reduction through Markovian projection (MP) and then

* Corresponding author.

E-mail address: wiechert@uq.rwth-aachen.de (S. Wiechert).

<https://doi.org/10.1016/j.cam.2024.115853>

Received 7 July 2023; Received in revised form 12 December 2023

Available online 23 February 2024

0377-0427/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

recovers the optimal IS controls of the full-dimensional SRNs as a mapping from the solution in lower-dimensional space, potentially one. To the best of our knowledge, we are the first to establish the MP framework for the SRN setting to solve an IS problem.

An SRN (refer to Section 1.1 for a brief introduction and [2] for more details) describes the time evolution of a set of species through reactions and can be found in a wide range of applications, such as biochemical reactions, epidemic processes [3,4], and transcription and translation in genomics and virus kinetics [5,6]. For a d -dimensional SRN, $\mathbf{X} : [0, T] \rightarrow \mathbb{N}^d$, with the given final time $T > 0$, we aim to determine accurate and computationally efficient MC estimations for the expected value $\mathbb{E}[g(\mathbf{X}(T))]$. The observable $g : \mathbb{N}^d \rightarrow \mathbb{R}$ is a given scalar function of \mathbf{X} , where indicator functions $g(\mathbf{x}) = \mathbb{1}_{\{\mathbf{x} \in B\}}$ are of interest to estimate the rare event probability $\mathbb{P}(\mathbf{X}(T) \in B) \ll 1$, where $B \subset \mathbb{R}^d$.

The quantity of interest, $\mathbb{E}[g(\mathbf{X}(T))]$, is the solution to the corresponding Kolmogorov backward equations [7]. Solving these ordinary differential equations (ODEs) in closed form is infeasible for most SRNs; thus, numerical approximations based on discretized schemes are used to derive solutions. A drawback of these approaches is that, without using dimension reduction techniques, the computational cost scales exponentially with the number of species d . To avoid the curse of dimensionality, we propose estimating $\mathbb{E}[g(\mathbf{X}(T))]$ using MC methods.

Numerous schemes have been developed to simulate the exact sample paths of SRNs. These include the stochastic simulation algorithm introduced by Gillespie in [8] and the modified next reaction method proposed by Anderson in [9]. However, when SRNs involve reaction channels with high reaction rates, simulating exact realizations of the system can be computationally expensive. To address this issue, Gillespie [10] and Aparicio and Solari [11] independently proposed the explicit-TL method (see Section 1.2), which approximates the paths of \mathbf{X} by evolving the process with fixed time steps while maintaining constant reaction rates within each time step. Additionally, other simulation schemes have been proposed to handle situations with well-separated fast and slow time scales [12–17].

In order to compute MC estimates of $\mathbb{E}[g(\mathbf{X}(T))]$ more efficiently, different variance reduction techniques have been proposed in the context of SRNs. In the spirit of the multilevel MC (MLMC) idea [18,19], various MLMC-based methods [16,17,20–22] have been introduced to overcome different challenges in this context. Moreover, as the naive MC and MLMC estimators have high computational costs when used for estimating rare event probabilities, different IS approaches [1,23–29] have been proposed.

To estimate various statistical quantities efficiently for SRNs (specifically rare event probabilities), we use the path-dependent IS approach originally introduced in [1]. This class of probability measure change is based on modifying the rates of the Poisson random variables used to construct the TL paths. In [1], it is shown how optimal IS controls are obtained by minimizing the second moment of the IS estimator (equivalently, the variance), representing the cost function of the associated SOC problem, and that the corresponding value function solves a dynamic programming relation (see Section 2.1 for revising these results). In this work, we generalize the discrete-time dynamic programming relation by a set of continuous-time ODEs, the Hamilton–Jacobi–Bellman (HJB) equations, allowing the formulation of optimal IS controls in continuous time. Compared to the discrete-time IS control formulation presented in [1], the continuous-time formulation offers the advantage that it provides a curve of IS controls over time instead of a discrete set. This allows its application for any time stepping in the IS-TL paths and thereby eliminates the need for ad-hoc interpolations often needed in the discrete setting.

In the multidimensional setting, the cost of solving the backward HJB equations increases exponentially with respect to the dimension d (curse of dimensionality). In [1], we proposed a learning-based approach to reduce this effect. In that approach, the value function is approximated using an ansatz function, the parameters of which are learned through a stochastic optimization algorithm (see Fig. 1.1 for a schematic illustration of the approach). In this work, we present an alternative method using a dimension reduction approach for SRNs (see Fig. 1.2 for a schematic illustration of the approach). The proposed methodology is to adapt the MP idea originally introduced in [30] for the setting of diffusion-type stochastic differential equations (SDEs) to the SRN framework, resulting in a significantly lower-dimensional process, preserving the marginal distribution of the original full-dimensional SRN. The propensities characterizing the lower-dimensional MP process can be approximated using L^2 regression. Using the resulting low-dimensional SRN, we derive an approximate value function and, consequently, near-optimal IS controls while reducing the effect of the curse of dimensionality. By mapping the IS controls to the original full-dimensional SRNs, we derive an unbiased IS-MC estimator for the TL scheme. Compared to the learning-based approach presented in [1], this novel MP-IS approach eliminates the need for an ansatz function to model the value function. This approach allows its application to general observables g that differ from indicator functions for rare event estimation, because no prior knowledge regarding the shape of the value function and suitable ansatz functions is required.

To the best of our knowledge, we are the first to establish the MP idea for SRNs and apply it to derive an efficient pathwise IS for MC methods. Initially, the MP idea was introduced for Itô stochastic processes in [30,31] and was later generalized to martingales and semimartingales [32,33]. In addition, MP has been widely applied for dimension reduction in SDEs [34], particularly in financial applications [35,36]. For instance, in [37], solving HJB equations for an MP process was pursued but in the setting of Itô SDEs with the application of pricing American options. In [38], MP was used for control problems and IS problems for rare events in high-dimensional diffusion processes with multiple time scales. In this work, we introduce the general dimension reduction framework of MP for SRNs such that it can be applied to other problems beyond the selected IS application. (e.g., solving the chemical master equation [39] or the Kolmogorov backward equations [40]).

The remainder of this work is organized as follows. Sections 1.1, 1.2, 1.3, and 1.4 recall the relevant SRN, TL, MC, and IS notations and definitions from [1]. Next, Section 2 reviews the connection between IS and SOC by introducing the IS scheme, the value function, and the corresponding dynamic programming theorem from [1] in Section 2.1. Then, Section 2.2 extends the framework to a continuous-time formulation leading to the continuous-time value function and deriving the corresponding HJB equations. Section 3 presents the MP technique for SRNs and shows how the projected dynamics can be computed using L^2 regression. Next,

Section 4 addresses the curse of dimensionality of high-dimensional SRNs occurring from the optimal IS scheme in Section 2 by combining the IS scheme with MP (Section 3) to derive near-optimal IS controls. Finally, Section 5 presents the numerical results for the rare event probability estimation to demonstrate the efficiency of the proposed MP-IS approach compared to a standard TL-MC estimator.

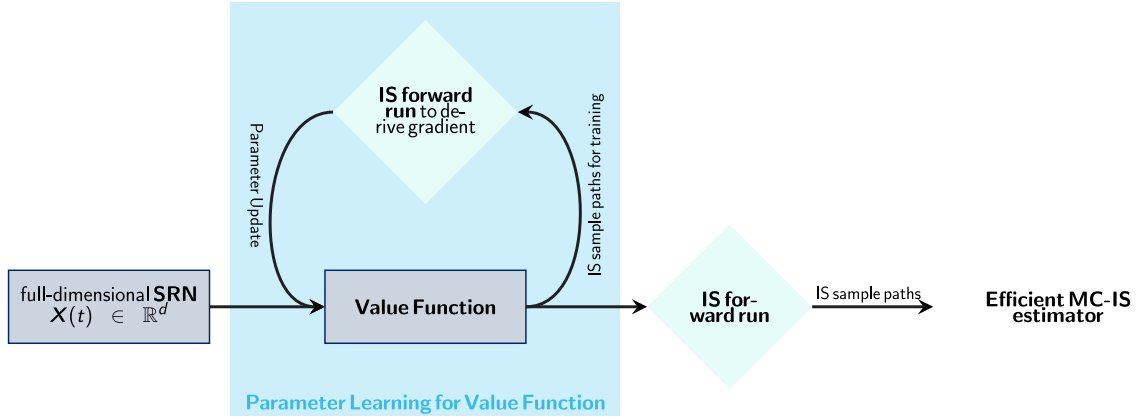


Fig. 1.1. Schematic diagram of the learning-based approach in [1].




Fig. 1.2. Schematic diagram of the MP-IS approach.

1.1. Stochastic reaction networks (SRNs)

We recall from [1] that an SRN describes the time evolution for a homogeneously mixed chemical reaction system, in which d distinct species interact through J reaction channels. Each reaction channel $\mathcal{R}_j, j = 1 \dots, J$, is given by the relation

$$\alpha_{j,1}S_1 + \dots + \alpha_{j,d}S_d \xrightarrow{\theta_j} \beta_{j,1}S_1 + \dots + \beta_{j,d}S_d, \tag{1.1}$$

where $\alpha_{j,i}$ molecules of species S_i are consumed and $\beta_{j,i}$ molecules are produced. The positive constants $\{\theta_j\}_{j=1}^J$ represent the reaction rates.

This process can be modeled by a Markovian pure jump process, $\mathbf{X} : [0, T] \times \Omega \rightarrow \mathbb{N}^d$, where $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space. We are interested in the time evolution of the state vector,

$$\mathbf{X}(t) = (X_1(t), \dots, X_d(t)) \in \mathbb{N}^d, \tag{1.2}$$

where the i th component, $X_i(t)$, describes the abundance of the i th species present in the system at time t . The process \mathbf{X} is a continuous-time, discrete-space Markov process characterized by Kurtz’s random time change representation [41]:

$$\mathbf{X}(t) = \mathbf{x}_0 + \sum_{j=1}^J Y_j \left(\int_0^t a_j(\mathbf{X}(s)) ds \right) \mathbf{v}_j, \tag{1.3}$$

where $Y_j : \mathbb{R}_+ \times \Omega \rightarrow \mathbb{N}$ are independent unit-rate Poisson processes and the stoichiometric vector is defined as $\mathbf{v}_j = (\beta_{j,1} - \alpha_{j,1}, \dots, \beta_{j,d} - \alpha_{j,d}) \in \mathbb{Z}^d$.

For each reaction channel \mathcal{R}_j , the propensity function $a_j : \mathbb{N}^d \rightarrow \mathbb{R}_+$ obeys the *non-negativity assumption* (i.e. the system cannot attain negative states), which is that $a_j(\mathbf{X}) = 0$ for \mathbf{x} such that $\mathbf{X} + \mathbf{v}_j \notin \mathbb{N}^d$. In our numerical simulations, we consider a propensity derived from the *stochastic mass-action kinetic* principle

$$a_j(\mathbf{x}) := \theta_j \prod_{i=1}^d \frac{x_i!}{(x_i - \alpha_{j,i})!} \mathbf{1}_{\{x_i \geq \alpha_{j,i}\}}, \tag{1.4}$$

where x_i is the counting number for species S_i . However, the approach presented in this work is not restricted to the particular structure of the propensity function in (1.4) (see Remark 3.4).

1.2. Explicit Tau-Leap approximation

The explicit-TL scheme is a pathwise approximate method based on Kurtz’s random time change representation (1.3) [10,11]. It was originally introduced to overcome the computational drawbacks of exact methods, which become computationally expensive when many reactions fire during a short time interval. For a uniform time mesh $\{t_0 = 0, t_1, \dots, t_N = T\}$ with step size $\Delta t = \frac{T}{N}$ and a given initial value $\mathbf{X}(0) = \mathbf{x}_0$, the explicit-TL approximation for \mathbf{X} is defined by

$$\begin{aligned} \hat{\mathbf{X}}_0 &:= \mathbf{x}_0 \\ \hat{\mathbf{X}}_k^{\Delta t} &:= \max \left(\mathbf{0}, \hat{\mathbf{X}}_{k-1}^{\Delta t} + \sum_{j=1}^J \mathcal{P}_{k-1,j} \left(a_j(\hat{\mathbf{X}}_{k-1}^{\Delta t}) \Delta t \right) \mathbf{v}_j \right), \quad 1 \leq k \leq N, \end{aligned} \tag{1.5}$$

where $\{\mathcal{P}_{k,j}(r_{k,j})\}_{1 \leq j \leq J}$ are independent Poisson random variables with respective rates $r_{k,j} := a_j(\hat{\mathbf{X}}_k^{\Delta t}) \Delta t$ conditioned on the current state $\hat{\mathbf{X}}_k^{\Delta t}$. The maximum in (1.5) is applied entry-wise. In each TL step, the current state is projected to zero to prevent the process from exiting the lattice (i.e., producing negative values).

1.3. Biased Monte Carlo estimator

We let \mathbf{X} be an SRN and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar observable. For a given final time T , we estimate $\mathbb{E}[g(\mathbf{X}(T))]$ using the standard MC-TL estimator:

$$\mu_M := \frac{1}{M} \sum_{m=1}^M g(\hat{\mathbf{X}}_{[m]}^{\Delta t}(T)), \tag{1.6}$$

where $\{\hat{\mathbf{X}}_{[m]}^{\Delta t}(T)\}_{m=1}^M$ are independent TL samples.

The global error for the proposed MC estimator has the following error decomposition:

$$|\mathbb{E}[g(\mathbf{X}(T))] - \mu_M| \leq \underbrace{|\mathbb{E}[g(\mathbf{X}(T))] - \mathbb{E}[g(\hat{\mathbf{X}}^{\Delta t}(T))]|}_{\text{Bias}} + \underbrace{|\mathbb{E}[g(\hat{\mathbf{X}}^{\Delta t}(T))] - \mu_M|}_{\text{Statistical Error}}. \tag{1.7}$$

Under some assumptions, the TL scheme has a weak order, $\mathcal{O}(\Delta t)$ [42], that is, for sufficiently small Δt ,

$$\left| \mathbb{E} \left[g(\mathbf{X}(T)) - g(\hat{\mathbf{X}}^{\Delta t}(T)) \right] \right| \leq C \Delta t \tag{1.8}$$

where $C > 0$.

The bias and statistical error can be bound equally using $\frac{TOL}{2}$ to achieve the desired accuracy, TOL, with a confidence level of $1 - \alpha$ for $\alpha \in (0, 1)$, which can be achieved by the step size:

$$\Delta t(TOL) = \frac{TOL}{2 \cdot C} \quad \text{and} \quad M^*(TOL) = C_\alpha^2 \frac{4 \cdot \text{Var}[g(\hat{\mathbf{X}}^{\Delta t}(T))]}{TOL^2} \tag{1.9}$$

sample paths, where the constant C_α is the $(1 - \frac{\alpha}{2})$ -quantile for the standard normal distribution. We select $C_\alpha = 1.96$ for a 95% confidence level corresponding to $\alpha = 0.05$.

When estimating rare event probabilities, we are interested in the relative error

$$\frac{|\mathbb{E}[g(\mathbf{X}(T))] - \mu_M|}{|\mathbb{E}[g(\mathbf{X}(T))]|}.$$

In this context, to achieve a prescribed relative tolerance TOL_{rel} , we use step size

$$\Delta t_{rel}(TOL_{rel}) = \frac{TOL_{rel} |\mathbb{E}[g(\mathbf{X}(T))]|}{2 \cdot C} \tag{1.10}$$

and

$$M^*_{rel}(TOL_{rel}) = C_\alpha^2 \frac{4 \cdot \text{Var}[g(\hat{\mathbf{X}}^{\Delta t}(T))]}{TOL_{rel}^2 |\mathbb{E}[g(\mathbf{X}(T))]|^2} \tag{1.11}$$

sample paths.

Given that the computational cost to simulate a single path is $\mathcal{O}(\Delta t^{-1})$, the expected total computational complexity is $\mathcal{O}(TOL^{-3})$ and $\mathcal{O}(TOL_{rel}^{-3})$ for the absolute and relative errors, respectively.

1.4. Importance sampling

Using IS techniques [23,24,26–28,28,29] can improve the computational costs for the crude MC estimator through variance reduction in (1.9). For a general motivation, we refer to [1] Section 1.4. For illustrating the IS method, let us consider the general problem of estimating $\mathbb{E}[g(Y)]$, where g is a given observable and Y is a random variable taking values in \mathbb{R} with the probability density function ρ_Y . We let $\hat{\rho}_Z$ be the probability density function for an auxiliary real random variable Z . The MC estimator under the IS measure is

$$\mu_M^{IS} = \frac{1}{M} \sum_{j=1}^M L(Z_{[j]}) \cdot g(Z_{[j]}), \tag{1.12}$$

where $Z_{[j]}$ are independent and identically distributed samples from $\hat{\rho}_Z$ for $j = 1, \dots, M$ and the likelihood factor is given by $L(Z_{[j]}) := \frac{\rho_Y(Z_{[j]})}{\hat{\rho}_Z(Z_{[j]})}$. The IS estimator retains the expected value of (1.6) (i.e., $\mathbb{E}[L(Z)g(Z)] = \mathbb{E}[g(Y)]$), but the variance can be reduced due to a different second moment $\mathbb{E}[(L(Z) \cdot g(Z))^2]$.

Determining an auxiliary probability measure that substantially reduces the variance compared with the original measure is challenging and strongly depends on the structure of the considered problem. In addition, the derivation of the new measure must come with a moderate additional computational cost to ensure an efficient IS scheme. This work uses the path-dependent change of probability measure introduced in [1], employing an IS measure derived from changing the Poisson random variable rates in the TL paths. Section 2.1 recalls the SOC formulation for optimal IS parameters from [1] and extends it with a novel HJB formulation. We conclude this consideration in Section 4, combining the IS scheme with a dimension reduction approach to reduce the computational cost.

2. Importance sampling via stochastic optimal control formulation

2.1. Dynamic programming for importance sampling parameters

This section revisits the connection between optimal IS measure determination within a class of probability measures, and the SOC formulated originally in [1]. We let \mathbf{X} be an SRN as defined in Section 1.1 and let $\hat{\mathbf{X}}^{Dr}$ denote its TL approximation as given by (1.5). Then, the goal is to derive a near-optimal IS measure to estimate $\mathbb{E}[g(\mathbf{X}(T))]$. We limit ourselves to the parameterized class of IS schemes used in [1,22]:

$$\begin{aligned} \bar{\mathbf{X}}_{n+1}^{Dr} &= \max \left(\mathbf{0}, \bar{\mathbf{X}}_n^{Dr} + \sum_{j=1}^J \bar{P}_{n,j} \mathbf{v}_j \right), \quad n = 0, \dots, N - 1, \\ \bar{\mathbf{X}}_0^{Dr} &= \mathbf{x}_0, \end{aligned} \tag{2.1}$$

where the measure change is obtained by modifying the Poisson random variable rates of the TL paths:

$$\bar{P}_{n,j} = \bar{P}_{n,j} \left(\delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr}) \Delta t \right), \quad n = 0, \dots, N - 1, j = 1, \dots, J. \tag{2.2}$$

In (2.2), $\delta_{n,j}^{Dr}(\mathbf{x}) \in \mathcal{A}_{\mathbf{x},j}$ is the control parameter at time step n , under reaction j , and in state $\mathbf{x} \in \mathbb{N}^d$. In addition, $\{\bar{P}_{n,j}(r_{n,j})\}_{\{1 \leq j \leq J\}}$ are independent Poisson random variables, conditioned on $\bar{\mathbf{X}}_n^{Dr}$, with the respective rates $r_{n,j} := \delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr}) \Delta t$. The set of admissible controls is

$$\mathcal{A}_{\mathbf{x},j} = \begin{cases} \{0\} & , \text{if } a_j(\mathbf{x}) = 0 \\ \{y \in \mathbb{R} : y > 0\} & , \text{otherwise.} \end{cases} \tag{2.3}$$

In the following, we use the vector notation $(\delta_n^{Dr}(\mathbf{x}))_j := \delta_{n,j}^{Dr}(\mathbf{x})$ and $(\bar{\mathbf{P}}_n)_j := \bar{P}_{n,j}$ for $j = 1, \dots, J$.

The corresponding likelihood ratio of the path $\{\bar{\mathbf{X}}_n^{Dr} : n = 0, \dots, N\}$ for the IS parameters $\delta_n^{Dr}(\mathbf{x}) \in \times_{j=1}^J \mathcal{A}_{\mathbf{x},j}$ is

$$L \left((\bar{\mathbf{P}}_0, \dots, \bar{\mathbf{P}}_{N-1}), (\delta_0^{Dr}(\bar{\mathbf{X}}_0^{Dr}), \dots, \delta_{N-1}^{Dr}(\bar{\mathbf{X}}_{N-1}^{Dr})) \right) = \prod_{n=0}^{N-1} L_n(\bar{\mathbf{P}}_n, \delta_n^{Dr}(\bar{\mathbf{X}}_n^{Dr})), \tag{2.4}$$

where the likelihood ratio update at time step n is

$$L_n(\bar{\mathbf{P}}_n, \delta_n^{Dr}(\bar{\mathbf{X}}_n^{Dr})) = \prod_{j=1}^J \exp \left(-(a_j(\bar{\mathbf{X}}_n^{Dr}) - \delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr})) \Delta t \right) \left(\frac{a_j(\bar{\mathbf{X}}_n^{Dr})}{\delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr})} \right)^{\bar{P}_{n,j}} = \exp \left(- \left(\sum_{j=1}^J a_j(\bar{\mathbf{X}}_n^{Dr}) - \delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr}) \right) \Delta t \right) \cdot \prod_{j=1}^J \left(\frac{a_j(\bar{\mathbf{X}}_n^{Dr})}{\delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr})} \right)^{\bar{P}_{n,j}}. \tag{2.5}$$

To simplify the notation, we use the convention that $\frac{a_j(\bar{\mathbf{X}}_n^{Dr})}{\delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr})} = 1$, whenever both $a_j(\bar{\mathbf{X}}_n^{Dr}) = 0$ and $\delta_{n,j}^{Dr}(\bar{\mathbf{X}}_n^{Dr}) = 0$ in (2.5). From (2.3), this results in a factor of 1 in the likelihood ratio for reactions where $a_j(\bar{\mathbf{X}}_n^{Dr}) = 0$.

Using the introduced change of measure (2.2), the quantity of interest can be expressed with respect to the new measure:

$$\mathbb{E}[g(\widehat{\mathbf{X}}_N^{\Delta t})] = \mathbb{E} \left[L \left(\left(\overline{\mathbf{P}}_0, \dots, \overline{\mathbf{P}}_{N-1} \right), \left(\delta_0^{\Delta t}(\overline{\mathbf{X}}_0), \dots, \delta_{N-1}^{\Delta t}(\overline{\mathbf{X}}_{N-1}) \right) \right) \cdot g(\overline{\mathbf{X}}_N^{\Delta t}) \right], \tag{2.6}$$

with the expectation on the right-hand side of (2.6) taken with respect to the dynamics in (2.1).

Next, we recall the connection between the optimal second moment minimizing IS parameters $\{\delta_n^{\Delta t}(\mathbf{x})\}_{n=0, \dots, N-1; \mathbf{x} \in \mathbb{N}^d}$ and the corresponding discrete-time dynamic programming relation from [1]. We revisit the definition of the discrete-time value function $u_{\Delta t}(\cdot, \cdot)$ in Definition 2.1, allowing the formulation of the dynamic programming equations in Theorem 2.2. The proof and further details for Theorem 2.2 are provided in [1].

Definition 2.1 (Value Function). For a given $\Delta t > 0$, the discrete-time value function $u_{\Delta t}(\cdot, \cdot)$ is defined as the optimal (infimum) second moment for the proposed IS estimator. For time step $0 \leq n \leq N$ and state $\mathbf{x} \in \mathbb{N}^d$,

$$u_{\Delta t}(n, \mathbf{x}) = \inf_{\{\delta_k^{\Delta t}\}_{k=n, \dots, N-1} \in \mathcal{A}^{N-n}} \mathbb{E} \left[g^2(\overline{\mathbf{X}}_N^{\Delta t}) \prod_{k=n}^{N-1} L_k^2(\overline{\mathbf{P}}_k, \delta_k^{\Delta t}(\overline{\mathbf{X}}_k)) \mid \overline{\mathbf{X}}_n^{\Delta t} = \mathbf{x} \right], \tag{2.7}$$

where $\mathcal{A} = \times_{\mathbf{x} \in \mathbb{N}^d} \times_{j=1}^J \mathcal{A}_{\mathbf{x},j} \in \mathbb{R}^{\mathbb{N}^d \times J}$ is the admissible set for the IS parameters, and $u_{\Delta t}(N, \mathbf{x}) = g^2(\mathbf{x})$, for any $\mathbf{x} \in \mathbb{N}^d$.

Theorem 2.2 (Dynamic Programming for IS Parameters [1]). For $\mathbf{x} \in \mathbb{N}^d$ and the given step size $\Delta t > 0$, the discrete-time value function $u_{\Delta t}(n, \mathbf{x})$ fulfills the dynamic programming relation:

$$u_{\Delta t}(N, \mathbf{x}) = g^2(\mathbf{x})$$

and for $n = N - 1, \dots, 0$, and $\mathcal{A}_{\mathbf{x}} := \times_{j=1}^J \mathcal{A}_{\mathbf{x},j}$,

$$u_{\Delta t}(n, \mathbf{x}) = \inf_{\delta_n^{\Delta t}(\mathbf{x}) \in \mathcal{A}_{\mathbf{x}}} \exp \left(\left(-2 \sum_{j=1}^J a_j(\mathbf{x}) + \sum_{j=1}^J \delta_{n,j}^{\Delta t}(\mathbf{x}) \right) \Delta t \right) \times \sum_{\mathbf{p} \in \mathbb{N}^J} \left(\prod_{j=1}^J \frac{(\Delta t \cdot \delta_{n,j}^{\Delta t}(\mathbf{x}))^{p_j}}{p_j!} \left(\frac{a_j(\mathbf{x})}{\delta_{n,j}^{\Delta t}(\mathbf{x})} \right)^{2p_j} \right) \cdot u_{\Delta t}(n+1, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}\mathbf{p})), \tag{2.8}$$

where $\mathbf{v} = (v_1, \dots, v_J) \in \mathbb{Z}^{d \times J}$.

Analytically solving the minimization problem (2.8) is challenging due to the infinite sum. In [1], the problem is solved by approximating the value function (2.7) using a truncated Taylor expansion of the dynamic programming (2.8). To overcome the curse of dimensionality, a learning-based approach for the value function was proposed. Instead, in this work, we utilize a continuous-time SOC formulation, leading to a set of coupled d -dimensional ODEs, the HJB equations (refer to Section 2.2). We deal with the curse of dimensionality issue by using a dimension reduction technique, namely the MP, as explained in Section 3.

2.2. Derivation of Hamilton–Jacobi–Bellman (HJB) Equations

In Corollary 2.3, the discrete-time dynamic programming relation in Theorem 2.2 is replaced by its analogous continuous-time relation, resulting in a set of ODEs known as the HJB equations. The continuous-time value function $\tilde{u}(\cdot, \mathbf{x}) : [0, T] \rightarrow \mathbb{R}$, $\mathbf{x} \in \mathbb{N}^d$, is the limit of the discrete value function $u_{\Delta t}(\cdot, \mathbf{x})$ as the step size Δt approaches zero. In addition, the IS controls $\delta(\cdot, \mathbf{x}) : [0, T] \rightarrow \mathcal{A}_{\mathbf{x}}$ become time-continuous curves for $\mathbf{x} \in \mathbb{N}^d$.

Corollary 2.3 (HJB Equations for IS Parameters). For all $\mathbf{x} \in \mathbb{N}^d$, the continuous-time value function $\tilde{u}(t, \mathbf{x})$ fulfills (2.9) for $t \in [0, T]$:

$$\tilde{u}(T, \mathbf{x}) = g^2(\mathbf{x})$$

$$-\frac{d\tilde{u}}{dt}(t, \mathbf{x}) = \inf_{\delta(t, \mathbf{x}) \in \mathcal{A}_{\mathbf{x}}} \left(-2 \sum_{j=1}^J a_j(\mathbf{x}) + \sum_{j=1}^J \delta_j(t, \mathbf{x}) \right) \tilde{u}(t, \mathbf{x}) + \sum_{j=1}^J \frac{a_j(\mathbf{x})^2}{\delta_j(t, \mathbf{x})} \tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}_j)), \tag{2.9}$$

where $\delta_j(t, \mathbf{x}) := (\delta(t, \mathbf{x}))_j$.

Proof. The proof of the corollary is presented in Appendix A. \square

If $\tilde{u}(t, \mathbf{x}) > 0$ for all $\mathbf{x} \in \mathbb{N}^d$ and $t \in [0, T]$, we can solve the minimization problem in (2.9) in closed form, such that the optimal controls are given by

$$\tilde{\delta}_j(t, \mathbf{x}) = a_j(\mathbf{x}) \sqrt{\frac{\tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}_j))}{\tilde{u}(t, \mathbf{x})}} \tag{2.10}$$

and (2.9) simplifies to

$$\frac{d\tilde{u}}{dt}(t, \mathbf{x}) = -2 \sum_{j=1}^J a_j(\mathbf{x}) \left(\sqrt{\tilde{u}(t, \mathbf{x}) \tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}_j))} - \tilde{u}(t, \mathbf{x}) \right). \tag{2.11}$$

To estimate rare event probabilities with an observable $g(\mathbf{x}) = \mathbb{1}_{\{x_i > \gamma\}}$, we encounter $\bar{u}(t, \mathbf{x}) = 0$ for some $\mathbf{x} \in \mathbb{N}^d$; therefore, we modify (2.9) by approximating the final condition $g(\mathbf{x})$ using a sigmoid:

$$\bar{g}(\mathbf{x}) = \frac{1}{1 + \exp(-b - \beta x_i)} > 0 \tag{2.12}$$

with appropriately chosen parameters $b \in \mathbb{R}$ and $\beta \in \mathbb{R}$. By incorporating the modified final condition, we obtain an approximate value function by solving (2.11) using an ODE solver (e.g. ode23s from MATLAB). When using the numerical solver, we truncate the infinite state space \mathbb{N}^d using sufficiently large upper bounds. The approximated near-optimal IS controls are then expressed by (2.10). By the truncation of the infinite state space and the approximation of the final condition g by \bar{g} , we introduce a bias to the value function. This can impact the amount of variance reduction in the IS-MC forward run, however, the IS-MC estimator is bias-free.

The cost for the ODE solver scales exponentially with the dimension d of the SRNs, making this approach infeasible for high-dimensional SRNs. Section 3 presents a dimension reduction approach for SRNs employed in Section 4 to derive suboptimal IS controls for a lower-dimensional SRN. We later demonstrate how these controls are mapped to the full-dimensional SRN system.

Remark 2.4 (Continuous-time IS Controls). In Corollary 2.3 and Theorem 2.2, we present two alternative methods to express the value function (2.7) and the IS controls. Utilizing the HJB framework, we can derive continuous controls across time. This allows any time stepping Δt in the IS-TL forward run and eliminates the need for ad-hoc interpolations.

3. Markovian projection for stochastic reaction networks

3.1. Formulation

To address the curse of dimensionality problem when deriving near-optimal IS controls, we project the SRN to a lower-dimensional network while preserving the marginal distribution of the original high-dimensional SRN system. We adapt the MP idea originally introduced in [30] for the setting of diffusion type stochastic differential equations to the SRNs framework. For an d -dimensional SRN state vector, $\mathbf{X}(t)$, we introduce a projection to a \bar{d} -dimensional space such that $1 \leq \bar{d} \ll d$:

$$P : \mathbb{R}^d \rightarrow \mathbb{R}^{\bar{d}} : \mathbf{x} \mapsto \mathbf{P}\mathbf{x},$$

where $\mathbf{P} \in \mathbb{R}^{\bar{d} \times d}$ is a given matrix. This section develops a general MP framework for arbitrary projections with $\bar{d} \geq 1$. However, the choice of the projection depends on the quantity of interest. In particular, when considering rare event probabilities with an observable $g(\mathbf{x}) = \mathbb{1}_{\{x_i > \gamma\}}, \gamma \in \mathbb{R}$ as we do in Section 4, the projection operator is of the form

$$P(\mathbf{x}) = \left(0, \dots, 0, \underset{i-1}{1}, \underset{i}{0}, \dots, \underset{i+1}{0}, \dots, 0 \right) \mathbf{x}. \tag{3.1}$$

In the derivation of the MP, we assume that

$$\mathbb{E} [a_j(\mathbf{X}(t)) \mid P(\mathbf{X}(t)) = \mathbf{s}, \mathbf{X}(0) = \mathbf{x}_0] < \infty \tag{3.2}$$

for $\mathbf{s} \in \mathbb{N}^{\bar{d}}, t \in [0, T]$ and $1 \leq j \leq J$.

Theorem 3.1 shows that a \bar{d} dimensional SRN, $\bar{\mathcal{S}}(t)$ exists that follows the same conditional distribution as $\mathcal{S}(t) := P(\mathbf{X}(t))$ conditioned on the initial state $\mathbf{X}(0) = \mathbf{x}_0$ for all $t \in [0, T]$.

Theorem 3.1 (MP for SRNs). Under the assumption in Eq. (3.2), we let $\bar{\mathcal{S}}(t)$ be a \bar{d} -dimensional stochastic process whose dynamics are given by

$$\bar{\mathcal{S}}(t) = P(\mathbf{x}_0) + \sum_{j=1}^J \bar{Y}_j \left(\int_0^t \bar{a}_j(\tau, \bar{\mathcal{S}}(\tau)) d\tau \right) \underbrace{P(\mathbf{v}_j)}_{=: \bar{\mathbf{v}}_j}, \tag{3.3}$$

for $t \in [0, T]$, where \bar{Y}_j denotes independent unit-rate Poisson processes and $\bar{a}_j, j = 1, \dots, J$, are characterized by

$$\bar{a}_j(t, \mathbf{s}) := \mathbb{E} [a_j(\mathbf{X}(t)) \mid P(\mathbf{X}(t)) = \mathbf{s}, \mathbf{X}(0) = \mathbf{x}_0], \text{ for } 1 \leq j \leq J, \mathbf{s} \in \mathbb{N}^{\bar{d}}. \tag{3.4}$$

Thus, $\mathcal{S}(t) \mid_{\{\mathbf{X}(0)=\mathbf{x}_0\}} = P(\mathbf{X}(t)) \mid_{\{\mathbf{X}(0)=\mathbf{x}_0\}}$ and $\bar{\mathcal{S}}(t) \mid_{\{\mathbf{X}(0)=\mathbf{x}_0\}}$ have the same distribution for all $t \in [0, T]$.

Proof. The proof for Theorem 3.1 is given in Appendix B. \square

In Theorem Theorem 3.1, we require the assumption in (3.2) to hold in order to ensure that the MP propensity in (3.4) is well-defined. Assumption (3.2) does not hold for all SRNs. However, in Remark 3.1, we present sufficient conditions guaranteeing (3.2).

Remark 3.2 (Sufficient conditions for assumption (3.2)). If the propensity functions $\{a_j(\cdot)\}_{j=1}^J$ are bounded, i.e., it exist bounds $K_j \in \mathbb{R}^+$ such that for any $\mathbf{x} \in \mathbb{N}^d$ $a_j(\mathbf{x}) < K_j$ for $1 \leq j \leq J$. Then, the expectation in (3.2) is bounded by the same bounds.

An alternative condition is that the set

$$\mathcal{K}_{t,s} := \{y \in \mathbb{N}^d : P(y) = s \text{ and } \mathbb{P}(\mathbf{X}(t) = y) > 0\} \tag{3.5}$$

is finite for all $s \in \mathbb{N}^d$ and $t \in [0, T]$. Consequently, the expectation in Eq. (3.2) corresponds to a finite sum and is finite.

An in-depth analysis is required to establish sharp conditions under which (3.2) holds. Our numerical analysis show that the MP propensity is well defined in many examples.

The mimicking process $\{\bar{S}(t)\}_{t \in [0, T]}$ is a SRN and consequently Markovian, whereas the projected full-dimensional process $\{S(t)\}_{t \in [0, T]}$ is non-Markovian. Moreover, the propensities of the full-dimensional process $\{a_j\}_{j=1}^J$ are time-homogeneous functions of the state, whereas the resulting propensities, $\{\bar{a}_j\}_{j=1}^J$ of the MP-SRN \bar{S} are time-dependent (see (3.4)). Reactions with $P(v_j) = 0$ do not contribute to the MP propensity in (3.4). For reactions with $P(v_j) \neq 0$, it may occur that their corresponding projected propensity is known analytically. We denote the index set of reactions requiring an estimation of (3.4) (e.g., via a L^2 regression as described in Section 3.2) by \mathcal{J}_{MP} . This index set is described as follows:

$$\mathcal{J}_{MP} := \left\{ 1 \leq j \leq J : P(v_j) \neq 0 \text{ and } \underbrace{a_j(\mathbf{x}) \neq f(P(\mathbf{x}))}_{(*)} \text{ for all functions } f : \mathbb{R}^d \rightarrow \mathbb{R} \right\}, \tag{3.6}$$

where condition (*) excludes reaction channels for which the MP propensity is only dependent on s and given in closed form by $\bar{a}_j(t, s) = f(s)$ for the function f .

3.2. Discrete L^2 regression for approximating projected propensities

To approximate the Markovian propensity \bar{a}_j for $j \in \mathcal{J}_{MP}$, we reformulate (3.4) as a minimization problem and then use discrete L^2 regression as described below.

We let $V := \{f : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R} : \int_0^T \mathbb{E}[f(t, P(\mathbf{X}(t)))]^2 dt < \infty\}$. Then, the projected propensities via the MP for $j \in \mathcal{J}_{MP}$ are approximated by

$$\begin{aligned} \bar{a}_j(\cdot, \cdot) &= \operatorname{argmin}_{h \in V} \int_0^T \mathbb{E} \left[(a_j(\mathbf{X}(t)) - h(t, P(\mathbf{X}(t))))^2 \right] dt \\ &\approx \operatorname{argmin}_{h \in V} \mathbb{E} \left[\frac{1}{N} \sum_{n=0}^{N-1} \left(a_j(\hat{\mathbf{X}}_n^{\Delta t}) - h(t_n, P(\hat{\mathbf{X}}_n^{\Delta t})) \right)^2 \right] \\ &\approx \operatorname{argmin}_{h \in V} \frac{1}{M} \sum_{m=1}^M \frac{1}{N} \sum_{n=0}^{N-1} \left(a_j(\hat{\mathbf{X}}_{[m],n}^{\Delta t}) - h(t_n, P(\hat{\mathbf{X}}_{[m],n}^{\Delta t})) \right)^2, \end{aligned} \tag{3.7}$$

where $\{\hat{\mathbf{X}}_{[m]}^{\Delta t}\}_{m=1}^M$ are M independent TL paths with a uniform time grid $0 = t_0 < t_1 < \dots < t_N = T$ with step size Δt .

To solve (3.7), we use a discrete L^2 regression approach. For the case $\bar{d} = 1$, we employ a set of basis functions in V , $\{\phi_p(\cdot, \cdot)\}_{p \in \Lambda}$, where $\Lambda \subset \mathbb{N}^2$ is a finite index set. In Remark 3.3, we provide more details on the choice of the basis. Consequently, the projected propensities via MP are approximated by

$$\bar{a}_j(t, s) \approx \sum_{p \in \Lambda} c_p^{(j)} \phi_p(t, s), j \in \mathcal{J}_{MP} \tag{3.8}$$

where the coefficients $c_p^{(j)}$ must be derived for $j \in \mathcal{J}_{MP}$ and $p \in \Lambda$.

Next, we derive the linear systems of equations, solved by $\{c_p^{(j)}\}_{p \in \Lambda}$ from (3.8) for $j \in \mathcal{J}_{MP}$. For a given one-dimensional indexing of $\{1, \dots, M\} \times \{0, \dots, N - 1\}$, the corresponding design matrix $\mathbf{D} \in \mathbb{R}^{MN \times |\Lambda|}$ is given by

$$D_{k,p} = \phi_p(t_n, P(\hat{\mathbf{X}}_{[m],n}^{\Delta t})), \text{ for } k = (m, n) \in \{1, \dots, M\} \times \{0, \dots, N - 1\}, \quad p \in \Lambda.$$

Further, we set $\psi_k^{(j)} = a_j(\hat{\mathbf{X}}_{[m],n}^{\Delta t})$ ($\psi^{(j)} \in \mathbb{R}^{MN}$) for $k \in \{1, \dots, M\} \times \{0, \dots, N - 1\}$, and $j \in \mathcal{J}_{MP}$.

Then, the minimization problem in (3.7) becomes

$$\begin{aligned} c^{(j)} &= \operatorname{argmin}_{\{c_p\}_{p \in \Lambda}} \frac{1}{MN} \sum_{m=1}^M \sum_{n=0}^{N-1} \left(a_j(\hat{\mathbf{X}}_{[m],n}^{\Delta t}) - \sum_{p \in \Lambda} c_p \phi_p(t_n, P(\hat{\mathbf{X}}_{[m],n}^{\Delta t})) \right)^2 \\ &= \operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^{|\Lambda|}} (\boldsymbol{\psi}^{(j)} - \mathbf{D}\mathbf{c})^\top (\boldsymbol{\psi}^{(j)} - \mathbf{D}\mathbf{c}) \\ &= \operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^{|\Lambda|}} \underbrace{\boldsymbol{\psi}^{(j)\top} \boldsymbol{\psi}^{(j)} - 2\mathbf{c}^\top \mathbf{D}^\top \boldsymbol{\psi}^{(j)} + \mathbf{c}^\top \mathbf{D}^\top \mathbf{D}\mathbf{c}}_{=: I(\mathbf{c})}. \end{aligned}$$

We minimize $I(\mathbf{c})$ with respect to \mathbf{c} by solving

$$\frac{\partial I(\mathbf{c})}{\partial \mathbf{c}} = -2\mathbf{D}^\top \boldsymbol{\psi}^{(j)} + 2\mathbf{D}^\top \mathbf{D} \mathbf{c} = 0$$

and obtain the normal equation for $j \in \mathcal{J}_{MP}$:

$$(\mathbf{D}^\top \mathbf{D}) \mathbf{c}^{(j)} = \mathbf{D}^\top \boldsymbol{\psi}^{(j)}. \tag{3.9}$$

Remark 3.3 (Orthonormal Basis Approach Via Empirical Inner Product). For the case $\bar{d} = 1$, the normal equation with a set of polynomials $\{\phi_p\}_{p \in \Lambda}$ in \mathbb{R}^2 can be used to derive the MP propensity \bar{a}_j for $j \in \mathcal{J}_{MP}$. We use the standard basis $\{\phi_{(i_1, i_2)}\}_{(i_1, i_2) \in \Lambda}$ for a two-dimensional index set Λ , where

$$\phi_{(i_1, i_2)} : \mathbb{R}^2 \rightarrow \mathbb{R}, (t, x) \mapsto t^{i_1} x^{i_2}.$$

For better stability [43], we use the Gram–Schmidt orthogonalization algorithm to determine an orthonormal set of functions for the empirical scalar product:

$$\langle \phi_i, \phi_j \rangle_M = \frac{1}{N} \sum_{n=0}^{N-1} \frac{1}{M} \sum_{m=1}^M \phi_i(t_n, P\hat{\mathbf{X}}_{[m],n}^{\Delta t}) \phi_j(t_n, P\hat{\mathbf{X}}_{[m],n}^{\Delta t}) \tag{3.10}$$

to find an orthonormal set of functions. We base the empirical scalar product and the normal Eq. (3.9) on the same set of TL paths, $\{\hat{\mathbf{X}}_{[m]}^{\Delta t}\}_{m=1, \dots, M}$, such that the matrix condition number becomes $\text{cond}(\mathbf{D}^\top \mathbf{D}) = 1$ and $\mathbf{D}^\top \mathbf{D} = \text{diag}\left(\frac{T}{\Delta t} M, \dots, \frac{T}{\Delta t} M\right)$ [43].

Remark 3.4 (Propensity Functions). The proof of Theorem 3.1 and the derivation of the normal Eq. (3.9) is general with respect to the structure of the propensity function. The Markovian projection can be applied to SRNs with arbitrary propensity functions, i.e., not restricted to the mass-action-kinetics principle introduced in (1.4) (e.g. Hill-type reaction rate law [44]). The choice of suitable basis functions $\{\phi_p\}_{p \in \Lambda}$ for the L^2 regression depends on the given example and also on the type of propensity.

3.3. Computational cost of Markovian projection

The computational work to derive an MP for an SRN with J reactions based on a time stepping Δt based on M TL paths and an orthonormal set of polynomials (see Remark 3.3) of size $\#\Lambda$ splits into three types of costs:

$$W_{MP}(\#\Lambda, \Delta t, M) \approx M \cdot W_{TL}(\Delta t) + W_{G-S}(\#\Lambda, \Delta t, M) + W_{L^2}(\#\Lambda, \Delta t, M), \tag{3.11}$$

where W_{TL} , W_{G-S} , and W_{L^2} denote the computational costs to simulate a TL path, derive an orthonormal basis (as described in Remark 3.3), and derive and solve the normal equation in (3.9), respectively. The dominant terms of these costs contribute as follows:

$$W_{TL}(\Delta t) \approx \frac{T}{\Delta t} \cdot J \cdot C_{Poi}, \quad W_{G-S}(\#\Lambda, \Delta t, M) \approx M \cdot \frac{T}{\Delta t} \cdot (\#\Lambda)^3, \quad W_{L^2}(\#\Lambda, \Delta t, M) \approx M \cdot \frac{T}{\Delta t} \cdot ((\#\Lambda)^2 + \#\mathcal{J}_{MP} \cdot \#\Lambda),$$

where C_{Poi} represents the cost to simulate one realization of a Poisson random variable. The main computational cost results from deriving an orthonormal basis (see Remark 3.3). A more detailed derivation of the cost terms is provided in Appendix C. For many applications, such as the MP-IS approach presented in Section 4, the MP must be computed only once, such that the computational cost $W_{MP}(\#\Lambda, \Delta t, M)$ can be regarded as an off-line cost.

Remark 3.5 (Simulating MP Paths). The MP-SRN $\bar{\mathbf{S}}(t)$ can be simulated as a SRN with inhomogeneous propensity function. The explicit TL scheme in (1.5) can be naturally adapted to this setting. However, we note that the IS approach presented in Section 4 does not require explicitly simulating paths of the MP-SRN.

4. Importance sampling for higher-dimensional stochastic reaction networks via Markovian projection

Next, we employ MP to overcome the curse of dimensionality when deriving IS controls from solving (2.9). Specifically, we solve the HJB equations in (2.11) for a reduced-dimensional MP system (refer to Fig. 1.2 for a schematic illustration of the approach). Given a suitable projection $P : \mathbb{R}^d \rightarrow \mathbb{R}^{\bar{d}}$ and a corresponding final condition $\bar{g} : \mathbb{N}^{\bar{d}} \rightarrow \mathbb{R}$ with $\bar{g}(P(\mathbf{x})) = g(\mathbf{x})$, the HJB Eqs. (2.11) for the MP process are

$$\begin{aligned} \bar{u}_{\bar{d}}(T, s) &= \bar{g}^2(s), \quad s \in \mathbb{N}^{\bar{d}} \\ \frac{d\bar{u}_{\bar{d}}}{dt}(t, s) &= -2 \sum_{j=1}^J \bar{a}_j(t, s) \left(\sqrt{\bar{u}_{\bar{d}}(t, s) \bar{u}_{\bar{d}}(t, \max(0, s + \bar{\mathbf{v}}_j))} - \bar{u}_{\bar{d}}(t, s) \right), \quad t \in [0, T], s \in \mathbb{N}^{\bar{d}}. \end{aligned} \tag{4.1}$$

For observables of the type $g(\mathbf{x}) = \mathbb{1}_{\{x_i > \gamma\}}$, we use an MP to a $(\bar{d} = 1)$ -dimensional process via projection (3.1), and the final condition is approximated by a positive sigmoid (see (2.12)). The solution of (4.1) is the value function $\bar{u}_{\bar{d}}$ of the \bar{d} -dimensional




Fig. 4.1. Schematic diagram MP-IS-MC. The costs of the operations in the first line (blue boxes) are off-line.

MP process. To obtain continuous-time IS controls for the d -dimensional SRN, we substitute the value function $\bar{u}(t, \mathbf{x})$ of the full-dimensional process in (2.10) with the value function $\bar{u}_d(t, P(\mathbf{x}))$ of the MP-SRN:

$$\bar{\delta}_j(t, \mathbf{x}) = a_j(\mathbf{x}) \sqrt{\frac{\bar{u}_d(t, \max(0, P(\mathbf{x} + \mathbf{v}_j)))}{\bar{u}_d(t, P(\mathbf{x}))}} \text{ for } \mathbf{x} \in \mathbb{N}^d, t \in [0, T]. \tag{4.2}$$

Remark 4.1 (Alternative MP-IS Approach). In the presented approach, we map the value function of the \bar{d} -dimensional MP process to the full-dimensional SRNs. Alternatively, one could also map the optimal controls from the \bar{d} -dimensional MP-SRN to the full-dimensional SRNs, leading to the following controls:

$$\bar{\delta}_j^d(t, \mathbf{x}) = \bar{a}_j(t, P(\mathbf{x})) \sqrt{\frac{\bar{u}_d(t, \max(0, P(\mathbf{x} + \mathbf{v}_j)))}{\bar{u}_d(t, P(\mathbf{x}))}}, \text{ for } \mathbf{x} \in \mathbb{N}^d, t \in [0, T]. \tag{4.3}$$

The numerical experiments demonstrate that this approach results in a comparable variance reduction to the approach presented in (4.2).

Remark 4.2 (Adaptive MP for $\bar{d} > 1$). In (4.2), when utilizing \bar{u}_d as the value function for the d -dimensional control, we introduce a bias to the optimal IS controls by approximating $\bar{u}(t, \mathbf{x})$ by $\bar{u}_d(t, P(\mathbf{x}))$ for $\mathbf{x} \in \mathbb{N}^d$ and $t \in [0, T]$. For the case $\bar{d} = d$, we have $\bar{u}_d(t, P(\mathbf{x})) = \bar{u}(t, \mathbf{x})$ and the MP produces the optimal IS control for the full-dimensional SRNs. For $\bar{d} < d$, this equality does not hold, since the interaction (correlation effects) between non-projected species are not taken into account in the MP SRNs, because the MP only ensures that the marginal distributions of $P(\mathbf{X}(t))|_{\{\mathbf{X}(0)=\mathbf{x}_0\}}$ and $\bar{S}(t)|_{\{\mathbf{X}(0)=\mathbf{x}_0\}}$ are identical. This can be seen in examples in which reactions occur with $P(\mathbf{v}_j) = 0$. Those reactions, are not present in the MP and; thus, are not included in the IS scheme. For the extreme case, $\bar{d} = 1$, we expect to achieve the least variance reduction which could be already substantial and satisfactory for many examples as we show in our numerical experiments.

However, examples could exist where a projection to dimension $\bar{d} = 1$ is insufficient to achieve a desired variance reduction. In this case, we can adaptively choose a better projection with increased dimension $\bar{d} = 1, 2, \dots$ until a sufficient variance reduction is achieved. This will imply an increased computational cost in the MP and in solving the HJB Eqs. (4.1) for $\mathbf{x} \in \mathbb{N}^d$. Investigating the effect of \bar{d} on improving the variance reduction of our approach is left for a future work.

To derive an MP-IS-MC estimator for a given uniform time grid $0 = t_0 \leq t_1 \leq \dots \leq t_N = T$ with step size Δt , we generate IS paths using the scheme in (2.1) with IS control parameters $\delta_{n,j}^{\Delta t}(\mathbf{x}) = \delta_j(t_n, \mathbf{x})$, as in (4.2), for $j = 1, \dots, J, \mathbf{x} \in \mathbb{R}^d, n = 0, \dots, N - 1$. Fig. 4.1 presents a schematic illustration of the entire derivation of the MP-IS-MC estimator.

This computational work consists of three cost contributions:

$$W_{MP-IS-MC}(\#A, \Delta t, M, M_{fw}) \approx W_{MP}(\#A, \Delta t, M) + W_{HJB}(\#A) + W_{forward}(\Delta t, M_{fw}),$$

where $W_{MP}(\#A, \Delta t, M)$ denotes the off-line cost to derive the MP (see (3.11)), $W_{HJB}(\#A)$ represents the cost to solve the HJB (4.1) for the \bar{d} -dimensional MP-SRN, and $W_{forward}(\Delta t, M_{fw})$ indicates the cost of deriving M_{fw} IS paths. The cost to solve the HJB (4.1) $W_{HJB}(\#A)$ depends on the used solver, and the cost for the forward run has the following dominant terms:

$$W_{forward}(\Delta t, M_{fw}) \approx M_{fw} \cdot \frac{T}{\Delta t} \cdot (J \cdot C_{Poi} + C_{lik} + \#J_{MP} \cdot C_{\delta}), \tag{4.4}$$


where C_{δ} is the cost to evaluate (4.2).

Remark 4.3 (Further Applications of MP in the Context of SRNs). In this work, we use the described MP for dimension reduction to derive a sub-optimal change of measure for IS, but the same MP framework can be used for other applications, such as solving the chemical master equation [39] or the Kolmogorov backward equations [40]. We intend to explore these directions in a future work.

5. Numerical experiments and results

Through **Examples 5.1** and **5.2**, we demonstrate the advantages of the proposed MP-IS approach compared with the standard MC approach for rare event estimations. We numerically demonstrate that the proposed approach achieves a substantial variance reduction compared with standard MC estimators when applied to SRNs with various dimensions.

Example 5.1 (Michaelis–Menten Enzyme Kinetics [45]). The Michaelis–Menten enzyme kinetics are enzyme-catalyzed reactions describing the interaction of an enzyme E with a substrate S , resulting in a product P :



where $\theta = (0.001, 0.005, 0.01)^\top$. We consider the initial state $\mathbf{X}_0 = (E(0), S(0), C(0), P(0))^\top = (100, 100, 0, 0)^\top$ and the final time $T = 1$. The corresponding propensity and the stoichiometric matrix are given by


$$a(\mathbf{x}) = \begin{pmatrix} \theta_1 E S \\ \theta_2 C \\ \theta_3 C \end{pmatrix}, \quad \mathbf{v} = \begin{pmatrix} -1 & 1 & 1 \\ -1 & 1 & 0 \\ 1 & -1 & -1 \\ 0 & 0 & 1 \end{pmatrix}.$$

The observable of interest is $g(\mathbf{x}) = \mathbf{1}_{\{x_3 > 22\}}$.

Example 5.2 (Goutsias’s Model of Regulated Transcription [46,47]). The model describes a transcription regulation through the following six molecules:

Protein monomer	(M) ,	Transcription factor	(D) ,
mRNA	(RNA) ,	Unbound DNA	(DNA) ,
DNA bound at one site	$(DNA \cdot D)$,	DNA bounded at two sites	$(DNA \cdot 2D)$.

These species interact through the following 10 reaction channels



where $(\theta_1, \dots, \theta_{10}) = (0.043, 0.0007, 0.0715, 0.0039, 0.0199, 0.479, 0.000199, 8.77 \times 10^{-12}, 0.083, 0.5)$. As the initial state, we use $\mathbf{X}_0 = (M(0), D(0), RNA(0), DNA(0), DNA \cdot D(0), DNA \cdot 2D(0)) = (2, 6, 0, 0, 2, 0)$, and the final time is $T = 1$. We aim to estimate the rare event probability $\mathbb{P}(D(T) > 8)$.

5.1. Markovian projection results

Through simulations for **Examples 5.1** and **5.2**, we numerically demonstrate that the distribution of the MP process $\bar{S}(T) |_{\{\mathbf{X}_0 = \mathbf{x}_0\}}$ matches the conditional distribution of the projected process $S(t) |_{\{\mathbf{X}_0 = \mathbf{x}_0\}} = P(\mathbf{X}(t)) |_{\{\mathbf{X}_0 = \mathbf{x}_0\}}$, as shown in **Theorem 3.1**. For both examples, we use an MP projection with $\bar{d} = 1$ using the projection given in (3.1), where the projected species is indexed as $i = 3$ in **Example 5.1** and as $i = 2$ in **Example 5.2**.

The MP is based on $M = 10^4$ TL sample paths with a step size of $\Delta t = 2^{-8}$ and uses the orthonormal basis of polynomials described in **Remark 3.3** with $\Lambda = \{0, 1, 2\} \times \{0, 1, 2\}$ for the L^2 regression. For the numerical comparison of the full-dimensional SRN with the MP-SRN, we use explicit TL paths with step size $\Delta t = 2^{-8}$ (see **Remark 3.5**). In **Fig. 5.1**, we show sample paths of the original SRN and the mimicking MP-SRN. **Fig. 5.2** shows the relative occurrences of states at final time T with $M_{fw} = 10^4$ sample paths, comparing the TL estimate of $P(\mathbf{X}(t)) |_{\{\mathbf{X}_0 = \mathbf{x}_0\}}$ and the TL-MP estimate of $\bar{S}(T) |_{\{\mathbf{X}_0 = \mathbf{x}_0\}}$. In both examples, the one-dimensional MP process mimics the distribution of the state of interest $X_i(T)$ of the original SRNs. Further quantification and analyses of the MP error are left for future work. In this work, a detailed analysis of the MP error is less relevant because the MP is used as a tool to derive IS controls for the full-dimensional process, and the IS is bias-free with respect to the TL scheme.

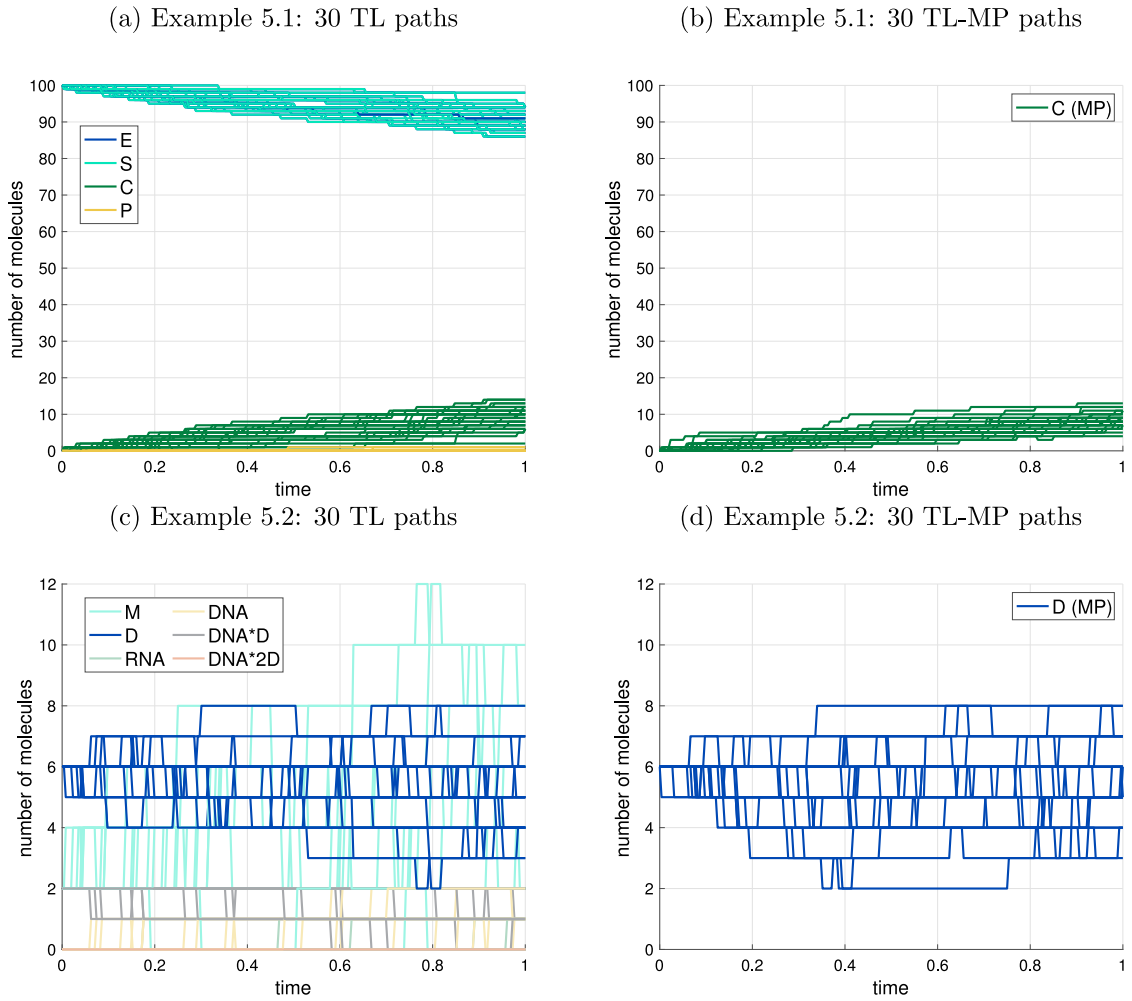


Fig. 5.1. Sample paths of TL and MP-TL for a step size of $\Delta t = 2^{-8}$.




Fig. 5.2. Relative occurrences of states at final time T in $M_{test} = 10^4$ sample paths, comparing the TL estimate of $P(\mathbf{X}(t))|_{\mathbf{X}_0=\mathbf{x}_0}$ and the MP estimate of $\bar{S}(T)|_{\mathbf{X}_0=\mathbf{x}_0}$. We set the step size to $\Delta t = 2^{-8}$ for the sample paths. The MP is based on $M = 10^4$ TL paths with a step size of $\Delta t = 2^{-8}$ and uses the orthogonal basis of polynomials described in Remark 3.3, where $\Lambda = \{0, 1, 2\} \times \{0, 1, 2\}$.

5.2. Markovian projection-importance sampling results

For the numerical experiments, we use a four-dimensional and a six-dimensional SRNs with the observable $g(\mathbf{x}) = \mathbb{1}_{\{x_i > \gamma\}}$, where i and γ are specified in Examples 5.1 and 5.2. Fig. 5.2 indicates that this observable leads to a rare event probability estimation for which an MC estimate is insufficient. We use the workflow in Fig. 4.1 with separate simulations for various Δt values for the MP-IS simulations. The MP is based on $M = 10^4$ TL sample paths each. The MP-IS-MC estimator, the sample variance, and the kurtosis estimate are based on $M_{fw} = 10^6$ IS sample paths.

The relative error is more relevant than the absolute error for rare event probabilities. Therefore, we display the squared coefficient of variation [1,48] in the simulations results, which is given by the following for a random variable X :

$$Var_{rel}[X] = \frac{Var[X]}{\mathbb{E}[X]^2}. \quad (5.1)$$

The kurtosis is a good indicator of the robustness of the variance estimator (see [1,22] for the connection between the sample variance and kurtosis).

Fig. 5.3 shows the simulation results for the four-dimensional Example 5.1 for different step sizes Δt . The quantity of interest is a rare event probability with a magnitude of 10^{-5} . For a step size of $\Delta t = 2^{-10}$, the proposed MP-IS approach reduces the squared coefficient of variation by a factor of 10^6 compared to the standard MC-TL approach. The third plot in Fig. 5.3 indicates that the kurtosis of the proposed MP-IS approach is below the kurtosis for standard TL for all observed step sizes Δt , confirming that the proposed approach results in a robust variance estimator. In Fig. 5.3(d), we show the required number of sample paths to reach a prescribed relative tolerance (see (1.10) and (1.11)). We compare the required number of sample paths for standard TL with the sample paths in the forward run of the IS-TL estimator, M_{fw} , and the total number of paths $M + M_{fw}$ including $M = 10^4$ TL paths to derive the MP (see Section 3.3). We observe that for small tolerances, the additional paths to derive the MP become negligible compared to the cost of the forward IS-TL run. Further, we reduce the number of paths significantly compared to the standard TL approach. For small tolerances, the reduction of sample paths is up to a factor of approximately 10^6 compared to the standard TL approach.

The second application of the proposed IS approach is the six-dimensional Example 5.2. Fig. 5.4 shows that this rare event probability has a magnitude of 10^{-3} . We observe that, for $\Delta t \leq 2^{-3}$, the squared coefficient of variation of the proposed MP-IS approach is reduced compared to the standard TL-MC approach. For a step size of $\Delta t = 2^{-10}$, this is a variance reduction of a factor of approximately 500. Note that, this example achieves less variance reduction than Example 5.1 due to a less rare quantity of interest. For most step sizes Δt , the kurtosis of the proposed IS approaches is moderately increased compared to the standard TL estimator, with decreasing kurtosis for smaller Δt . This outcome indicates a potentially insatiable variance estimator for coarse time steps of $\Delta t > 2^{-7}$. For finer time steps, we expect a robust variance estimator. For small tolerances the total number of paths is reduced by a factor of approximately 500.

6. Conclusion

In conclusion, this work presented an efficient IS scheme for estimating rare event probabilities for SRNs. We utilized a class of parameterized IS measure changes originally introduced in [1], for which near-optimal IS controls can be derived through a SOC formulation. We showed that the value function associated with this formulation can be expressed as a solution of a set of coupled ODEs, the HJB equations. One challenge encountered in solving the HJB equations is the curse of dimensionality, arising from the high-dimensional SRN. To address this issue, we introduced a dimension reduction approach for the setting of SRNs, namely MP. Then, we used a discrete L^2 regression to approximate the propensity and the stoichiometric vector of the MP-SRN. We demonstrated how the MP-SRN can be used for solving a significantly lower-dimensional HJB system, and how the resulting parameters are then mapped back to the full-dimensional SRNs to derive near-optimal IS controls. Our numerical simulations showed substantial variance reduction for the MP-IS-MC estimator compared to the standard MC-TL estimator for rare event probability estimations.

Data availability

The authors are unable or have chosen not to specify which data has been used.

Acknowledgments

This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-2019-CRG8-4033. This work was performed as part of the Helmholtz School for Data Science in Life, Earth and Energy (HDS-LEE) and received funding from the Helmholtz Association of German Research Centres and the Alexander von Humboldt Foundation. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

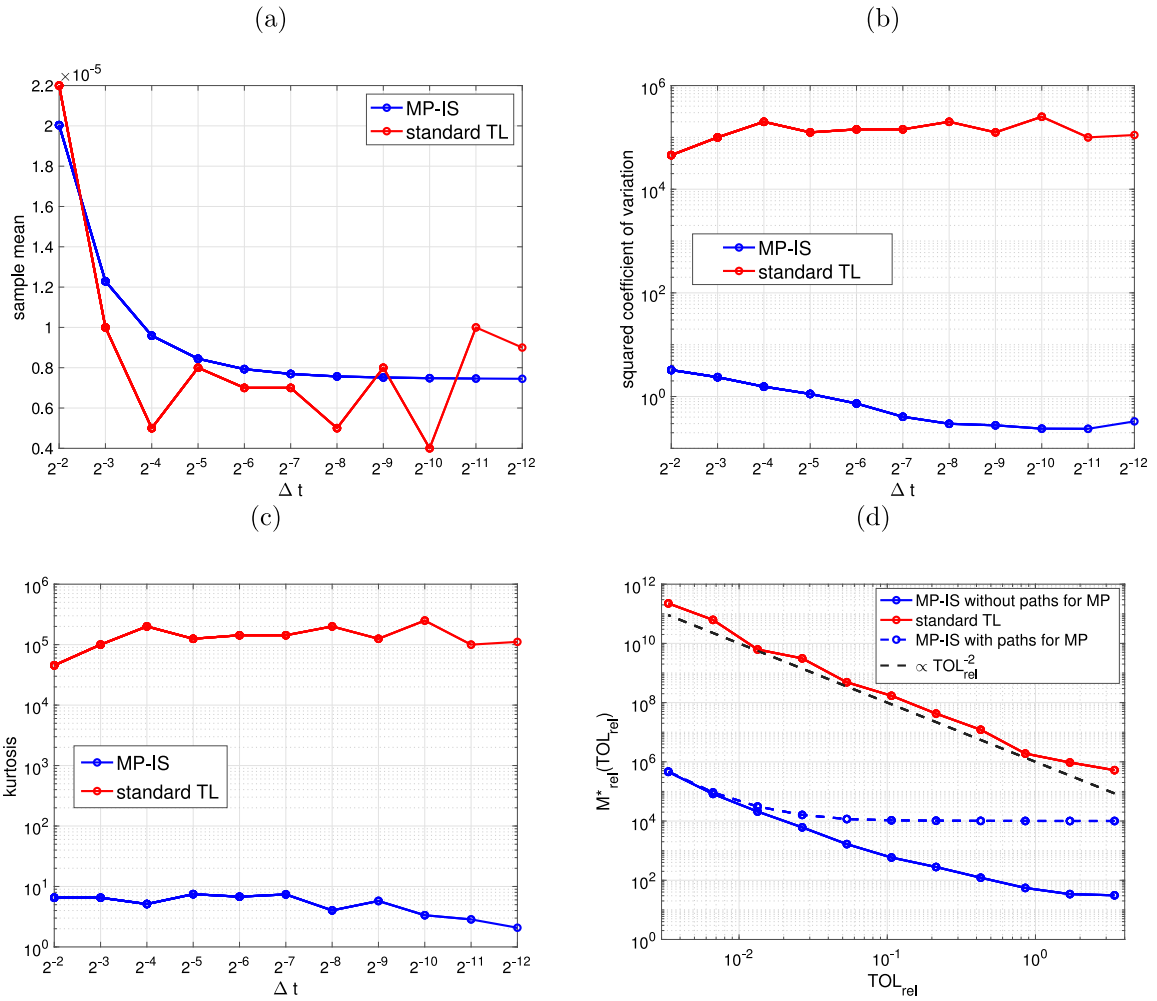


Fig. 5.3. Example 5.1 with separate simulations for different step sizes Δt for the proposed IS methods and the standard TL-MC estimator: (a) sample mean, (b) squared coefficient of variation, (c) kurtosis, and (d) number of required sample paths to reach a prescribed relative tolerance TOL_{rel} . The used MP is based on $M = 10^4$ TL sample paths each, and the estimators in (a), (b) and (c) are derived based on $M_{fw} = 10^6$ sample paths. For (d), we use (1.10) and (1.11) and the dashed line represents the total number of sample paths including the derivation of the MP (i.e. $M + M_{rel}^*(TOL_{rel})$).

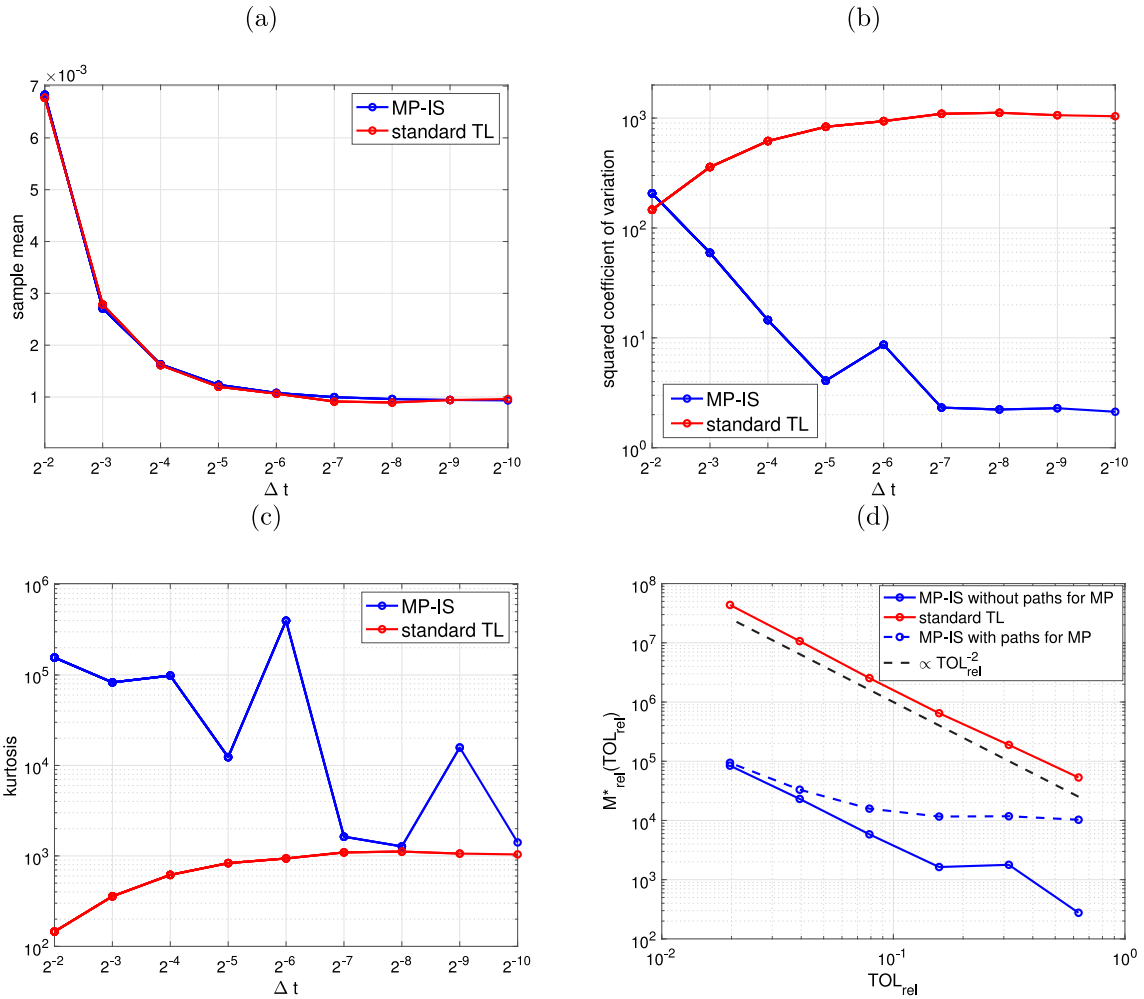


Fig. 5.4. Example 5.2 with separate simulations for different step sizes Δt for the proposed IS methods and the standard TL-MC estimator: (a) sample mean, (b) squared coefficient of variation, (c) kurtosis, and (d) number of required sample paths to reach a prescribed relative tolerance TOL_{rel} . The used MP is based on $M = 10^4$ TL sample paths each, and the estimators in (a), (b) and (c) are derived based on $M_{fw} = 10^6$ sample paths.

Appendix A. Proof for Corollary 2.3

Proof. For $\mathbf{x} \in \mathbb{N}^d$, we define $\tilde{u}(\cdot, \mathbf{x}; \Delta t)$ as the continuous smooth extension of $u_{\Delta t}(\cdot, \mathbf{x})$ (defined in (2.8)) on $[0, T]$. Consequently, we denote the continuous-time IS controls by $\delta(\cdot, \mathbf{x}) : [0, T] \rightarrow \mathcal{A}_{\mathbf{x}}$ for $\mathbf{x} \in \mathbb{N}^d$. Then, the Taylor expansion of $\tilde{u}(t + \Delta t, \mathbf{x}; \Delta t)$ in t results in the following:

$$\tilde{u}(t + \Delta t, \mathbf{x}; \Delta t) = \tilde{u}(t, \mathbf{x}; \Delta t) + \Delta t \partial_t \tilde{u}(t, \mathbf{x}; \Delta t) + \mathcal{O}(\Delta t^2), \mathbf{x} \in \mathbb{N}^d. \tag{A.1}$$

By the definition of the value function 2.1, the final condition is given by

$$\tilde{u}(T, \mathbf{x}; \Delta t) = g^2(\mathbf{x}), \mathbf{x} \in \mathbb{N}^d.$$

For $t = T - \Delta t, \dots, 0$, we apply to (2.8) from Theorem 2.2 a Taylor expansion around $\Delta t = 0$ to the exponential term and (A.1) to $\tilde{u}(t + \Delta t, \mathbf{x}; \Delta t)$:

$$\begin{aligned} \tilde{u}(t, \mathbf{x}; \Delta t) &= \inf_{\delta(t, \mathbf{x}) \in \mathcal{A}_{\mathbf{x}}} \exp \left(\left(-2 \sum_{j=1}^J a_j(\mathbf{x}) + \sum_{j=1}^J \delta_j(t, \mathbf{x}) \right) \Delta t \right) \times \sum_{\mathbf{p} \in \mathbb{N}^J} \left(\prod_{j=1}^J \frac{(\Delta t \cdot \delta_j(t, \mathbf{x}))^{p_j}}{p_j!} \left(\frac{a_j(\mathbf{x})}{\delta_j(t, \mathbf{x})} \right)^{2p_j} \right) \cdot \tilde{u}(t + \Delta t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}\mathbf{p}); \Delta t) \\ &= \inf_{\delta(t, \mathbf{x}) \in \mathcal{A}_{\mathbf{x}}} \left(1 + \left(-2 \sum_{j=1}^J a_j(\mathbf{x}) + \sum_{j=1}^J \delta_j(t, \mathbf{x}) \right) \Delta t + \mathcal{O}(\Delta t^2) \right) \times \left[\sum_{\mathbf{p} \in \mathbb{N}^J} \left(\prod_{j=1}^J \frac{(\Delta t \cdot \delta_j(t, \mathbf{x}))^{p_j}}{p_j!} \left(\frac{a_j(\mathbf{x})}{\delta_j(t, \mathbf{x})} \right)^{2p_j} \right) \right. \\ &\quad \left. \cdot \left(\tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}\mathbf{p}); \Delta t) + \Delta t \partial_t \tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}\mathbf{p}); \Delta t) + \mathcal{O}(\Delta t^2) \right) \right] \\ &\stackrel{(*)}{\implies} -\partial_t \tilde{u}(t, \mathbf{x}; \Delta t) = \inf_{\delta(t, \mathbf{x}) \in \mathcal{A}_{\mathbf{x}}} \left(-2 \sum_{j=1}^J a_j(\mathbf{x}) + \sum_{j=1}^J \delta_j(t, \mathbf{x}) \right) \tilde{u}(t, \mathbf{x}; \Delta t) + \mathcal{O}(\Delta t) + (1 + \mathcal{O}(\Delta t)) \left[\sum_{\mathbf{p} \neq \mathbf{0}} \Delta t^{|\mathbf{p}|-1} \left(\prod_{j=1}^J \frac{a_j(\mathbf{x})^{2p_j}}{p_j! \cdot (\delta_j(t, \mathbf{x}))^{p_j}} \right) \right. \\ &\quad \left. \cdot \left(\tilde{u}(t, \max(\mathbf{0}, \mathbf{x} + \mathbf{v}\mathbf{p}); \Delta t) + \mathcal{O}(\Delta t) \right) \right], \end{aligned} \tag{A.2}$$

where $|\mathbf{p}| := \sum_{i=1}^J p_i$, and $\delta_j(t, \mathbf{x}) := (\delta(t, \mathbf{x}))_j$. In (*), we split the sum, rearrange the terms, divide by Δt , and collect the terms of $\mathcal{O}(\Delta t)$.

The limit for $\Delta t \rightarrow 0$ in (A.2) is denoted by $\tilde{u}(t, \mathbf{x})$, leading to (2.9) for $0 < t < T$ and $\mathbf{x} \in \mathbb{N}^d$. \square

Appendix B. Proof for Theorem 3.1

Proof. We let $f : \mathbb{R}^{\bar{d}} \rightarrow \mathbb{R}$ be an arbitrary bounded continuous function and $\bar{\mathcal{S}}$ be defined in (3.3). We consider the following weak approximation error:

$$\varepsilon_T := \mathbb{E} [f(P(\mathbf{X}(T))) | \mathbf{X}(0) = \mathbf{x}_0] - \mathbb{E} [f(\bar{\mathcal{S}}(T)) | \bar{\mathcal{S}}(0) = P(\mathbf{x}_0)]. \tag{B.1}$$

For $t \in [0, T]$, we define the cost to go function as

$$\bar{v}(t, s) := \mathbb{E} [f(\bar{\mathcal{S}}(T)) | \bar{\mathcal{S}}(t) = s].$$

Then, we can represent the weak error in (B.1) as follows:

$$\varepsilon_T = \mathbb{E} [\bar{v}(T, P(\mathbf{X}(T))) | \mathbf{X}(0) = \mathbf{x}_0] - \bar{v}(0, P(\mathbf{x}_0)). \tag{B.2}$$

Using Dynkin's formula [49] and (1.3), we can express the first term in (B.1) as follows:

$$\begin{aligned} &\mathbb{E} [\bar{v}(T, P(\mathbf{X}(T))) | \mathbf{X}(0) = \mathbf{x}_0] \\ &= \bar{v}(0, P(\mathbf{x}_0)) + \int_0^T \mathbb{E} \left[\partial_t \bar{v}(\tau, P(\mathbf{X}(\tau))) + \sum_{j=1}^J a_j(\mathbf{X}(\tau)) (\bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) - \bar{v}(\tau, P(\mathbf{X}(\tau)))) | \mathbf{X}(0) = \mathbf{x}_0 \right] d\tau. \end{aligned}$$

The Kolmogorov backward equations [49] of $\bar{\mathcal{S}}$ are given as

$$\partial_\tau \bar{v}(\tau, s) = - \sum_{j=1}^J \bar{a}_j(\tau, s) (\bar{v}(\tau, s + \bar{\mathbf{v}}_j) - \bar{v}(\tau, s)), s \in \mathbb{N}^{\bar{d}},$$

implying that the weak error simplifies to

$$\varepsilon_T = \sum_{j=1}^J \int_0^T \mathbb{E} [a_j(\mathbf{X}(\tau)) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau))) \bar{v}(\tau, P(\mathbf{X}(\tau)) + \bar{\mathbf{v}}_j) | \mathbf{X}(0) = \mathbf{x}_0] d\tau$$

$$- \mathbb{E} \left[(a_j(\mathbf{X}(\tau)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau)))) \bar{v}(\tau, P(\mathbf{X}(\tau))) \mid \mathbf{X}(0) = \mathbf{x}_0 \right] d\tau. \tag{B.3}$$

Next, we choose \bar{a}_j and \bar{v}_j for $j = 1, \dots, J$ such that $\epsilon_T = 0$ for any function f . We consider the second term in (B.3) and use the tower property to obtain

$$\begin{aligned} & \mathbb{E} \left[(a_j(\mathbf{X}(\tau)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau)))) \bar{v}(\tau, P(\mathbf{X}(\tau))) \mid \mathbf{X}(0) = \mathbf{x}_0 \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[(a_j(\mathbf{X}(\tau)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau)))) \bar{v}(\tau, P(\mathbf{X}(\tau))) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0 \right] \mid \mathbf{X}(0) = \mathbf{x}_0 \right] \\ &= \mathbb{E} \left[(\mathbb{E} [a_j(\mathbf{X}(\tau)) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0] - \bar{a}_j(\tau, P(\mathbf{X}(\tau)))) \bar{v}(\tau, P(\mathbf{X}(\tau))) \mid \mathbf{X}(0) = \mathbf{x}_0 \right]. \end{aligned} \tag{B.4}$$

To ensure that (B.4) = 0 for any function f , we obtain the following:

$$\bar{a}_j(\tau, P(\mathbf{X}(\tau))) = \mathbb{E} [a_j(\mathbf{X}(\tau)) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0], j = 1, \dots, J. \tag{B.5}$$

Applying (B.5) and the tower property for the first term, we derive

$$\begin{aligned} & \mathbb{E} [a_j(\mathbf{X}(\tau)) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau))) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) \mid \mathbf{X}(0) = \mathbf{x}_0] \\ &= \mathbb{E} \left[\mathbb{E} [a_j(\mathbf{X}(\tau)) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) - \bar{a}_j(\tau, P(\mathbf{X}(\tau))) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0] \mid \mathbf{X}(0) = \mathbf{x}_0 \right] \\ &= \mathbb{E} \left[\mathbb{E} [a_j(\mathbf{X}(\tau)) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0] \bar{v}(\tau, P(\mathbf{X}(\tau) + P(\mathbf{v}_j))) - \bar{a}_j(\tau, P(\mathbf{X}(\tau))) \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j)) \mid \mathbf{X}(0) = \mathbf{x}_0 \right] \\ &= \mathbb{E} \left[\mathbb{E} [a_j(\mathbf{X}(\tau)) \mid P(\mathbf{X}(\tau)), \mathbf{X}(0) = \mathbf{x}_0] \cdot (\bar{v}(\tau, P(\mathbf{X}(\tau) + P(\mathbf{v}_j))) - \bar{v}(\tau, P(\mathbf{X}(\tau) + \mathbf{v}_j))) \mid \mathbf{X}(0) = \mathbf{x}_0 \right]. \end{aligned}$$

Moreover, Eq. (B.6) becomes zero for any function f using

$$\bar{v}_j = P(\mathbf{v}_j), j = 1, \dots, J. \tag{B.6}$$

With this choice for \bar{a}_j and \bar{v}_j , we derive $\epsilon_T = 0$. The derivation holds for arbitrary bounded and smooth functions f , for all fixed times T ; thus, the process $\mathcal{S}(t) = P(\mathbf{X}(t))$ has the same conditional distribution as $\bar{\mathcal{S}}(t)$ conditioned on the initial value $\mathbf{X}(0) = \mathbf{x}_0$. \square

Appendix C. Markovian projection cost derivation

We present details on the computational cost of MP, as provided in (3.11):

- The number of operations to generate one TL paths is given by

$$W_{TL}(\Delta t) = \frac{T}{\Delta t} \cdot (C_{prop} + J \cdot C_{poi} + d(J + 2)), \tag{C.1}$$

where C_{prop} is the cost of one evaluation of the propensity function (1.4). The dominant cost in (C.1) is C_{poi} (the cost of generating a Poisson random variable).

- The number of operations for the Gram–Schmidt algorithm, as described in Remark 3.3, is given by

$$W_{G-S}(\#\Lambda, \Delta t, M) = \#\Lambda \cdot (C_{inner} + \#\Lambda + 1) + \frac{(\#\Lambda - 1)\#\Lambda}{2} (2\#\Lambda + C_{inner}), \tag{C.2}$$

where C_{inner} is the cost of the evaluation of the empirical inner product (3.10) given by

$$C_{inner} = \frac{T}{\Delta t} \cdot M(2 + 2C_{pol}) + 3 = \mathcal{O}\left(\frac{T}{\Delta t} \cdot M \cdot \#\Lambda\right). \tag{C.3}$$

The cost C_{pol} in (C.3) is the computational cost for one evaluation of a polynomial in the space $\langle \phi_p \rangle_{p \in \Lambda}$, which is $\mathcal{O}(\#\Lambda)$.

In the simulations, we apply the setting $\#\Lambda \ll \frac{T}{\Delta t} \cdot M$ (see Section 5.2, using the parameter $\#\Lambda = 9, M = 10^4, \frac{T}{\Delta t} = 2^4$). Therefore, the dominant cost in (C.2) is $\mathcal{O}(M \cdot \frac{T}{\Delta t} \cdot (\#\Lambda)^3)$.

- The cost $W_{L^2}(\#\Lambda, \Delta t, M)$ is split into two: the cost to (i) derive and (ii) solve the normal Eq. (3.9). The number of operations to derive the design matrix D is $M \cdot \frac{T}{\Delta t} \cdot \#\Lambda \cdot C_{pol}$, and the cost to derive one right-hand side $(\Psi^{(j)})_{j \in \mathcal{J}_{MP}}$ is $M \cdot \frac{T}{\Delta t} \cdot C_{prop}$. In (3.9), the cost for the matrix product $D^T D$ is $\mathcal{O}(\#\Lambda^2 \cdot M \cdot \frac{T}{\Delta t})$, and the cost for $\#\mathcal{J}_{MP}$ matrix–vector products is $\mathcal{O}(\#\mathcal{J}_{MP} \cdot \#\Lambda \cdot M \cdot \frac{T}{\Delta t})$. Finally, solving (3.9) costs $\mathcal{O}(\#\mathcal{J}_{MP} \cdot \#\Lambda^3)$, which is a nondominant term under the given setting, $\#\Lambda \ll \frac{T}{\Delta t} \cdot M$.

References

- [1] Chiheb Ben Hammouda, Nadhir Ben Rached, Raúl Tempone, Sophia Wiechert, Learning-based importance sampling via stochastic optimal control for stochastic reaction networks, Stat. Comput. 33 (3) (2023) 58.
- [2] Chiheb Ben Hammouda, Hierarchical Approximation Methods for Option Pricing and Stochastic Reaction Networks (Ph.D. thesis), 2020.
- [3] Fred Brauer, Carlos Castillo-Chavez, Mathematical Models in Population Biology and Epidemiology, vol. 40, Springer.
- [4] David F. Anderson, Thomas G. Kurtz, Stochastic Analysis of Biochemical Systems, vol. 1, Springer, 2015.
- [5] Ranjan Srivastava, Lingchong You, Jesse Summers, John Yin, Stochastic vs. deterministic modeling of intracellular viral kinetics, J. Theoret. Biol. 218 (3) (2002) 309–321.
- [6] Sebastian C. Hensel, James B. Rawlings, John Yin, Stochastic kinetic modeling of vesicular stomatitis virus intracellular growth, Bull. Math. Biol. 71 (7) (2009) 1671–1692.
- [7] Christian Bayer, Alvaro Moraes, Raúl Tempone, Pedro Vilanova, An efficient forward–reverse expectation-maximization algorithm for statistical inference in stochastic reaction networks, Stoch. Anal. Appl. 34 (2) (2016) 193–231.

- [8] Daniel T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, *J. Comput. Phys.* 22 (4) (1976) 403–434.
- [9] David F. Anderson, A modified next reaction method for simulating chemical systems with time dependent propensities and delays, *J. Chem. Phys.* 127 (21) (2007) 214107.
- [10] Daniel T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems, *J. Chem. Phys.* 115 (4) (2001) 1716–1733.
- [11] Juan P. Aparicio, Hernán G. Solari, Population dynamics: Poisson approximation and its relation to the Langevin process, *Phys. Rev. Lett.* 86 (18) (2001) 4183.
- [12] Yang Cao, Linda Petzold, Trapezoidal tau-leaping formula for the stochastic simulation of biochemical systems, in: *Proceedings of Foundations of Systems Biology in Engineering, FOSBE 2005*, 2005, pp. 149–152.
- [13] Muruhan Rathinam, Hana El Samad, Reversible-equivalent-monomolecular tau: A leaping method for “small number and stiff” stochastic chemical systems, *J. Comput. Phys.* 224 (2) (2007) 897–923.
- [14] Assyr Abdulle, Yucheng Hu, Tiejun Li, Chebyshev methods with discrete noise: the τ -ROCK methods, *J. Comput. Math.* (2010) 195–217.
- [15] Tae-Hyuk Ahn, Adrian Sandu, Xiaoying Han, Implicit simulation methods for stochastic chemical kinetics, 2013, arXiv preprint arXiv:1303.3614.
- [16] Alvaro Moraes, Raúl Tempone, Pedro Vilanova, A multilevel adaptive reaction-splitting simulation method for stochastic reaction networks, *SIAM J. Sci. Comput.* 38 (4) (2016) A2091–A2117.
- [17] Chiheb Ben Hammouda, Alvaro Moraes, Raúl Tempone, Multilevel hybrid split-step implicit tau-leap, *Numer. Algorithms* 74 (2) (2017) 527–560.
- [18] Michael B. Giles, Multilevel Monte Carlo path simulation, *Oper. Res.* 56 (3) (2008) 607–617.
- [19] Michael B. Giles, Multilevel Monte Carlo methods, *Acta Numer.* 24 (2015) 259–328.
- [20] David F. Anderson, Desmond J. Higham, Multilevel Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics, *Multiscale Model. Simul.* 10 (1) (2012) 146–179.
- [21] Christopher Lester, Christian Adam Yates, Michael B Giles, Ruth E Baker, An adaptive multi-level simulation algorithm for stochastic biological systems, *J. Chem. Phys.* 142 (2) (2015) 01B612_1.
- [22] Chiheb Ben Hammouda, Nadhir Ben Rached, Raúl Tempone, Importance sampling for a robust and efficient multilevel Monte Carlo estimator for stochastic reaction networks, *Stat. Comput.* 30 (2020) 1665–1689.
- [23] Hiroyuki Kuwahara, Ivan Mura, An efficient and exact stochastic simulation method to analyze rare events in biochemical systems, *J. Chem. Phys.* 129 (16) (2008) 10B619.
- [24] Dan T. Gillespie, Min Roh, Linda R. Petzold, Refining the weighted stochastic simulation algorithm, *J. Chem. Phys.* 130 (17) (2009) 174103.
- [25] Min K. Roh, Dan T. Gillespie, Linda R. Petzold, State-dependent biasing method for importance sampling in the weighted stochastic simulation algorithm, *J. Chem. Phys.* 133 (17) (2010) 174106.
- [26] Bernie J. Daigle Jr., Min K. Roh, Dan T. Gillespie, Linda R. Petzold, Automated estimation of rare event probabilities in biochemical systems, *J. Chem. Phys.* 134 (4) (2011) 01B628.
- [27] Youfang Cao, Jie Liang, Adaptively biased sequential importance sampling for rare events in reaction networks with comparison to exact solutions from finite buffer dCME method, *J. Chem. Phys.* 139 (2) (2013) 07B605_1.
- [28] Colin S. Gillespie, Andrew Golightly, Guided proposals for efficient weighted stochastic simulation, *J. Chem. Phys.* 150 (22) (2019) 224103.
- [29] Min K. Roh, Data-driven method for efficient characterization of rare event probabilities in biochemical systems, *Bull. Math. Biol.* 81 (8) (2019) 3097–3120.
- [30] István Gyöngy, Mimicking the one-dimensional marginal distributions of processes having an Itô differential, *Probab. Theory Related Fields* 71 (4) (1986) 501–516.
- [31] Nikolai Vladimirovich Krylov, *Nonlinear Elliptic and Parabolic Equations of the Second Order*, Springer, 1987.
- [32] Thomas Kurtz, Richard Stockbridge, Stationary solutions and forward equations for controlled and singular martingale problems, 2001.
- [33] Amel Bentata, Rama Cont, Mimicking the marginal distributions of a semimartingale, 2009, arXiv preprint arXiv:0910.3992.
- [34] Frédéric Legoll, Tony Lelièvre, Effective dynamics using conditional expectations, *Nonlinearity* 23 (9) (2010) 2131.
- [35] Vladimir Piterberg, Markovian projection method for volatility calibration, 2006, Available at SSRN 906473.
- [36] Boualem Djehiche, Björn Löfdahl, Risk aggregation and stochastic claims reserving in disability insurance, *Insurance Math. Econom.* 59 (2014) 100–108.
- [37] Christian Bayer, Juho Häppölä, Raúl Tempone, Implied stopping rules for American basket options from Markovian projection, *Quant. Finance* 19 (3) (2019) 371–390.
- [38] Carsten Hartmann, Christof Schütte, Wei Zhang, Model reduction algorithms for optimal control and importance sampling of diffusions, *Nonlinearity* 29 (8) (2016) 2298.
- [39] Linar Mikeev, Werner Sandmann, Approximate numerical integration of the chemical master equation for stochastic reaction networks, 2019, arXiv preprint arXiv:1907.10245.
- [40] Alvaro Moraes, *Simulation and Statistical Inference of Stochastic Reaction Networks with Applications to Epidemic Models* (Ph.D. thesis), 2015.
- [41] Stewart N. Ethier, Thomas G. Kurtz, *Markov Processes : Characterization and Convergence*, in: *Wiley Series in Probability and Mathematical Statistics*, J. Wiley & Sons, New York, Chichester, 1986.
- [42] Tiejun Li, Analysis of explicit tau-leaping schemes for simulating chemically reacting systems, *Multiscale Model. Simul.* 6 (2) (2007) 417–436.
- [43] Albert Cohen, Mark A. Davenport, Dany Leviatan, On the stability and accuracy of least squares approximations, *Found. Comput. Math.* 13 (2013) 819–834.
- [44] James N. Weiss, The hill equation revisited: uses and misuses, *FASEB J.* 11 (11) (1997) 835–841.
- [45] Christopher V. Rao, Adam P. Arkin, Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the gillespie algorithm, *J. Chem. Phys.* 118 (11) (2003) 4999–5010.
- [46] John Goutsias, Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems, *J. Chem. Phys.* 122 (18) (2005) 184102.
- [47] Hye-Won Kang, Thomas G. Kurtz, Separation of time-scales and model reduction for stochastic reaction networks, 2013.
- [48] Nadhir Ben Rached, Abdul-Lateef Haji-Ali, Gerardo Rubino, Raúl Tempone, Efficient importance sampling for large sums of independent and identically distributed random variables, *Stat. Comput.* 31 (6) (2021) 1–13.
- [49] Floyd B. Hanson, *Applied Stochastic Processes and Control for Jump-Diffusions: Modeling, Analysis and Computation*, SIAM, Philadelphia, PA, 2007.